ISSN 1471-0498

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

SEARCH ADVERTISING

Alexandre de Corniere

Number 649
March 2013

Manor Road Building, Manor Road, Oxford OX1 3UQ

Search Advertising

Alexandre de Cornière *

January 18, 2013

Abstract

Search engines enable advertisers to target consumers based on the query they have entered. In a framework with horizontal product differentiation, imperfect product information and in which consumers incur search costs, I study a game in which advertisers have to choose a price and a set of relevant keywords. The targeting mechanism brings about three kinds of efficiency gains, namely lower search costs, better matching, and more intense product market price-competition. A monopolistic search engine charges advertisers too high a price, and has incentives to provide a suboptimal matching quality. Competition among search engines eliminates the latter distortion, but exacerbates the former.

Keywords: search engine, targeted advertising, consumer search.

JEL Classification: D43, D83, L13, M37.

1 Introduction

Search engines are arguably the most important actors of the digital economy. More than four billion search queries are processed every day by search engines such as Google, Yahoo or Bing, to help users find all sorts of information. It is not a surprise that the development of these actors has generated interest from advertisers, to the point that search advertising is nowadays a multi-billion dollar industry.[1]

Search advertising designates the display of "sponsored links" on a search engine results page, alongside "organic links". Whereas organic links are free, sponsored links are the main source of

*Department of Economics and Nuffield College, University of Oxford. Email: adecorniere@gmail.com. I am grateful to Simon Anderson, Mark Armstrong, Heski Bar-Isaac, Luc Bridet, Bernard Caillaud, Olivier Compte, Jacques Crémer, Gianni De Fraja, Gabrielle Demange, Romain de Nijs, Bruno Jullien, Frederic Koessler, Dina Mayzlin, Régis Renault and Joana Resende for very useful suggestions. Previous versions of this paper have circulated under the title "Targeting with consumer search: an economic analysis of keyword advertising".

[1]See Evans (2008) for an interesting presentation of the online advertising industry, with a special emphasis on search engines

revenue for search engines. The standard pricing scheme in the industry is *per-click pricing*: search engines collect fees from advertisers every time a consumer clicks on their link.

Besides the size of search engines' user base, several factors explain the success of this advertising format. Unlike, say, TV advertising, search advertising reaches consumers at a point at which they are actively looking for information, and is therefore less of a nuisance. This is all the more so that, as advertisers can select precise keywords to target, sponsored links are generally relevant to the queries and thus valuable to consumers.

The questions that I address in the paper are the following: how does the mechanism composed of keyword targeting and per-click pricing affect the market outcomes (profits, welfare)? What are the strategic incentives of a search engine? Is competition between search engines desirable?

To answer these questions, I present (sections 2 and 3) a model of targeted advertising through a search engine, with differentiated products, which includes the main features mentioned above. Firms are horizontally differentiated *à la* Salop (1979), and consumers do not have prior knowledge of firms' prices or products' characteristics. The search engine is an intermediary between firms and consumers: firms choose which keywords they want to target, and consumers enter keywords and then search sequentially at random through the links that appear. Firms incur a fixed cost to be registered on the search engine, and they pay the search engine on a per-click basis.

The main findings are the following: in equilibrium, search expenses are minimized, since firms only target consumers who find it optimal not to search further. With respect to a benchmark in which firms cannot target consumers, I also find that the quality of the matching between firms and consumers is higher (i.e the average distance in the product space is smaller). Perhaps more surprisingly, another consequence of firms' ability to target consumers is an increase in the intensity of price competition. This result stems from the fact that targeting endogenously reduces the perceived cost of an additional search, because consumers know that with targeting they draw firms from a better pool. The intensification of price competition thus lowers firms' mark-up, which is the third way through which targeting may improve efficiency on the market. However, allowing firms to target their advertising leads them to regard the per-click fee as a marginal cost, and to pass it through in the price of their product. The optimal fee charged by the search engine is thus too high with respect to the social optimum, because it excludes some consumers from the market. On the other hand, without targeting, the per-click fee is analogous to a fixed cost, which has no bearing on the equilibrium price chosen by firms.

In practice, if search engines possess superior information, say, about the quality of the match between a firm and a keyword, they will most likely try to use it so as to optimally design the matching mechanism. For instance, Google sorts firms using a weighted average of the firms' bids and of a "quality score" index. Consumers are also sometimes provided with additional information on the results page, such as a map showing the locations of different vendors. On the other hand, the "broad match" technology enables search engines to expand the set of keywords corresponding to a given advertisement. In section 4, I study a situation in which the search engine can more finely design the matching mechanism, and can also influence consumer search costs. The analysis reveals that, even if the search engine could implement the perfect matching at no cost, it would not be optimal to do so. Indeed, implementing an accurate matching mechanism would lead to a hold-up situation (the Diamond paradox), in which firms charge such a high price that consumers prefer not to participate. It can even be optimal for the search engine to implement a matching that is less accurate than the laissez-faire outcome (in which accuracy is the result of equilibrium behavior by firms and consumers). The reason is that offering a noisy matching mechanism makes consumers more willing to accept high-prices on the product market, because it is now more costly to refuse an offer and to search again, as the next firm is less likely to be a good match. Since the search engine cannot charge consumers, it may then be optimal to use such a strategy. It is not *always* optimal, because it results in a decrease in the number of active consumers, and so the search engine trades off per-consumer profit and number of consumers. Another, related, result is that the relation between search costs and search engine's profit is non-monotonic. In fact, for low values of the search costs, firms are not able to set their price above marginal cost, and therefore the search engine has to charge a low per-click fee in order for them to recover their fixed costs. For high values of the search costs, this effect no longer plays, but an increase in the search cost leads to a decrease in consumer participation. Therefore the search engine would prefer consumers to have intermediate search costs if possible.

In section 5, I build upon my baseline model to incorporate the issue of competition between search engines. First, I show that the presence of another search engine can lead to a rise in the advertising price. Indeed, if firms cannot price-discriminate consumers based on the search engine they use, a firm's marginal cost is the average of the two search engines' fees. Therefore, when a search engine raises its price, the pass-through rate is lower than under monopoly, and so the elasticity of consumer demand is lower. Competition thus exacerbates the distortion that occurs with a monopoly. Second, and in contrast with the first point, when search engines can design the quality of their matching mechanism, competition between search engines mitigates the incentives of search engines to lower the accuracy of their matching. An economist's assesment of the overall effect of competition will therefore depend on which effect is more prevalent in practice.

Related literature

This paper develops a new framework to provide an economic analysis of search engine advertising. The key features of the model (targeted advertising, consumer search, two-sided market) each have been extensively studied in the economic literature, but the combination of the three generates new insights.

Targeted advertising has received increased attention in recent years. Esteban, Gil, and Hernandez (2001) show that in a monopoly framework, firms' ability to target consumers reduces both consumer welfare and total surplus. Roy (2000), Galeotti and Moraga-Gonzalez (2008) or Iyer, Soberman, and Villas-Boas (2005) show how targeted advertising may generate market segmentation in a duopoly, with homogenous and heterogenous products. In this paper I will focus on the interplay between targeting and search, and show that it tends to make the market more competitive. Other recent works on targeted advertising include Van Zandt (2004) who shows that targeted advertising can lead to information overload, Johnson (forthcoming), who examines ad avoidance behavior, or Bergemann and Bonatti (2011) and Athey and Gans (2010), who study competition between medias with different targeting technologies. An important paper on advertising in the presence of horizontal differentiation is Grossman and Shapiro (1984). The product space is similar to the one in my paper, and the difference lies in the fact that in their model advertising is perfectly informative and there is no search.

The seminal paper on consumer search is Diamond (1971). In a model with several firms producing an homogenous good, and in which consumers incur a positive cost to obtain price information, firms necessarily charge the monopoly price in equilibrium. The reason for that is that demand is inelastic with respect to price, because a rise in the price inferior to the search cost does not drive consumers away from a firm. With heterogenous consumers, demand becomes price elastic and the "Diamond paradox" disappears. Such heterogeneity can lie in the level of information of consumers (e.g. Varian (1980), Stahl (1989)) or in their tastes. In the present paper I use the latter source of heterogeneity, building on Wolinsky (1983) who models preferences using Salop (1979) circular city model. Wolinsky (1986) and Anderson and Renault (1999) also deal with heterogenous preferences, modeling match values as i.i.d shocks.[2]

Some models of consumer search are more directly relevant to the search engine industry. Athey and Ellison (2011) focus on the design of the auction to allocate advertisement slots, given that consumers search strategically through the slots. However their analysis does not include competition between firms on the product market. Armstrong, Vickers, and Zhou (2009) deal with price competition between firms, in a model in which one firm is made prominent, meaning that although consumers search strategically, they always visit the prominent firm first. Chen and He (2011) and Haan and Moraga-Gonzalez (2011) endogenize prominence by including an advertising stage prior to firms' pricing decision and consumer search. In Chen and He (2011) this advertising stage is an auction in which the more relevant firms submit higher bids, making it rational for consumers to sample them first. Haan and Moraga-Gonzalez (2011) assume that consumers are boundedly rational, in the sense that the probability that a consumer remembers a firm is proportional to that firm's advertising expenses. None of these papers study the strategic choice of keywords by advertisers, nor

[2]The two approaches would yield qualitativey similar results.

the role of the search engine.

Finally, my paper is related to the growing literature on two-sided markets, with the seminal papers of Armstrong (2006), Caillaud and Jullien (2003), or Rochet and Tirole (2006). My approach is different from these papers, in the sense that I do not use a reduced-form way of modeling interactions between agents on the platform, in order to account for some important details. Neither do I allow complete flexibility in terms of pricing, focusing instead on the design of the matching process as a way to increase the platform's profit. Other papers have a similar approach: Baye and Morgan (2001) model an intermediary who acts as an information gatekeeper on a homogenous product market, and look at the optimal two-sided pricing, taking into account subsequent price setting by firms and consumer search. Hagiu and Jullien (2011) focus on the design of a platform in terms of search diversion, and highlight several reasons why an intermediary does not want to provide the highest quality matching, even when the technology is costless. Eliaz and Spiegler (2011), in a related paper, also show that a search engine wants to implement a matching with a suboptimal quality. White (2009) and Taylor (Forthcoming) examine the trade-off faced by a search engine between providing quality organic results (which tend to attract users) and generating clicks on sponsored links (through which the search engine makes money).[3] Gomes (2011) characterizes the optimal mechanism to sell an advertising slot when consumers and advertisers are heterogenous. In an auction context, de Corniere and de Nijs (2011) study the incentives of a platform to disclose information about consumers to firms, and show that the decision has implications on the market equilibrium. Disclosing information generate targeted advertising, but can result in a rise in the equilibrium price that make the policy inefficient.

2 The model

2.1 Description of the market and of preferences

The framework is based on Wolinsky (1983). Consider a market where a continuum of mass 1 of firms produce a differentiated good at zero marginal cost. Each product may be described by a single keyword. Keywords are uniformly distributed on a circle, whose perimeter is normalized to one.[4] Thus a firm is characterized by the position of its product's keyword on the circle. The type of a firm, i.e the keyword that perfectly describes its product, will be denoted $\theta \in [0;1]$. θ is private information.[5]

There is a continuum of mass 1 of consumers. Consumers differ along two dimensions: (i) each

[3] I discuss how these papers relate to my model at the end of section 4.

[4] The assumptions of a continuum of firms and of a circular product space are made for analytical convenience. The main insights hold under an alternative specification with a finite number of firms and i.i.d. valuations for the products.

[5] In section 4 I assume that θ is also observed by the search engine.

consumer has a favorite product (or keyword), $\omega \in [0;1]$, uniformly distributed around the circle, and (ii) consumers differ with respect to their willingness to pay for their favorite product. This willingness to pay v is independent of ω, and across consumers. It is distributed on $[0, \overline{v}]$ according to a continuous and increasing cumulative distribution function F , with density f.[6]

Both ω and v are consumers' private information. Consumers have use for at most one unit, and the utility that a consumer located in ω gets from consuming product θ, with $d(\theta, \omega) = d$, is

$$u(v, d, p) = v - \phi(d) - p \tag{1}$$

where p is the price of the good and ϕ is increasing and twice continuously differentiable over $[0, 1/2]$. $\phi(d)$ is often referred to as a transportation cost in traditional models of spatial competition. Here, I will use the terminology "mismatch cost".

2.2 Advertising technology on the search engine

Consumers have imperfect information about firms' characteristics: they do not know firms' position on the circle (θ) nor their price, and thus have to search before buying.

A firm that wants to launch an online advertising campaign using the search engine incurs a fixed cost C. This cost corresponds to the marketing or monitoring expenses that accompany the advertising campaign, and is not a payment to the search engine.

The search engine plays the role of a matchmaker: on the one hand, firms select the set of keywords that they want to target. This set is assumed to be symmetric around θ and convex: $\mathcal{K}(\theta) = [\theta - D_\theta; \theta + D_\theta]$. On the other hand, consumers enter the keyword they are interested in $\mathcal{L}(\omega) = \{\omega\}$. If a certain keyword ω is entered by a consumer, the search engine randomly selects a firm θ such that $\omega \in \mathcal{K}(\theta)$.[7] The consumer incurs a search cost $s > 0$ and learns the price and position of this firm. s corresponds to the amount of time and effort that are necessary to examine a firm's offer. The firm θ pays a fee $a > 0$ to the search engine. At that point, the consumer has three options: (*i*) he can accept the offer and leave the market, (*ii*) he can refuse the offer and leave the market, (*iii*) he can hold the offer and continue searching. In that case, the search engine randomly selects another firm θ' such that $\omega \in \mathcal{K}(\theta')$, and the process starts over.

At any point, consumers can come back at no cost towards a firm they have previously visited (recall is costless). It is the case if for instance consumers open a new window every time they click on a link.

[6]Having consumers differ with respect to v allows to generate an elastic demand for the search engine.

[7]The random matching corresponds to the assumption that the search engine is non-strategic with respect to the matching mechanism.

Discussion The assumption that consumers do not observe anything before clicking on a link seems appropriate in many contexts. Indeed, firms can provide very little information with the text under their link on a search engine's page. Consumers have to click on the link to get more precise information. In this respect, advertising is not informative in the usual sense: it does not provide information in itself, but in equilibrium consumers correctly infer that a firm which targets them is not farther than a certain distance. The assumption is less relevant when consumers have a previous knowledge of the firms and/or products (if they bought in the past, of if they know the brand). I assume away these kinds of situations, which certainly deserve a proper analysis.

In the model, I also ignore the auction mechanism that is run in practice, and assume that all firms pay the same price and are visited with the same probability. The conceptual difference between an auction and this mechanism is the following: in an auction, firms reveal *how much* they are willing to pay to receive a click, whereas here they only reveal *whether* they are willing to pay a given price per click. Although there might be some subtle differences between the two mechanisms, I believe that most of the insights developped here would carry over to an auction set-up.

The model captures the complementarity between search and advertising that is inherent to this technology. Firms have to target a keyword for them to be visited by the consumers who enter that keyword, and consumers can infer something about the product offered by the firm and its price from knowing that the keyword they entered is targeted by the firm. Receiving an ad does not dispense consumers from searching, and neither does it provide "hard" information regarding the products' characteristics or their price. Rather, advertising acts more like a signal of relevance. This approach is very different from Robert and Stahl (1993), in which advertising and consumer search are substitute, in the sense that if a consumer has received an advertisement he does not need to search. The subtle interaction between search and advertising is reminiscent of, although different from, Anderson and Renault (2006). In that paper, a monopolist can give as much information as it wants in its advertising, but consumers still need to incur a search cost, which must then be interpreted as a cost to physically visit the store. Interestingly, they show that for some parameter range, it is optimal for the firm to reveal to consumers that the match value is above a given threshold, which is very close to what is achieved in the present paper.[8]

2.3 Strategies and equilibrium concept

Timing and strategies The timing of the game is the following:

1. **Search engine pricing:** The search engine chooses a per-click fee a, which is publicly observed by firms and consumers.

[8]For other papers exploring the links between search and informative advertising, see for instance Mayzlin and Shin (2011) and Bar-Isaac, Caruana, and Cunat (2010).

2. **Firms pricing and targeting:** Firms decide whether to register on the search engine. If they do so they incur the fixed cost C. A firm θ that decides to use the search engine chooses a price p_θ and an advertising strategy D_θ.

3. **Consumer search:** Consumers decide whether they want to use the search engine or not. If a consumer uses the search engine, he enters the keyword corresponding to his favorite product (ω), and starts a sequential search among firms such that $d(\theta, \omega) \leq D_\theta$. Firms are uniformly drawn from $\{\theta \quad s.t. \quad d(\theta, \omega) \leq D_\theta\}$.

A consumer faces two decisions: whether to participate, and, if so, how to search. Both decisions will involve cutoff rules. First, let $EU(v)$ be the expected utility of a consumer of type v if he uses the search engine. If he does not search, his utility is normalized to zero. Let $v^*(a)$ (sometimes noted v^*) be such that $EU_{SE}(v^*(a)) = 0$. Consumers with $v \geq v^*(a)$ use the search engine, while consumers with $v < v^*(a)$ do not.[9]

Second, once a consumer has decided to use the search engine, he faces a sequential search problem. We know, from Kohn and Shavell (1974), that the optimal strategy is a stationary decision rule as long as there is at least one firm that has not been sampled. If, at any point, the best available offer comes from a firm located at a distance $\hat{d}$ from ω, with a price of $\hat{p}$, the consumer continues to search if and only if $v - \phi(\hat{d}) - \hat{p} < U_R$. The strategy of a consumer thus consists in the choice of the reservation utility U_R, or, alternatively, in the choice of a reservation distance $R(\hat{p}) \equiv \phi^{-1}(v - \hat{p} - U_R)$. Notice that R will depend on the expected future prices and locations if the consumer keeps on searching.

The equilibrium concept used is perfect Bayesian equilibrium: the search engine optimally chooses its fee a. Given a per-click fee a, advertisers set their participation decision, their price and their advertising policies so as to maximize their profit given the other firms' strategies and the stopping rule used by consumers.

The stopping rule R^* is a best-response to firms' strategies. I will focus on symmetric equilibria in pure strategies $(a^*, R^*, v^*, p^*, D^*)$. To highlight the fact that $R^*(.)$ depends on the expectation about future prices and locations, I will use the notation $R^*(p, p^*, D^*)$ where (p^*, D^*) refer to what consumers expect other firms to play.

Consumers have passive beliefs in the following sense: if a firm deviates from the equilibrium strategy (p^*, D^*), and this deviation is observed by a consumer, this consumer does not update his beliefs regarding other firms' strategy.

[9] In practice, consumers most likely do not observe the per-click fee paid by advertisers. My interpretation of this assumption is, in a broad sense, that a higher per-click fee will eventually drive consumers away from the search engine, because, say, they experience that prices online are too high compared to their search costs.If a was not observed but consumers could form rational expectations, the market would unravel.

3 Equilibrium analysis

Solving the game can be done in three steps. First, given equilibrium behavior by firms, and given the per click fee a, one can determine consumers' optimal stopping rule. Next, given this rule, we can find firms' equilibrium strategy in terms of pricing and advertising. Finally, given the equilibrium of the subgame, we can find the search engine's optimal per click fee a.

3.1 Consumer search

In equilibrium, when a consumer of type (v, ω) clicks on a link, the expected utility he gets from this click if he buys is

$$\int_{\omega-D^*}^{\omega+D^*} \frac{u(v, d(\omega, \theta), p^*)}{2D^*} d\theta = \int_0^{D^*} \frac{u(v, x, p^*)}{D^*} dx$$

Consumers regard each click as a random draw of a location θ from a uniform distribution, whose support is $[\omega - D^*; \omega + D^*]$. Indeed a firm located at a distance greater than D^* from ω would not appear on the results' page in equilibrium (the consumer would not be targeted). Suppose for now that all firms set the equilibrium price p^*. Then, after the first visit, the only way a consumer can improve his utility is by finding a firm that is a better match, i.e that is closer to him. For $R^* \equiv R(p^*, p^*, D^*)$ to be a reservation distance it must be such that a consumer is indifferent between continuing to search and buying the product:

$$\int_0^{R^*} \frac{u(v, x, p^*) - u(v, R^*, p^*)}{D^*} dx = s \tag{2}$$

The left-hand side of this equality is the expected improvement if a consumer decides to keep on searching after being offered a product at a price p^* and at a distance R^*. This expected improvement equals the search cost, so that the consumer is indifferent between buying or searching again. By totally differentiating (2), one gets

$$\frac{dR^*}{ds} = -\frac{D^*}{R^* u_2(v, R^*, p^*)} > 0, \quad \frac{dR^*}{dD^*} = -\frac{s^*}{R^* u_2(v, R^*, p^*)} > 0 \tag{3}$$

where u_2 is the partial derivative of u with respect to the second argument. R^* is an increasing function of the equilibrium reach of advertising D^*: if consumers expect firms to try to reach a wide audience (by targeting many keywords), they adjust their stopping rule by being less demanding, because the expected improvement after a given offer is lower than with more precise targeting. R^* is also an increasing function of search costs: consumers are less demanding if it costs more to continue searching. Note also that R^* does not depend on the equilibrium price p^*, because in equilibrium the expected price improvement due to an extra sample is always zero with quasi-linear utility functions. Indeed we have the following result:

Lemma 1 *For every D, p and p', we have $R(p,p,D) = R(p',p',D)$ when the utility is given by (1).*

Proof: From (2), $R(p,p,D)$ is given by $\int_0^{R(p,p,D)} \frac{\phi(R(p,p,D))-\phi(x)}{D} dx = s$. We see that it does not depend on p.$\square$

Now, when a consumer samples a firm which has set an out-of-equilibrium price $p \neq p^*$, his belief about other firms' strategy and position does not change, and therefore his optimal stopping rule (and thus the firm's demand) $R(p,p^*,D^*)$ is such that accepting a price p at a distance $R(p,p^*,D^*)$ gives the same utility as accepting a price p^* at a distance R^*, i.e $v - \phi(R(p,p^*,D^*)) - p = v - \phi(R^*) - p^*$. Thus we have the following result:

Lemma 2 *Given other firms' expected strategy (p^*,D^*), a consumer accepts to buy a good at price p if and only if the selling firm is located at a distance less than $R(p,p^*,D^*)$, with $R(p,p^*,D^*)$ such that*

$$v - \phi(R(p,p^*,D^*)) - p = v - \phi(R^*) - p^*$$

where R^* is given by (2).

Moreover, by the implicit function theorem, R is continuously differentiable and

$$\frac{dR(p,p^*,D^*)}{dp} = -\frac{dR(p,p^*,D^*)}{dp^*} = -\frac{1}{\phi'(R(p,p^*,D^*))} < 0 \quad (4)$$

Thus we have the natural property that a firm's demand decreases with its own price and increases with the expected price of other firms.

3.2 Equilibrium

Advertising Now that we know consumers' search behavior, it is possible to characterize firms' optimal targeting strategy. It turns out that this optimal strategy is surprisingly simple: a firm should target a consumer if and only if the distance between the two is smaller than the reservation distance. Indeed, suppose that firm θ sets a price p. Since it only has to pay for consumers who actually visit its link, firm θ's optimal targeting strategy is to appear to every consumer ω such that the expected profit made by θ through a sale to ω conditionally on ω clicking on θ's link is positive, i.e

$$p.Pr(\omega \text{ buys } \theta\text{'s product}|\omega \text{ clicks on } \theta\text{'s link}) - a \geq 0 \quad (5)$$

where a is the per-click fee paid to the search engine. With a continuum of firms, consumers' stopping rule is stationary, and a consumer never comes back to a firm he previously visited. Thus the conditional probability is either 0 (when $d(\omega,\theta) > R(p,p^*,D^*)$) or 1 (when $d(\omega,\theta) \leq R(p,p^*,D^*)$). Thus we have the following result, the proof of which is in the appendix:

Lemma 3 *Any symmetric equilibrium must involve $D^* = R^*(p^*, p^*, D^*)$. Therefore, if an equilibrium exists, it must be the case that consumers do not search more than once.*

This result, which relies on the assumption that all consumers have the same search rule and that targeting can be arbitrarily accurate, is counterfactual in the sense that in practice some consumers search more than once. This apparent paradox is still useful in that it clearly illustrates that targeting through keywords is a powerful instrument to reduce some inefficiencies due to the presence of search costs. However, notice that the equilibrium outcome is not the perfect matching, which would mean that firms target only the consumers for whom the product they offer is the ideal one. There is still some noise in the matching, due to the existence of search costs, but the level of noise is endogenously determined so as to cancel consumers' incentives to visit more than one firm.[10]

Pricing Thanks to Lemma 3, it is straightforward to find the per-(search engine)-user profit function of a firm if the other firms and consumers play their respective equilibrium strategies.[11] Indeed, if that firm wants to set a price p different from the candidate equilibrium price p^*, it must also change the set of consumers that it targets. By the same argument as in Lemma 3, the optimal advertising strategy is to target consumers if and only if they are located at a distance smaller than the new reservation distance $R(p, p^*, D^*)$. Since every consumer within this reservation distance is targeted by a mass $2R(p^*, p^*, D^*)$ of firms,[12] the per-user demand for the firm's product is $\frac{R(p,p^*,D^*)}{R(p^*,p^*,D^*)}$. Conditional on visiting the firm, all consumers buy without searching further, and this implies that a is formally equivalent to the firm's marginal cost of production. Therefore, if all the other players (firms and consumers) follow the equilibrium strategy profile, a firm's per-user profit function is

$$\pi(p, p^*, a) = (p - a)\frac{R(p, p^*, D^*)}{R(p^*, p^*, D^*)} \tag{6}$$

The previous reasonning does not rely on p^* being an equilibrium price, and so the profit function is defined for any price p^* that is played by all the other firms. The only restriction is that the profit function is defined only for $D^* = R(p^*, p^*, D^*)$. But it should be clear that if firms expect all the firms to play a price p^*, it is indeed optimal to choose $D^* = R(p^*, p^*, D^*)$.

Given firms' profit function when their rivals play the equilibrium targeting strategy D^* and charge the same price p^*, standard arguments will ensure the existence of a price equilibrium. Notice first that there always exists a "trivial"' equilibrium, in which firms set $D^* = 0$ and $p^* = \overline{v}$, and in which consumers do not search at all. I shall assume that when there is another equilibrium in which

[10] The first part of the lemma also relies on there being a continuum of firms. This assumption is however not crucial to obtain the result that consumers only search once. A model with a finite number of firms would deliver the same result.

[11] Since consumers cannot observe prices prior to using the search engine, firms cannot affect the number of search engine users and we can focus on the per-user profit.

[12] $R(p^*, p^*, D^*)$ coming from his left and the same amount from his right.

trade takes place, agents coordinate on the latter. Two additional assumptions ensure existence (Assumption 1) and uniqueness (Assumption 2) of an equilibrium.

Assumption 1 *For any p,* $R(p,p,1/2) < 1/2$.

Under Assumption 1, if firms do not target specific keywords (i.e they target the whole circle), some consumers search more than once before buying. In particular, this assumption requires search costs not to be too large. It seems a rather weak assumption, for if it was not satisfied there would be little point in studying the implications of a targeting mechanism (since firms would target every keyword). A more restrictive assumption is the following:

Assumption 2 *For all* $d \in [0,1/2]$, $\phi'(d) + d\phi''(d) \geq 0$.

Assumption 2 guarantees that the function $D \mapsto R(p,p,D)$ is concave, which is a sufficient condition for the uniqueness of equilibrium. Assumption 2 is satisfied for any increasing and convex function ϕ, that is if consumers are risk averse with respect to the quality of the match. It does not rule out risk-loving behavior of consumers ($\phi'' < 0$), but restricts the extent of risk-loving. For instance, if $\phi(d) = 1 - e^{-\alpha d}$, assumption 2 implies $\alpha \in (0,2]$

Proposition 1 *Under Assumption 1, there exists a non trivial equilibrium of the subgame in which the search engine has chosen a. If Assumption 2 holds, there is a unique non-trivial equilibrium, given by:*

$$s = \int_0^{R^*} \frac{\phi(R^*) - \phi(x)}{D^*} dx \tag{7}$$

$$R^* = D^* \tag{8}$$

$$p^* - a = \phi'(R^*)R^* \tag{9}$$

Proof: The proof of the existence and uniqueness is provided in the appendix. Equation (7) is simply a rewriting of equation (2), while (8) comes directly from Lemma 3. Equation (9) obtains by taking the first-order condition at a symmetric equilibrium in the expression of profit (equation (6)). This FOC writes $(p-a)R_1 + R = 0$ which, after using (4), gives the solution.□

Equation (9) gives the mark-up in equilibrium. By Assumption 2 and by (3), one can see that the mark-up is an increasing function of the search costs. As s increases, the option to search further becomes less valuable for consumers, and firms can therefore charge a higher price. As s goes to zero, the mark-up vanishes.

One should note that the results would also hold if payments were made on a per-impression basis, i.e every time a consumer enters a keyword, instead of a per-click basis. Indeed, in that case the per-user profit function of a firm would be $\pi(p,p^*,a) = \left(p\frac{R(p,p^*,D^*)}{R(p^*,p^*,D^*)} - aR(p,p^*,D^*)\right)$, and one would just need to replace a by $aR(p^*,p^*,D^*)$ in the expression of the equilibrium price (9).[13]

[13]If firms paid a to the search engine on a per-sale basis, (7), (8) and (9) would still constitute an equilibrium, but

3.3 Platform pricing

I now turn to the pricing decision of the search engine. The search engine is constrained in its choice, since it only has one instrument, namely the per-click fee paid by firms. With this instrument the search engine must pursue several goals: (i) attract users, (ii) attract firms and (iii) extract revenue from the trade that it allows. While objectives (i) and (ii) give an incentive to lower the fee, the third one entails increasing it.

Recall that $v^*(a) = \underbrace{\phi'(R^*)R^* + a}_{p^*} + E[\phi(d)|d \leq R^*] + s$ is the lowest value of v such that a consumer participates. Since in equilibrium every consumer who uses the search engine clicks only once, the search engine's profit is

$$\Pi^{SE}(a) = a\,(1 - F(v^*(a)))$$

Let $\hat{a}$ be the fee that maximizes this quantity. $\hat{a}$ is given by

$$\hat{a} = \frac{1 - F(v^*(\hat{a}))}{v^{*\prime}(\hat{a})f(v^*(\hat{a}))} = \frac{1 - F(v^*(\hat{a}))}{f(v^*(\hat{a}))} \tag{10}$$

This formula is reminiscent of the Lerner formula for monopoly pricing. Indeed, ignoring firms' participation constraint, the search engine earns a for every user, and raising a lowers consumer participation because the fee is passed through to consumers. The problem of the search engine is thus similar to a classical monopoly pricing problem.

However, the search engine faces the additional constraint that firms must make a non-negative profit. A firm's profit, in equilibrium, writes as

$$\pi_\theta(a) = (p^* - a)\,(1 - F(v^*(a))) - C = \phi'(R^*)R^*\,(1 - F(v^*(a))) - C$$

$\pi_\theta(.)$ is a decreasing function of a. Let $a_\theta \equiv sup\{a \quad / \pi_\theta(a) \geq 0\}$

Thanks to the assumption that firms are homogenous with respect to their fixed costs, the solution is straightforward to obtain. To further facilitate the analysis, the following assumption is needed:

Assumption 3 *f is log-concave.*

It is well-known, see Caplin and Nalebuff (1991), that the log concavity of the distribution of willingness to pay implies the quasi-concavity of the profit function, here $\Pi^{SE}(a)$.

there would also be an equilibrium without targeting (such that $D^* = 1/2$). See Taylor (2011) for a comparison of these payment schemes. Other mechanisms (two-part tariffs, first- or second-degree price-discrimination) would be more profitable for the search engine than the flat per-click fee used in this model. However, such schemes would, in my opinion, add complexity without doing as much in the way of realism.

Proposition 2 *Under Assumption 3, the optimal fee for the search engine is*

$$a^* = min\{a_\theta, \hat{a}\} \tag{11}$$

Proof: If $a_\theta \geq \hat{a}$, then firms' participation constraint is slack, and the search engine can charge its first best fee $\hat{a}$. If, on the other hand, $a_\theta < \hat{a}$, the constraint is binding. By quasi-concavity of the profit function Π^{SE}, the profit is increasing for $a \leq \hat{a}$, and so the optimal fee is a_θ. □

In any case, one easily sees that the optimal fee for the search engine is greater than the socially optimal fee, which is zero. Indeed, looking at (7), (8) and (9), one sees that a will have no impact on the quality of the matching in equilibrium, while a high a implies a higher price paid by consumers.[14] However, the search engine would not make any profit if this was the case. In order to increase its profit, the search engine imposes a distortion, because the higher fee results in a higher equilibrium price.

3.4 The effects of targeting

One of the main motivations for this paper is to understand the implications of the targeting technology on the product market equilibrium. In order to properly evaluate these implications, one needs a benchmark in which targeting is not possible. This benchmark, provided by Wolinsky (1983) (see also Bakos (1997)), consists in simply assuming that firms cannot specify a set of keywords, so that advertising is non-targeted. I use the subscript NT to index the equilibrium values corresponding to this case, and T for the case of targeting corresponding to the previous analysis.

The main results are the following:

Proposition 3 *Compared to a situation with random advertising, targeting:*

1. *reduces the expected number of clicks;*
2. *reduces the mismatch frictions;*
3. *has an ambiguous effect on the price of the final good.*

Proof: In the case of no-targeting, the equilibrium reservation distance for consumers, R_{NT}, is given by (7) with $D_{NT} = 1/2$. The first point of the proposition is a direct corollary of Lemma 3 and of Assumption 1.

The second point is a consequence of the fact that, as D increases without targeting, so does the reservation distance (by (3)), and therefore the expected mismatch cost $E[\phi(d)|d \geq R^*]$ decreases.

[14] When $a = 0$ there is also an equilibrium without targeting, but setting a arbitrarily close to zero is enough to discipline firms into targeting D^*.

The third point of the proposition is more subtle. To understand it, notice that, without targeting, the (per search engine user) profit of a firm that charges p is

$$\frac{1}{2R_{NT}}\left(p \times 2R(p, p_{NT}, D_{NT}) - a_{NT}\right)$$

Using equation (4), which gives the firm's demand derivative, and the first order condition at a symmetric equilibrium, one gets

$$p_{NT} = \phi'(R_{NT})R_{NT} \tag{12}$$

Comparing this latter expression with (9), one gets

$$p_T - p_{NT} = \underbrace{a_T}_{\text{pass-through effect,}>0} + \underbrace{\left(R_T\phi'(R_T) - R_{NT}\phi'(R_{NT})\right)}_{\text{semi-elasticity effect,}<0}.$$

The pass-through effect follows from the remark that, unlike in (6), advertising expenses without targeting do not vary with the firm's own price. Indeed, the number of clicks is independent of this price because the firm cannot adjust its advertising strategy along with its price.[15] This leads firms to regard these expenses as fixed costs, with no impact on the optimal product price.

By Assumption 2, which is satisfied for most of the usual functional forms in the literature, the function $x \mapsto \phi'(x)x$ is non-decreasing, which explains why the second effect is negative. The intuition for this effect is that targeting increases the continuation value of search for consumers, which makes them more willing to refuse high prices by increasing the (semi-) elasticity of demand. This in turn puts more pressure on firms to reduce their mark-up. □

4 Platform design

The assumption that the search engine does not behave strategically with respect to information revelation leaves aside interesting theoretical as well as practical issues. There is evidence that search engines pay a lot of attention to the way advertisements are displayed. The ranking of advertisements through a "quality score" illustrates this concern, as well as the use of a "broad match" technology aimed at matching consumers to firms when the keywords do not correspond exactly but are "close" enough. Basically, with broad match, the search engine will display an advertisement even if the keyword has not been selected by the firm, provided it is regarded as relevant by the search engine. For instance, suppose that a firm only selects one keyword, namely "web hosting". If a consumer enters the keyword "web hosting company" or "webhost", then the firm's advertisement will appear on the consumer's screen. Google argues that one of the benefits brought by such a practice is that

[15] See Dellarocas (2012) for a discussion of this point.

it saves time for firms: they no longer have to spend time and resources figuring out what are the right keywords to use. The search engine will do that for them, using the available information on past queries and results in order to find relevant keywords.

Such practices may be regarded as an attempt to choose the accuracy of the matching system. For instance, putting large weights on the most relevant websites to a query improves the quality of the matching process, whereas applying a very loose "broad match" policy introduces some additional noise. Another example is the display of maps, indicating the physical location of firms. In this section I study two ways in which a monopolistic search engine can design its matching mechanism so as to increase its profit. First, I assume that the search engine can influence the relevance of ads by choosing the value of D. Next, I look at the case in which the search engine can choose the value of the search cost. As White (2009) suggests, such an assumption can be related for instance to the quality of the non-paid links that are provided. If the search engine provides high quality links, consumers can gather more information about the products, and be able to evaluate the different offers more easily. In both cases, the search engine faces a trade-off between giving firms enough market power and ensuring sufficient consumer participation. This leads to the optimal choice of intermediate values of D and s. Interestingly, in the case in which the search engine chooses D, the optimal value is *always* at least as large as the equilibrium value obtained in section 3.

4.1 Optimal accuracy of the matching mechanism

Assume that the search engine is able to choose an accuracy level D, as well as a per-click fee a

Lemma 4 *If the search engine has the possibility to choose the accuracy of the matching, in equilibrium it can entirely extract firms' profit.*

Proof: Let $v^*(D)$ be the consumer who is indifferent between using the searching and his outside value of zero. Let $R(p, p^*, D)$ be the reservation distance of a consumer who faces a price p if other firms set a price p^*, and if the search engine chooses a level of accuracy D. Then the firm's profit is

$$(1 - F(v^*(D))) \left(p\frac{R(p, p^*, D)}{R(p^*, p^*, D)} - a \max\{\frac{D}{R(p^*, p^*, D)}, 1\} \right)$$

Indeed, if $D \leq R(p^*, p^*, D)$ consumers will search only once, whereas otherwise they search on average $\frac{D}{R(p^*,p^*,D)}$ times. In equilibrium, by setting $a \max\{\frac{D}{R(p^*,p^*,D)}, 1\} = p^*$, the search engine extracts all the profit. □

As is the case when firms cannot target specific keywords, the search engine extracts the whole profit. It is now straightforward to see that the search engine will choose D so as to maximize firms' gross profit, because it cannot make consumers pay.

The following proposition gives the optimal matching accuracy for the search engine. Recall that D^* is the equilibrium distance in the game in which firms choose their targeting strategy. But first, I make an additional assumption.

Assumption 4 *ϕ is convex.*

Under this assumption, consumers are now risk-averse with respect to the sampling process.

Proposition 4 *The optimal matching accuracy, from the search engine's point of view, is $D^{SE} \geq D^*$. That is, the search engine will not improve the quality of the matching with respect to the "laissez-faire" situation.*

The complete proof of this proposition is in the appendix, but its logic is the following. When the search engine chooses a low value of D ($D < D^*$), the reservation distance $R(p, p, D)$ is strictly larger than D, so that sufficiently small price increases by firms do not result in a lower demand. But then firms increase their price at least up to the willingness to pay of the marginal user. Since the search cost is sunk, it does not enter into the willingness to pay, and the marginal user is thus left with a negative utility. This leads the market to collapse. This is a variant of the well-known Diamond paradox. Said differently, too much accuracy in the results protects firms from competition, but this extra market power dissuades consumers from participating.

On the other hand, when $D > D^*$, we have $R(p, p, D) < D$, so that some users find it optimal to visit several firms. This disciplines firms into charging "competitive" prices (as opposed to monopoly prices), and thus ensures user participation. Over the range $[D^*; 1/2]$, firms' per-user profit is increasing in D (by the same logic as the semi-elasticity effect discussed in section 3.4) but the number of users is decreasing. The search engine will choose the optimal D^{SE} so as to balance these two effects.

4.2 Optimal search costs

Suppose that, by a careful design of its website, the search engine is able to choose the value of consumers' search cost s. In order to make things as transparent as possible, assume also that there is no cost associated with the choice of s. Then we have the following proposition.

Proposition 5 *From the search engine point of view, there exists $\overline{s} \in (0, \overline{v})$ such that the search engine's profit is maximal.*

Proof: When $s = 0$, firms basically compete à la Bertrand, and make zero profit. Therefore firms do not wish to participate, and the search engine's profit is zero. If, on the other hand, $s = \overline{v}$, consumers do not want to participate. Any s such that firms and some consumers participate allows the search engine to make some positive profit, therefore there must exist an interior $\overline{s}$ such that the search engine's profit is maximal. □

The idea that a platform should pay attention to the competition that takes place between merchants is present in Armstrong (2006). In particular, that paper shows, in a very stylized way, that when the platform cannot charge one side of the market, it is better-off restricting the intensity of competition, by granting exclusivity to one merchant.[16] In the present model, I focus on instruments that differ from exclusivity contracts (although these could be incorporated without difficulties). Several recent contributions make related points. Hagiu and Jullien (2011) show that an intermediary has two important motives to divert consumer search: (i) it can generate visits that would not have occurred if consumers had been directed towards their favorite shop, (ii) diversion may increase the elasticity of demand that each shop faces, leading them to charge a lower price, which can result in more participation to the platform by consumers. The latter effect is present in my model, in a very stark way, and explains why the search engine never wants to implement $D < D^*$. Eliaz and Spiegler (2011), in a set-up closer to this one, underline the fact that a lower quality of the matching mechanism increases the effective search cost faced by consumers, so that firms can charge a higher price to consumers. Here, it is because of this effect that the search engine may want to implement a matching of even lower quality than D^*.

Another paper that looks at the design of platform in a related spirit is White (2009). In that paper, the search engine can choose the value of search costs, by choosing the quality of the organic results. The trade-off is then that some of the clicks on sponsored links will be diverted to organic links, which bring no profits to the search engine. Taylor (Forthcoming) shows that even with competition among search engines, there exist equilibria in which both search engines provide sub-optimal quality for their organic links, in order to generate more revenues from the sponsored links.

5 Competing search engines

Looking at the model with a monopolistic search engine allowed to exhibit some interesting aspects of the targeting mechanism, but one may ask how the introduction of competition in the model would affect the results.

Suppose that there are two search engines on the market. Search engines are *ex ante* identical, so that all consumers have the same ordering of preferences between the two. If consumers are indifferent between the search engines, search engine i receives a share n_i of traffic. I assume that firms incur no cost to use search engines, so that they multi-home.

In this section I study two stylized games of competition between search engines, that extend the analysis of the two previous sections. In the first one, targeting is determined by firms, and the

[16] See also Dukes and Gal-Or (2003) for a model of TV advertising and exclusivity.

search engines only choose per-click fees a_1 and a_2. In the second game, the search engines can each choose an accuracy of targeting (D_1 and D_2), just as in section 4.1. In both games, the timing is the following: (1) both search engines choose their strategy (a_i or a_i and D_i); (2) firms observe search engines' strategies and choose a price p that is uniform across search engines; (3) consumers observe search engines' strategies, choose one search engine (or none at all) and search sequentially.

An important assumption is that sellers cannot charge a different price to users across search engines. This assumption is consistent with casual empiricism, and its implications are discussed below.

Given that, in both set-ups, a monopolistic search engine imposes a distortion, a natural question to ask is whether competition between search engines improves the market outcomes. Perhaps surprizingly, the answer differs between the two cases. When the search engine cannot control the accuracy of targeting, competition leads to an increase in per-click fees that are then passed-through to consumers. However, when search engines control the accuracy of targeting, competition eliminates the distortion outlined in Proposition 4.

5.1 Firms choose targeting.

When search engines cannot control the accuracy of targeting, the assumption that firms charge a uniform price implies that consumers derive the same utility from trade on both search engines, and thus that market shares are n_1 and n_2.[17] Without loss of generality, let us assume that search engine 1 is "dominant", i.e that $n_1 \geq n_2$.

For each search engine user, a firm's expected profit, if it charges p when other firms charge p^* and play D^*, is

$$\pi(p, p^*, a_1, a_2) = (p - (n_1 a_1 + n_2 a_2)) \frac{R(p, p^*, D^*)}{R(p^*, p^*, D^*)} \tag{13}$$

As in Proposition 1, the symmetric price charged by firms is

$$p^*(a_1, a_2) = n_1 a_1 + n_2 a_2 + R^* \phi'(R^*) \tag{14}$$

A consumer will use a search engine if and only if $v \geq E[\phi(d)|d \leq R^*] + p^*(a_1, a_2) + s \equiv v^*(n_1 a_1 + n_2 a_2)$. Each search engine i maximizes $n_i a_i (1 - F(v^*))$. The first-order condition of the profit-maximization problem is thus

$$n_i a_i = \frac{1 - F(v^*)}{f(v^*)} \tag{15}$$

We therefore obtain the following results:

Proposition 6 *Under assumptions 1 to 3, there exists a unique equilibrium in the game. Moreover,*

[17] Indeed, the price of the good will be the same on both search engines, and so will be the reservation distances, which do not depend on a_i, per (7) and (8).

(i) the dominant search engine charges a lower per-click fee than its rival, and (ii) the expected per-click fee $n_1a_1^* + n_2a_2^*$ is higher than the monopoly per-click fee a^* (see 11).

Proof: First, from assumption 3, f is log-concave. Theorem 2 in Bagnoli and Bergstrom (2005) thus ensures that $1-F$ is also log-concave, which implies that $\frac{1-F}{f}$ is decreasing. To prove uniqueness, notice that (15) leads to $a_2 = \frac{n_1}{n_2}a_1$. Thus we can rewrite $v^* = E[\phi(d)|d \leq R^*]+2n_1a_1+R^*\phi'(R^*)+s$. It is clear that the solution to (15) for $i = 1$ exists and is unique, since it is the intersection between a decreasing curve ($\frac{1-F(v^*)}{f(v^*)}$ as a function of a_1) and an increasing line (n_1a_1) that covers the whole range of positive reals.

The fact that $a_1 \leq a_2$ directly comes from $n_1a_1 = n_2a_2$. Finally, in order to show that $n_1a_1^* + n_2a_2^* \geq a^*$, rewrite (15) as

$$\frac{n_1a_1^* + n_2a_2^*}{2} = \frac{1 - F(v^*(n_1a_1^* + n_2a_2^*)))}{f(v^*(n_1a_1^* + n_2a_2^*))}$$

and compare with equation (10):[18]

$$a^* = \frac{1 - F(v^*(a^*))}{f(v^*(a^*))}$$

It is straightforward to see that the solution to the first equation ($n_1a_1^* + n_2a_2^*$) must be larger than the solution to the second (a^*). □

Proposition 6 crucially relies on the assumption that firms cannot set a different price depending on the search engine. Indeed, such an assumption implies that there is no "real" price-competition between search engines, in the sense that a decrease in a_i does not steal business from search engine j. The relevant demand for search engine i is thus $n_i(1 - F(v^*(n_1a_1 + n_2a_2)))$, and its elasticity is lower than that of the monopoly demand $1 - F(v^*(a))$, because an increase in a_i is passed through to consumers at a rate $n_i < 1$.[19][20]

The analysis underlines the fact that the impossibility for search engines to charge users prevents them from actually competing with each other. Indeed, any cut on the per-click fee benefits users of both search engines, making such a strategy ineffective. This suggests that competition between search engines is more likely to take place on non-price dimensions, such as the quality of the organic links.

5.2 Competing in accuracy

Suppose now that, as in section 4, the search engines can commit to a matching accuracy D_i^{SE} before consumers make their participation decision. As before, allowing the search engines to choose the targeting accuracy leads firms to view advertising as a fixed cost, and therefore the fees a_1 and a_2 will

[18]When there are no fixed costs $a^* = \hat{a}$ in (11).

[19]This intuition is also present in Wright (2002), although in a different set-up.

[20]If firms could price-discriminate between search engines, they would set prices equal to $p_i = a_i + R^*\phi'(R^*)$. Search engines would then compete by lowering their fees, à la Bertrand, which would lead to an efficient outcome.

not be passed through to consumers. Search engines are thus able to extract all the industry profit. Let us now focus on the choice of D_i^{SE} by each search engine. Recall that D^* is the equilibrium targeting accuracy given by $D^* = R(p, p, D^*)$. Then we have the following:

Proposition 7 *When search engines compete in their choice of D, the only equilibrium with consumer participation is for both search engines to offer an accuracy equal to D^*.*

The proof, which uses the proof of Proposition 4, is in the appendix. In short, competition leads the search engines to offer as accurate a matching as possible, which explains why $D \leq D^*$ is equilibrium. The reason for $D \geq D^*$ is that any lower D would not be consistent with consumer participation, because of the Diamond Paradox.

This outcome is constrained-efficient, given that firms choose their price, since it minimizes the expected number of clicks (1) and leads to the highest matching quality compatible with consumers' participation.

We see therefore that the effects of competition between search engines depend a lot on whether the accuracy of targeting is determined by firms' equilibrium behavior or by the search engines' design. In the first case, competition between search engine can have adverse effects, by providing search engines with incentives to raise their advertising fees. In the second case, competition eliminates search engines' incentives to downgrade the quality of the matching.

6 Conclusion

This paper presents a model of search engine advertising that incorporates targeted advertising and consumer search in a two-sided market framework. The main results show that the targeting technology potentially improves efficiency, by minimizing search costs, reducing mismatch costs, and increasing the competitive pressure among firms, with respect to a benchmark without targeting. However, the search engine's profit-maximizing behavior leads it to charge too high an advertising fee, which results in a rise in the equilibrium price of the good that can offset the efficiency gains. When the search engine determines the accuracy of targeting, the previous distortion is eliminated, as firms no longer pass through the advertising fee to consumers, but another distortion emerges, namely a suboptimal matching quality. The effects of competition between search engines are ambiguous, since it can solve the latter distortion but could actually worsen the former one.

Although the model provides insights regarding the links between the design of the platform and market outcomes, it ignores some dimensions that are potentially important, such as the presence of organic links or the auction mechanism used in practice. Future work will hopefully further improve our understanding of how these aspects interact with each other.

References

ANDERSON, S. P., AND R. RENAULT (1999): "Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model," *RAND Journal of Economics*, 30(4), 719–735.

——— (2006): "Advertising Content," *American Economic Review*, 96(1), 93–113.

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

ARMSTRONG, M., J. VICKERS, AND J. ZHOU (2009): "Prominence and consumer search," *RAND Journal of Economics*, 40(2), 209–233.

ATHEY, S., AND G. ELLISON (2011): "Position Auctions with Consumer Search," *Quarterly Journal of Economics*, 126.

ATHEY, S., AND J. S. GANS (2010): "The Impact of Targeting Technology on Advertising Markets and Media Competition," *American Economic Review*, 100(2), 608–13.

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

BAKOS, Y. (1997): "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science*, 43(12), 1676–1692.

BAR-ISAAC, H., G. CARUANA, AND V. CUNAT (2010): "Information Gathering and Marketing," *Journal of Economics & Management Strategy*, 19.

BAYE, M. R., AND J. MORGAN (2001): "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review*, 91(3), 454–474.

BERGEMANN, D., AND A. BONATTI (2011): "Targeting in Advertising Markets: Implications for Offline vs. Online Media," *RAND Journal of Economics*, 42(2), 414–443.

CAILLAUD, B., AND B. JULLIEN (2003): "Chicken & Egg: Competition among Intermediation Service Providers," *RAND Journal of Economics*, 34(2), 309–28.

CAPLIN, A., AND B. NALEBUFF (1991): "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, 59(1), 25–59.

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

DE CORNIERE, A., AND R. DE NIJS (2011): "Online Advertising and Privacy," mimeo.

Dellarocas, C. (2012): "Double Marginalization in Performance-Based Advertising: Implications and Solutions," *Management Science*, 58(6), 1178–1195.

Diamond, P. A. (1971): "A model of price adjustment," *Journal of Economic Theory*, 3(2), 156–168.

Dukes, A., and E. Gal-Or (2003): "Negociations and Exclusivity Contracts for Advertising," *Marketing Science*, 22(2), 222–245.

Eliaz, K., and R. Spiegler (2011): "A Simple Model of Search Engine Pricing," *Economic Journal*, 121.

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

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

Galeotti, A., and J. L. Moraga-Gonzalez (2008): "Segmentation, advertising and prices," *International Journal of Industrial Organization*, 26(5), 1106–1119.

Gomes, R. (2011): "Optimal Auction Design in Two-Sided Markets," Discussion paper.

Grossman, G. M., and C. Shapiro (1984): "Informative Advertising with Differentiated Products," *Review of Economic Studies*, 51(1), 63–81.

Haan, M., and J. L. Moraga-Gonzalez (2011): "Competing for Attention in a Consumer Search Model," *The Economic Journal*, 121, 552–579.

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

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

Johnson, J. (forthcoming): "Targeted Advertising and Advertising Avoidance," *RAND Journal of Economics*.

Kohn, M. G., and S. Shavell (1974): "The Theory of Search," *Journal of Economic Theory*, 9(2), 93–123.

Mayzlin, D., and J. Shin (2011): "Uninformative Advertising as an Invitation to Search," *Marketing Science*, 30(4), 666–685.

Robert, J., and D. O. Stahl (1993): "Informative Price Advertising in a Sequential Search Model," *Econometrica*, 61(3), 657–86.

Rochet, J.-C., and J. Tirole (2006): "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, 37(3), 645–667.

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

Salop, S. (1979): "Monopolistic competition with outside goods," *Bell Journal of Economics*, 10(1), 141–156.

Stahl, D. O. (1989): "Oligopolistic pricing with sequential consumer search," *American Economic Review*, 79(4), 700–712.

Taylor, G. (2011): "The Informativeness of On-line Advertising," *International Journal of Industrial Organization*, 29(6), 668–677.

——— (Forthcoming): "Search Quality and Revenue Cannibalisation by Competing Search Engines," *Journal of Economics & Management Strategy*.

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

Varian, H. R. (1980): "A Model of Sales," *American Economic Review*, 70(4), 651–59.

Vives, X. (2001): *Oligopoly Pricing: Old Ideas and New Tools*, vol. 1 of *MIT Press Books*. The MIT Press.

White, A. (2009): "Search Engines: Left Side Quality versus Right Side Profits," Discussion paper.

Wolinsky, A. (1983): "Retail Trade Concentration Due to Consumers' Imperfect Information," *Bell Journal of Economics*, 14(1), 275–282.

——— (1986): "True monopolistic competition as a result of imperfect information," *Quarterly Journal of Economics*.

Wright, J. (2002): "Access Pricing Under Competition: An Application to Cellular Networks," *Journal of Industrial Economics*, 50, 289–315.

A Proofs

A.1 Proof of Lemma 3

Before proving the proposition, it is useful to state an intermediary result.

For $v \geq v^*$, let $\delta(v, p^*) \equiv sup\{d \in [0, 1/2] \quad s.t. \quad u(v, d, p^*) \geq 0\}$. $\delta(v, p^*)$ is the largest distance d such that a consumer would buy at price p^* and at distance d if there was no other firm available.

Lemma 5 *In equilibrium, for every $v \geq v^*$, $\delta(v, p^*) \geq R^*(p^*, p^*, D^*)$.*

Proof: Suppose that there is a consumer of type (v, ω), with $v \geq v^*$ such that $\delta(v, p^*) < R^*(p^*, p^*, D^*)$. Let a firm be located in θ_1, with $\theta_1 \in (\omega + \delta(v, p^*), \omega + R^*(p^*, p^*, D^*))$. Suppose that the consumer faces firm θ_1. Because $d(\omega, \theta_1) > \delta(v, p^*)$, the consumer would rather leave the market than buy from θ_1. But since $d(\omega, \theta_1) < R^*(p^*, p^*, D^*)$, the consumer strictly prefers buying than visiting a new firm. This implies that the expected net value of a random search is negative for consumer (v, ω), which contradicts the fact that $v \geq v^*$, since v^* is such that the expected value of a random search is just zero. □

Now we can prove the proposition. The proof is in two stages: (1) if firms set $D^* < R(p^*, p^*, D^*)$, then a firm can profitably deviate by targeting more consumers, (2) if $D^* > R(p^*, p^*, D^*)$, there is always at least one firm that can profitably deviate and lower its targeting distance.

1. Suppose that all firms have a targeting distance D^* smaller than $R^*(p^*, p^*, D^*)$. Take a consumer ω and a firm θ such that $D^* < d(\theta, \omega) < R^*(p^*, p^*, D^*)$. If θ were to deviate and choose to appear to consumer ω, then it would sell the good with probability equal to $P[v \geq p^* + \phi(d(\theta, \omega))|v \geq v^*]$ if ω clicked on its link. Now, from lemma 5, and since $d(\omega, \theta) < R^*(p^*, p^*, D^*)$, we know that $P[v \geq p^* + \phi(d(\theta, \omega))|v \geq v^*] = P[\delta(v, p^*) \geq d(\theta, \omega)|v \geq v^*] = 1$. Thus it would be a profitable deviation.

2. Now suppose that all firms set $D^* > R^*(p^*, p^*, D^*)$. Take a consumer ω, and denote $\overline{\theta}$ the firm which is located at a distance D^* from him. Since $d(\overline{\theta}, \omega) > R^*(p^*, p^*, D^*)$, the probability that ω buys from $\overline{\theta}$ is zero. By reducing its reach, firm $\overline{\theta}$ can increase its profit. □

A.2 Proof of Proposition 1

The equilibrium is obtained through the following steps:

1. Existence and uniqueness of an equilibrium targeting distance $D^* > 0$.

   Lemma 6 *Under assumption 1, and for any price p, the function $r : D \mapsto R(p, p, D)$ has two fixed points: 0 and $D^* \in (0; 1/2)$.*

   Proof: From (2), we see that $r(D)$ is defined by

$$\int_0^{r(D)} \frac{\phi(r(D)) - \phi(x)}{D} dx = s$$

   Using the implicit functions theorem on the open interval $(0; 1/2)$, we get $r'(D) = \frac{s}{r(D)\phi'(r(D))}$. As D goes to zero, $r'(D)$ tends to $+\infty$, because $\lim_{D \to 0} r(D) = 0$ and $\phi'(.)$ is bounded and positive.[21] Moreover, $r(1/2) \leq 1/2$

[21] When $u(v, d, p) = v - td^b - p$ and $b < 1$, the assumption that ϕ' is bounded on $[0, 1]$ does not hold. Still, in that case, $r'(D) = D^{-\frac{b^2}{b+1}} \frac{s}{tb} \left(\frac{(b+1)s}{tb}\right)^{-\frac{b^2}{b+1}}$, and tends to $+\infty$ when D goes to 0.

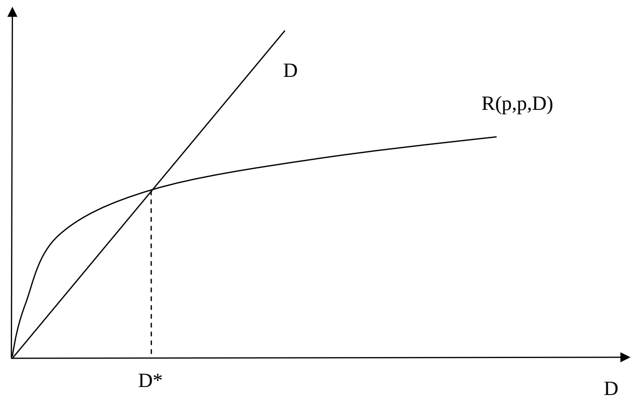

Figure 1: D versus $R(D)$

(by assumption 1), and therefore there must be a $D^* \in (0; 1/2)$ such that $D^* = r(D^*)$. Such a D^* is unique if $r(.)$ is concave. Differentiating $r(D)$ a second time , one gets

$$r''(D) = -sr'(D)[\phi'(r(D)) + r(D)\phi''(r(D))][r(D)\phi'(r(D))]^{-2} \tag{16}$$

By assumption 2, the second term in brackets is positive, and therefore $r(.)$ is concave. In that case, one can see that $r(D)$ is above D when $D < D^*$, and below D otherwise. □

2. Existence and uniqueness of an equilibrium price strategy.

   A firm's profit equals $(p - a)R(p, p^*, D^*) \times \frac{1-F(v^*(a))}{R(p,p^*,D^*)}$ if other firms play (p^*, D^*).

   First let's show that the profit is strictly quasi-concave in the firm's price if (C2) holds. A sufficient condition for that is that $1/R(p, p^*, D^*)$ is convex in p (see Vives (2001) p.149). For notational convenience let us drop the arguments in $R(p, p^*, D^*)$. From Lemma 2 and the implicit functions theorem, one gets $\frac{\partial R}{\partial p} = -\frac{1}{\phi'(R)}$. Straightforward computations show that $1/R(p, p^*, D^*)$ is convex in p if and only if $2\phi'(R) \geq -R\phi''(R)$, which is the case if assumption 2 holds.

   Now that we know that the profit is strictly quasi-concave, and thus that the best response is a function, the following contraction argument ensures uniqueness of a symmetric equilibrium:

   Let $\pi(p, p^*) \equiv (p - a_{SE})R(p, p^*, D^*)$. Since we are looking for symmetric equilibria only, uniqueness is ensured if the best response mapping is a contraction for every firm.

   Using the fact that $\frac{\partial R}{\partial p}(p, p^*, D^*) = -\frac{\partial R}{\partial p}(p, p^*, D^*)$, straightforward computations show that

$$\frac{\partial^2 \pi}{\partial p^2} + \frac{\partial^2 \pi}{\partial p \partial p^*} = \frac{\partial R}{\partial p} < 0$$

   which is a sufficient condition for the best response mapping to be a contraction (see Vives (2001), p.47). There is thus a unique symmetric equilibrium. □

A.3 Proof of Proposition 4

Suppose that a consumer is of type (v, ω), and that firm θ sets a price p_θ while other firms play p^*. Three conditions must be satisfied for trade to occur between the consumer and the firm:

$$d(\theta, \omega) \leq D \quad \text{(SED)}$$

$$v - \phi(d(\theta,\omega)) - p_\theta \geq 0 \quad \text{(IR)}$$

$$d(\theta,\omega) \leq R(p_\theta, p^*, D) \quad \text{(NS)}$$

Condition SED (for *search engine's D*) states that for a trade to happen, it must be the case that the firm is included in the pool of potential matches. Condition IR (*individual rationality*) ensures that buying the good provides a non-negative utility to the consumer. Finally, under condition NS (for *no-search*), the consumer prefers to buy than to continue searching.

Let v^* be the smallest value of v such that a consumer is willing to participate, given D. Let $\overline{x}(v,p,p^*,D)$ be the largest distance such that a consumer of type v buys at price p if other firms play p^*. $\overline{x}$ is the largest distance satisfying (SED), (IR) and (NS). Therefore $\overline{x}(v,p,p^*,D) = \min\{D, \phi^{-1}(v-p), R(p,p^*,D)\}$.

Firm θ's gross profit is then

$$\pi_\theta(p,p^*) = Dp\int_{v^*}^{\overline{v}}\int_0^{\overline{x}(v,p,p^*,D)} \frac{1}{D} f(v)dv = p\int_{v^*}^{\overline{v}} \overline{x}(v,p,p^*,D)f(v)dv \tag{17}$$

The next lemma simplifies the problem, by showing that $\overline{x}(v,p,p^*,D)$ cannot be equal to $\phi^{-1}(v-p)$ (unless it is also equal to D or $R(p,p^*,D)$).

Lemma 7 *For all $v \geq v^*$, if there exists $\overline{d} \leq D$ such that $v - \phi(\overline{d}) - p = 0$, then $\overline{d} \geq R(p,p^*,D)$.*

Proof: Suppose that $\overline{d} < R(p,p^*,D)$. Let $Z^*(v)$ be the expected value of a click (net of search costs) in equilibrium for a consumer of type v. Then

$$\overline{d} < R(p,p^*,D) \iff Z^*(v) < v - \phi(\overline{d}) - p$$

Indeed, $\overline{d} < R(p,p^*,D)$ means that the consumer strictly prefers to buy than to search again, i.e the expected value of a click is smaller than the utility he gets if he buys the product immediately.

Now, we have $v - \phi(\overline{d}) - p = 0$, which implies that $Z^*(v) < 0$. But this contradicts the fact that $v \geq v^*$, because v^* is such that $Z^*(v^*) = 0$ and Z^* is increasing in v. □

Therefore, (17) rewrites

$$\pi_\theta(p,p^*) = p\int_{v^*}^{\overline{v}} \min\left(D, R(p,p^*,D)\right) f(v)dv = p\min\left(D, R(p,p^*,D)\right)[1-F(v^*)] \tag{18}$$

Let D^* be the fixed point of the function $D \mapsto R(p,p,D)$. D^* is the equilibrium level of advertising from section 3, and does not depend on p.

Lemma 8 *If the search engine chooses $D < D^*$, in any symmetric equilibrium, consumers do not participate.*

Proof: Suppose that $D < D^*$. Then, for every $\tilde{p}$, $R(\tilde{p},\tilde{p},D) > D$. (see Lemma 6) Therefore, at any symmetric strategy profile p, demand is inelastic around p. Each firm has an incentive to raise the price by ϵ, since such a deviation is not enough to trigger an additional search by consumers. □

If $D > D^*$, then $\min\left(D, R(p,p^*,D)\right) = R(p^*,p^*,D)$. Therefore the equilibrium price p^* must be such that

$$p^* \in argmax_p pR(p,p^*,D)[1-F(v^*)]$$

Since v^* depends on D, a firm's profit is

$$\pi^*_\theta(D) = p^*(D)R(p^*(D),p^*(D),D)[1-F(v^*(D))]$$

By the envelope theorem,

$$\frac{\partial \pi_\theta^*(D)}{\partial D} = p^*(D)\frac{\partial R(p^*, p^*, D)}{\partial D}[1 - F(v^*(D))] - v^{*'}(D)f(v^*(D))p^*(D)R(p^*(D), p^*(D), D) \tag{19}$$

The first term is positive, and it corresponds to the fact that raising D enables firms to make a higher per-consumer profit. The second term takes into account the change in consumers' participation. We know that as D increases, both search costs and mismatch costs increase. The next lemma gives a sufficient condition for the equilibrium price to be increasing in D, in which case $v^{*'}(D) < 0$.

Lemma 9 *When $D > D^*$, if ϕ is convex, then the equilibrium price is an increasing function of D.*

Proof: The first order condition which determines the optimal price is

$$R(p(D), p(D), D) + p(D)\frac{\partial R}{\partial p}(p(D), p(D), D) = 0 \tag{20}$$

Given that $\frac{\partial R}{\partial p} = -\frac{\partial R}{\partial p(D)}$, totally differentiating (20) gives

$$\frac{dp(D)}{dD} = -\frac{\frac{\partial R}{\partial D}\left(1 + p(D)\phi''(R)(\phi'(R))^{-2}\right)}{\frac{\partial R}{\partial p}} \tag{21}$$

This last expression is non negative since $\frac{\partial R}{\partial D} > 0$ and $\frac{\partial R}{\partial p} < 0$.

A.4 Proof of Proposition 7

Suppose that $D_1^{SE} < D_2^{SE}$. Recall that I assumed that firms charge the same price whether consumers comes from search engine 1 or 2. Then, all the consumers who want to use a search engine will use search engine 1, since it offers them a better matching mechanism. Firms then set a price equal to $p^*(D_1^{SE})$, which maximizes (18). As long as some consumers are willing to participate if the matching is of quality D^{SE}, search engines have an incentive to deviate from a symmetric profile (D^{SE}, D^{SE}), by choosing a quality $D^{SE} - \epsilon$.

Now suppose that $D_1^{SE} < D_2^{SE}$ and $D_1^{SE} < D^*$. Because $D_1^{SE} < D_2^{SE}$, all the consumers who want to participate will chose search engine 1. But, as is shown in Lemma 8 (in the appendix), no consumer wants to participate if the monopolist search engine chooses $D^{SE} < D^*$. Therefore, a deviation from (D^*, D^*) to $(D^* - \epsilon, D^*)$ by search engine 1 is not profitable, for any $\epsilon > 0$. Neither is a deviation to $(D^* + \epsilon, D^*)$. □