Search Advertising

Alexandre de Cornière *

Job Market Paper

November 10, 2011

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 the equilibrium of a game in which firms who advertise on a search engine have to choose their pricing and targeting strategies.

The main results of the paper are the following: (1) in equilibrium, the targeting mechanism endogenously *minimizes* search costs, and improves the quality of the matching between consumers and firms. (2) Giving firms the opportunity to target queries fosters price competition, by improving the value of search for consumers. (3) The per-click fee chosen by a monopolistic search engine is too high with respect to the social optimum, and competition between search engines further *increases* the distortions if firms cannot price-discriminate consumers based on the search engine they use. (4) While designing its platform, a monopolistic search engine must solve a trade-off between attracting many users by offering them a high utility, and softening price-competition between advertisers in order to extract profit. A monopolisic search engine will thus choose a suboptimal matching quality, but competition between search engines eliminates this distortion.

Keywords: search engine, targeted advertising, consumer search.

JEL Classification: D43, D83, L13, M37.

*Paris School of Economics and ENSAE-ParisTech Email: adecorniere@gmail.com. I am grateful to Simon Anderson, Mark Armstrong, Heski Bar-Isaac, Luc Bridet, Bernard Caillaud, Olivier Compte, Gabrielle Demange, Romain de Nijs, Guido Friebel, 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 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]

Advertising through a search engine is the cheapest way of attracting new consumers: in 2005, the cost of attracting a new customer was \$ 8.5 with search advertising, \$20 with yellow pages, \$50 with banner advertising, \$60 with email advertising, and \$70 with traditional mail advertising.[2] Two aspects seem to be of particular importance for the success of search advertising: (i) advertising is *intent-related* and (ii) costs are paid on a *per click* basis.

Intent-related advertising, as opposed to content-related advertising, exploits the possibility of knowing what consumers are looking for, and of targeting them accordingly. Typically, suppose that a hotel located in Paris next to the Eiffel Tower wants to launch an advertising campaign. Using local newspapers may not be a good idea, since people who read them probably do not need a hotel in Paris. Alternatively, advertising through national TV or newspapers is probably too expensive. On the other hand, using a search engine appears like a very natural option, because the firm is able to target consumers who are searching for the keywords "hotel Paris", or, even better "hotel Paris Eiffel Tower". It might also choose not to target users who are searching for "hotel Paris CDG Airport".

"Per click" pricing is aimed at ensuring firms that their investments are not wasted, that is, that the consumers for whom they pay are those who actually see the ad *and* were looking for it. The Hilton Hotel in Paris is certainly not willing to pay every time a search engine user enters the query "Paris Hilton".

In this paper I present 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 model

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

[2]See Batelle (2005)

thus allows to clarify the relationship between consumer search and targeted advertising. The interplay between the two occurs at two stages: *ex ante*, firms and consumers must choose the same keyword to have a chance to be matched. *Ex post*, consumers search across the links knowing that the ads they see are targeted at them.

Section 2 introduces the model, and in section 3, I derive the equilibrium. Section 4 provides a benchmark in which firms are not allowed to use targeting, and must instead advertise randomly across the keywords. 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. I model 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 the 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 6, 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.

Section 7 contains some numerical results obtained for specific distributions and parameters, and in section 8 I present a closely related model, that can accomodate any number of firms, and that can be used as a robustness check.

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 (2010), 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. The two approaches would yield qualitativey similar results, as illustrated in section 8.

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 (forthcoming) 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 (forthcoming) 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 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 (forthcoming), in a related paper, also show that a search engine wants to implement a matching with a suboptimal quality. White (2009) and Taylor (2010) 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. De Corniere and de Nijs (2011) study the incentives of a platform to disclose information about consumers to firms prior to an auction, 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.

[3] I discuss how these papers relate to my model at the end of section 5.

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[4] of mass 1 of firms produce a differentiated good at a zero marginal cost. Each product may be described by a single keyword. Keywords are located on a circle, whose perimeter is normalized to one.[5] 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.

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

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 pays 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.

[4]The continuum assumption makes the derivation of results easier, but is not a necessary condition. See section 8 for an analysis of the case with a finite number of firms.

[5]The use of a circular product space is convenient because it allows to use a one-dimensional strategy to model targeted advertising. The most important results still hold in a model of Chamberlinian competition à la Wolinsky (1986) - Anderson and Renault (1999), see section 8.

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 $\omega \in \mathcal{K}(\theta)$.[6] 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.

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 that 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.

As discussed in the introduction, the model thus 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

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

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.[7]

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.

2. **Firms pricing and targeting:** Firms decide whether to register on the search engine. If they do so they pay 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^* be such that $EU_{SE}(v^*) = 0$. Consumers with $v \geq v^*$ use the search engine, while consumers with $v < v^*$ do not.[8]

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

[7]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).

[8]An alternative and equivalent interpretation of the model would be the following: consumers have to buy one product, and v measures their familiarity with respect to the search engine. If they don't use the search engine, they go to their local mall and get an identical surplus, normalized to zero.

consumer continues to search if and only if $v - \phi(\hat{d}) - \hat{p} < U_R$. Therefore, the strategy of a consumer 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.

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.

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.□

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 $R(p, p^*, D^*)$ 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 proposition.

Proposition 1 *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

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

Thus we have the natural property that the probability that a consumer buys from a firm, *ceteris paribus*, decreases with the firm's price and increases with the expected price of other firms.

3.2 Equilibrium

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.

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

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

where a is the per-click fee paid to the search engine.

Advertising The next proposition provides a characterization of any symmetric equilibrium, that links consumers' and firms' strategies. Existence and uniqueness of such an equilibrium are proven in Proposition 4 under additional assumptions.

Proposition 2 *Any symmetric profile of strategy (p^*, D^*) such that $D^* \neq R^*(p^*, p^*, D^*)$ cannot be an equilibrium.*

The proof of this proposition is provided in the appendix, but the intuition is very simple. Since there is an infinite number of firms, consumers never stop searching before they find a firm closer than the reservation distance. As a result, firms optimally decide to target the consumers that are closer to them than the reservation distance, and only these consumers.

As a corollary, we have the following proposition.

Proposition 3 *If an equilibrium exists, it must be the case that consumers do not search more than once.*

Proposition 3 underlines an important insight of the model, namely that targeting through keywords *minimizes* the amount of search costs incurred in equilibrium. It is the result of two sets of assumptions: (i) all consumers have the same stopping rule, and (ii) firms can target consumers based on their relative willingness to pay (up to a constant v). Part (i) comes from the fact that when the utility function takes the form (1), the optimal reservation distance does not depend on v. Indeed, what matters in the determination of the optimal stopping rule is the marginal rate of substitution between the match quality and income (ϕ') as well as the search cost. In (1), we see that consumers are identical along these lines, and will therefore have the same stopping rule.

It seems important to stress 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.

Also, although Proposition 2 relies on there being an infinite number of firms, Proposition 3 would still hold with any finite number of firms (see Proposition 12 in section 8).

Pricing Thanks to Proposition 2, it is straightforward to find the profit function of a firm if the other firms and consumers play their respective equilibrium strategies. The first order condition on the profit at a symmetric equilibrium will then give the equilibrium price.

Lemma 2 *Suppose that a firm expects that:*

- *all the other firms play the strategy* (p^*, D^*) *where* $D^* = R(p^*, p^*, D^*)$*, and*
- *consumers expect all firms to play* (p^*, D^*) *and thus play* $R(p, p^*, D^*)$*,*

then, the firm's profit function is

$$\pi(p, p^*, a) = (p - a)\frac{R(p, p^*, D^*)}{R(p^*, p^*, D^*)}\left(1 - F(v^*(a))\right) \tag{6}$$

Proof: If a 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 Proposition 2, 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, the 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. The term $1 - F(v^*(a))$ is the number of consumers who use the search engine (i.e such that $v \geq v^*(a)$). □

The previous proof 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 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: $D = 1/2$), some consumers search more than once before buying. In particular, this assumption requires search costs not to be too large.[9] 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 unicity 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 4 *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*

[9] For instance, if $\phi(d) = td^b$, Assumption 1 means $\left(\frac{b+1}{b}\frac{s}{2t}\right)^{\frac{1}{b+1}} \leq \frac{1}{2}$

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 Proposition 2. 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 profit function of a firm would be $\pi(p, p^*, a) = \left(p\frac{R(p,p^*,D^*)}{R(p^*,p^*,D^*)} - aR(p, p^*, D^*)\right)(1 - F(v^*(a)))$, and one would just need to replace a by $aR(p^*, p^*, D^*)$ in the expression of the equilibrium price (9).

Equations (7) and (8) respectively define consumers' optimal search rule and firms' optimal targeting strategy. We have the following:

Proposition 5 *Consumers' reservation distance and firms' targeting distance are strategic complements.*

Proof: We already saw, by (3), that consumers' optimal reservation distance R^* is increasing in D^*. Equation (8) immediately reveals that the optimal targeting distance is increasing in consumers' reservation distance.□

Proposition 5 clarifies the relationship between search and targeting: as advertising becomes more targeted, consumers have incentives to search less, and as consumers search less, firms wish to target consumers more finely. In order to study in more depth the interplay between search and targeting, one could model the precision of the query entered by consumers. Suppose for instance that consumers can enter keywords that corresponds to intervals. A consumer located in ω would then enter a keyword that covers $\mathcal{L}(\omega) = [\omega - L, \omega + L]$. L could be interpreted as a measure of the precision of language. If L is common to all consumers and known by firms, one can show that any symmetric equilibrium must be such that $D^* + L = R^*$, where D^* is replaced by $D^* + L$ in (7), so that the complementarity between search and targeting apply both *ex ante* (if consumers search in broader terms, advertisers target less), and *ex post* (as consumers become more demanding, advertisers must target more finely).

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.

Let $v^*(a)$ be the lowest value of v such that a consumer is willing to use the search engine if the per-click fee is a. $v^*(a)$ is such that $v^*(a) - E[\phi(d)|d \leq R^*] - s - p^*(a) = 0$, where $p^*(a) = a + \phi'(R^*)R^*$ is the equilibrium price. Therefore $v^{*'}(a) = 1$. To further facilitate the analysis, the following assumption is needed:

Assumption 3 *f is log-concave.*

The log-concavity of f implies the log-concavity of $1 - F$. It is satisfied for many common distributions (see Caplin and Nalebuff (1991)).

In equilibrium every consumer with $v \geq v^*(a)$ clicks only once, and so the search engine's profit is

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

Since $1 - F$ is log-concave and $v^*(a) = a + \beta$, with $\beta = E[\phi(d)|d \leq R^*] + s + \phi'(R^*)R^*$, Π^{SE} is also log-concave, which implies that the first order condition for profit maximization is necessary and sufficient.

Let $\hat{a}$ be the fee that maximizes $\Pi^{SE}(a)$. $\hat{a}$ is given by

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

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. To view it even more clearly, let $q = v(a)$ and $z = v(a) - a$. Then (10) rewrites as

$$q - z = \frac{1 - F(q)}{f(q)} \tag{11}$$

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.

Proposition 6 *Under Assumption 3, the optimal fee for the search engine is*

$$a^* = min\{a_\theta, \hat{a}\} \tag{12}$$

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 log-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. Indeed, we have

Proposition 7 *It is socially optimal to set the per click fee arbitrarily close to zero.*

Proof: 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.□

Here, in a model of pure horizontal differentiation with perfect targeting, a slightly positive fee is enough to discipline firms, and to prevent them from targeting too many keywords. 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.

4 A benchmark: non-targeted search

One of the main motivations of 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 is provided by Wolinsky (1983), and was also used in Bakos (1997) in the context of e-commerce. First, I will compare the outcome of the subgame, that is with a fixed. Next, I will look at the effect of targeting on the search engine's pricing decision.

Equilibrium in the subgame To avoid confusion, I will use an index T to refer to the situation in which targeting is allowed, and NT when it is not allowed. Not allowing targeting is formally equivalent to setting $D_{NT} = \frac{1}{2} \geq D_T$. Now, for any price p we have

$$R(p_T, p_T, D_T) = R(p, p, D_T) \leq R(p, p, D_{NT}) = R(p_{NT}, p_{NT}, D_{NT})$$

The two equalities come from Lemma 1 , while the inequality was proven in subsection 3.1. Since the average distance between a consumer and a firm who transact is $E[d|d \leq R]$, where R is the equilibrium reservation distance, we see that targeting reduces the mismatch costs.

More surprisingly, targeting also increases competitive pressure, which results in a lower mark-up. To see this, recall that, with targeting, the mark-up of a firm is

$$p_T - a_T = R_T \phi'(R_T)$$

If targeting is not allowed, the advertising decision is binary: either advertise (and thus target everyone) or not. Let us focus on the interesting case in which firms advertise. The profit of a firm is

$$\pi_{NT}(p, p_{NT}) = \left(p \frac{R(p, p_{NT}, D_{NT})}{R(p_{NT}, p_{NT}, D_{NT})} - \frac{a_{NT}}{2R(p_{NT}, p_{NT}, D_{NT})} \right) (1 - F(v^*_{NT})) - C \qquad (13)$$

Indeed, since firms cannot target consumers, the number of clicks that a firm receives is independent of its pricing strategy. Each click will lead to a purchase with probability $R_{NT}/(1/2)$, and so the expected number of clicks for each consumer is $1/(2R_{NT})$.

Using equation (4) and the first order condition at a symmetric equilibrium, one gets

$$p_{NT} = \phi'(R_{NT}) R_{NT} \qquad (14)$$

Notice that the equilibrium profit of a firm is $(p_{NT} - \frac{a_{NT}}{2R_{NT}})(1 - F(v^*_{NT})) - C$, so that $\frac{a_{NT}}{2R_{NT}}$ *looks like* a marginal cost. However, inspection of (13) reveals that firms do not regard it as a marginal cost, but rather as a cost that will be paid for each user of the search engine, irrespective of whether he buys from the firm or not.

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. Comparing (14) with (9), we see therefore that targeting introduces two effects. The first effect is an increase in firms' marginal cost, passed through to consumers. The other effect is a decrease in firms' mark-up $\phi'(R)R$.

The pass-through effect is discussed at length in Dellarocas (forthcoming). In that paper, he shows that any mechanism in which firms pay on a per-click basis , and in which the number of clicks depends on the price of a firm, leads firms to pass-through the per-click fee.[10] This is the case with targeting, since a firm that charges a low price will target more consumers and thus receives more clicks. On the other hand, without targeting, the probability of a click does not depend on the price, and so the advertising expenses are regarded as a fixed cost (more precisely, as a per-search engine user cost, instead of a per-sale cost, as is the case with targeting).

The decrease in the mark-up is a direct consequence of an improvement in the value of search for consumers. Indeed, with targeting, consumers have a lower reservation distance, and the price elasiticity of demand at this distance is higher. This insight seems quite general, and is discussed further in section 8.

Search engine pricing Recall that when targeting is allowed, the program of the search engine is to maximize $a_T(1-F(v^*(a_T))$ under the constraint that $(1-F(v^*(a_T))(R_T\phi'(R_T) \geq C$.

On the other hand, when firms cannot target consumers, their participation constraint becomes

$$(1 - F(v^*_{NT}))(R_{NT}\phi'(R_{NT}) - \frac{a_{NT}}{2R_{NT}}) \geq C \tag{15}$$

There are two differences between the two participation constraints: the first one was discussed above, and is related to the fact that consumers behave differently under targeting and under no-targeting (R_T versus R_{NT}). The second one is related to the nature of the distortion induced by the per-click fee a. Under targeting, a_T is entirely passed through to consumers, and affects firms only insofar as it reduces consumers participation. On the other hand, when targeting is not allowed, the fee is not passed through to consumers. Therefore the only effect of a_{NT} is to

[10]He actually shows this for a larger class of mechanisms.

reduce the per-unit profit that accrues to firms.

These differences have important implications in terms of optimal pricing for the search engine. Without targeting, the search engine's profit is

$$\Pi_{NT}^{SE}(a_{NT}) = \frac{a_{NT}}{2R_{NT}}\left(1 - F(v_{NT}^{*})\right)$$

It is clear that maximizing this quantity under firms' participation constraint involves a corner solution, namely the highest fee compatible with firms' participation. From this observation we can state the following proposition:

Proposition 8 *Under per-click pricing, firms can make positive profit only if targeting is allowed.*

We see therefore that there might be good reasons for a search engine not to allow firms to use targeted advertising. First, as we just saw, giving firms the ability to target advertising reduces the search engine's profit extraction capabilities. Second, targeting also increases competitive pressure, so that there is less profit to extract from firms. The only way targeting can be profitable for the search engine is if it is possible to extract enough surplus from consumers. This will be the case if consumers' participation is sufficiently price-inelastic (see section 7 for a numerical investigation of this point).

The insight that a search engine might want to restrict firms' ability to target advertisements is developped further in the next section.

5 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 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. For simplicity, I assume thereafter that $C = 0$.

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 3 *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 may 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 9 *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. If the search engine decides to improve the quality of the matching, that is, to set a lower value of D, a hold-up problem (the Diamond paradox) emerges. In this situation, firms set a price at least as high as the lowest value of v among participating consumers. Therefore these consumers, who also have to pay search costs, do not participate, a contradiction.

If the search engine sets a higher value of D, there are two competing effects. On the one hand, competition between firms is less intense, which leads to a higher price and hence a higher per-consumer profit. To understand why competition is less intense, it may be helpful to look at the situation as a bargaining situation. When a firm makes a take-it-or-leave-it offer to a consumer, the consumer compares the offer to his outside-option. This outside option is in fact the continuation value of the search process. If the matching accuracy is low (as is the case with a high D), the continuation value is also low, because the consumer expects to be matched with a firm that is of lower relevance to his query.

On the other hand, because consumers face higher prices, search and mismatch costs, the number of participating consumers decreases. If this drop in consumers participation is not too steep at $D = D^*$, then the search engine will optimally decide to lower the accuracy of the matching process.

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 10 *From the search engine point of view, there exists $\overline{s} > 0$ 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 allow the search engine to make some positive profit, therefore there must exist $\overline{s}$ such that the search engine's profit is maximal. □

Propositions 9 and 10 underline the connection that exists between the search cost s and the targeting disantance D. In fact, s and D play a very similar role from consumers' perspective, as illustrated by this rewriting of equation (2):

$$\int_0^{R^*} (\phi(R^*) - \phi(x))\, dx = sD \tag{16}$$

s and D jointly determine the effective search cost of a consumer, and so it should come as no surprise that the search engine will in both cases choose an interior solution. I discuss the relation between targeting and effective search costs in more depth in section 8.

Discussion : 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. [11] In the present model, I focus on instruments that differ from exclusivity contracts (although these could be incorporated without difficulties). There are several recent contributions that 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

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

Spiegler (forthcoming), 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 level of search costs of consumers, 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 (2010) show 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.

6 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, each search engine receives half the total traffic. I assume that firms can multi-home, and that there are no fixed costs to advertise on search engines.[12] By doing so, I avoid any coordination problem between the two sides of the market.[13]

I assume that search engines have two tools at their disposal: they can allow firms to target advertisements or not, and they choose the per-click fees a_1 and a_2. Regarding firms pricing, I assume that firms have to choose a single price, i.e they are not able to condition their price on the search engine through which consumers reach them. This assumption seems in line with casual empiricism. The timing of the game is the following: (1) both search engines commit to a technology (targeting or no targeting), (2) given the technology choices, they simultaneously choose their per-click fees a_1 and a_2, (3) firms choose a price and a set of keywords (if allowed), (4) consumers choose between the two search engines and their outside option (whose value is zero), (5) the consumers who use a search engine start a sequential search.

[12]One could assume that firms pay the fixed cost only for one search engine, and that registration on the second one is costless.

[13]See Caillaud and Jullien (2003) for a general treatment of this issue.

This game of competition between search engines is very stylized, and will certainly not allow for an exhaustive analysis of the question, but two insights emerge that shed a new light on the issue. The first insight is that although a monopolistic search engine may not want to allow targeting, both platforms adopting the targeting technology is always an equilibrium. The second insight is that in the equilibrium in which both search engines allow targeting, competition will actually raise the price of the good, due to an increase in the price of advertising.

The adoption of targeting. Suppose that p^* is the equilibrium price of the good. Since firms cannot price discriminate, consumers expect to face the same price irrespective of the search engine they use. As a result, they will choose the search engine that minimizes the expected search and mismatch costs. If one search engine lets firms target their ads and the other search engine does not, the former will attract all the consumers. It is then straightforward to see that it is always an equilibrium for both search engines to allow targeting, because a deviation by one search engine from *targeting* to *no targeting* would result in zero profit.

While it is still possible for (*no targeting*, *no targeting*) to be an equilibrium, the set of parameter values such that it is the case is smaller than under monopoly. Indeed, under the candidate strategy profile (*no targeting*, *no targeting*), each search engine receives half the consumers, whereas a deviation to *targeting* attracts all of them.

The effect of competition on prices Suppose that we are at the equilibrium in which both search engines allow firms to target consumers. Then each search engine receives half the traffic, and a firm's profit function is

$$\pi(p, p^*, a_1, a_2) = (p - \frac{a_1 + a_2}{2})\frac{R(p, p^*, D^*)}{R(p^*, p^*, D^*)} \tag{17}$$

and the equilibrium price is therefore

$$p^*(a_1, a_2) = \frac{a_1 + a_2}{2} + R^*\phi'(R^*) \tag{18}$$

At a symmetric equilibrium, the search engine's fee $\tilde{a}$ is given by the first order condition

$$\frac{1}{2}\tilde{a} = \frac{1 - F(v^*(\tilde{a}))}{f(v^*(\tilde{a}))} \tag{19}$$

Proposition 11 *If $\frac{1-F}{f}$ is monotonic, competition between search engines leads to an increase of the per-click fee and of the equilibrium price of the good.*

Proof: The proof relies on the comparison of equations (10) and (19). Since $(1-F)/f \geq 0$, $2(1-F)/f \geq (1-F)/f$, and so the solution to (10) is necessarily smaller than the solution to (19). $\square$

Basically, since only half the consumers use a given search engine, the effect of an increase in this search engine's per-click fee on the equilibrium price is lower than under monopoly. This leads search engines to charge a higher fee in equilibrium, in a *tragedy of the commons* fashion.

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.

Competing in accuracy Suppose now that, as in section 5, the search engines can commit to a matching accuracy D_i^{SE} before consumers make their participation decision. The timing is then: (1) Both search engines simultaneously choose D_i^{SE} and a_i. Firms and consumers can observe them. (2) Firms make their participation and pricing decisions. Firms can multi-home at no cost, and will do it if no search engines provides them with a non negative profit.(3) Consumers choose which search engine to use, if any, and start a sequential search.

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 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.

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 their profit (given by (28) in the appendix). 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 $(D^* + \epsilon, D^*)$. The only equilibrium that allows consumer participation is thus for both search engines to offer an accuracy equal to D^*. This outcome is constrained-efficient, given that firms choose their price, since it minimizes the expected number of clicks (one) 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.

7 An example: the linear-uniform case

In this section I present an example based on specific functional forms. This will allow me to discuss in more depth, although with a loss of generality, the role of certain primitives of the model, such as the search costs, the mismatch costs and the distribution of consumers' willingness to pay.

Let $\phi(d) = td$, so that the utility of a consumer is $v - td - p$ if he buys a product a a distance d and at a price p. Furthermore, let's assume that v is uniformly distributed on $[0; V]$.

The derivation of the equilibrium are provided in the appendix. Below are some figures that illustrate how the search engine's profit and social welfare vary with the different parameters, both with and without targeting.[14]

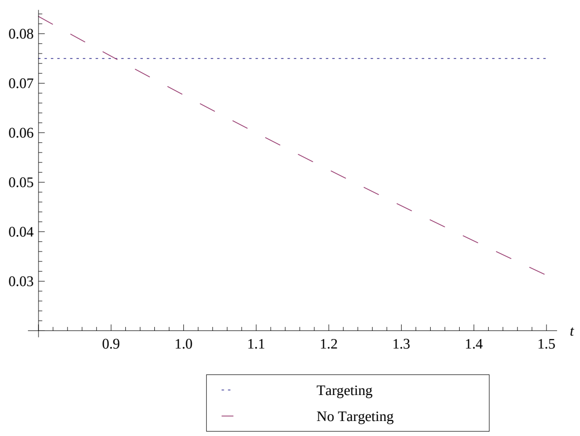

Figure 1: $\Pi^{SE}(t)$

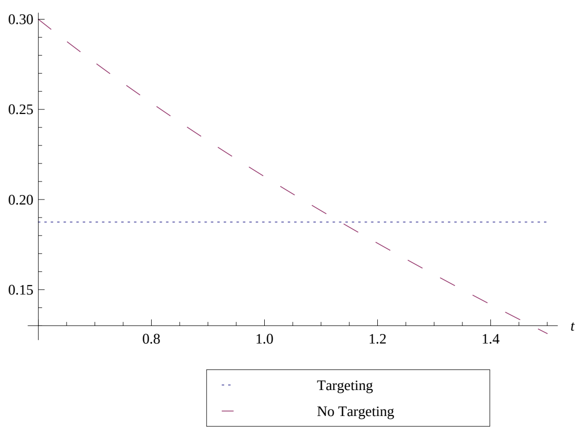

Figure 2: $W(t)$

Figures 1 and 2 represent the profit and the welfare as a function of the mismatch cost t. We see that targeting is desirable, both from a private and a social point of view, when

[14] The benchmark values of the parameters, chosen for their illustrative character, are $s = .15$, $t = 1$, $C = .05$, and $V = 1.2$.

mismatch costs are large, because of its "rescaling" property. In the uniform-linear case, firm's advertising strategy and consumer's search behavior react to a change in t so as to cancel the effect of the change. Without targeting, a change in t reduces both the gains from trade and consumers' participation.

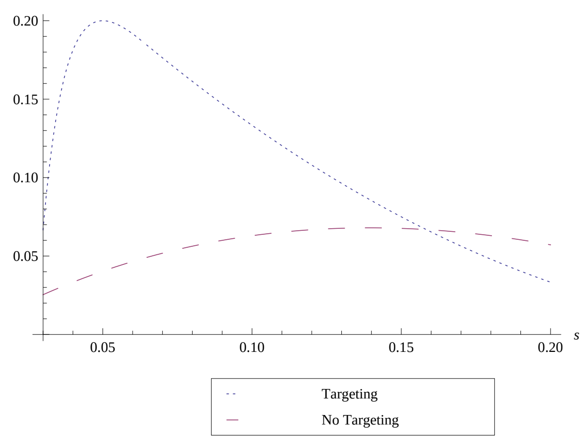

Figure 3: $\Pi^{SE}(s)$

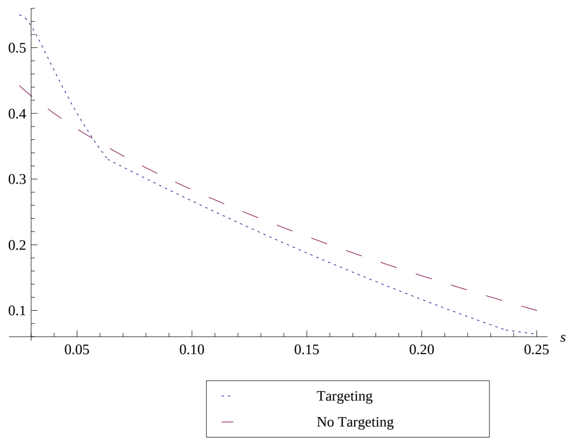

Figure 4: $W(s)$

Perhaps surprisingly, we see that targeting is superior relative to no targeting for small values of the search cost s, both for the search engine and social welfare, but less so for large values of s. This is due to the fact that as s increases, the elasticity of consumers participation with respect to the per-click fee a, which equals $a/(V-4s-a)$, also increases, leading to a drop in participation.

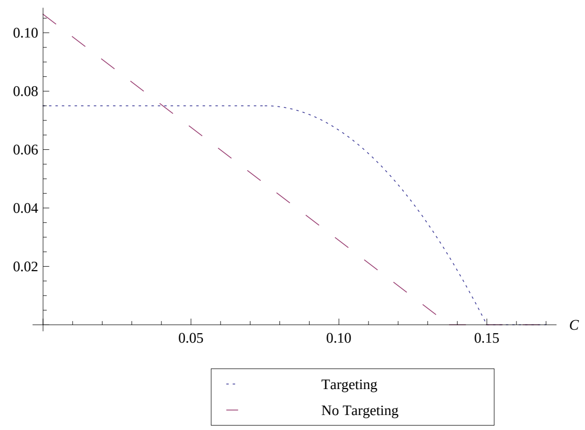

Figure 5: $\Pi^{SE}(C)$

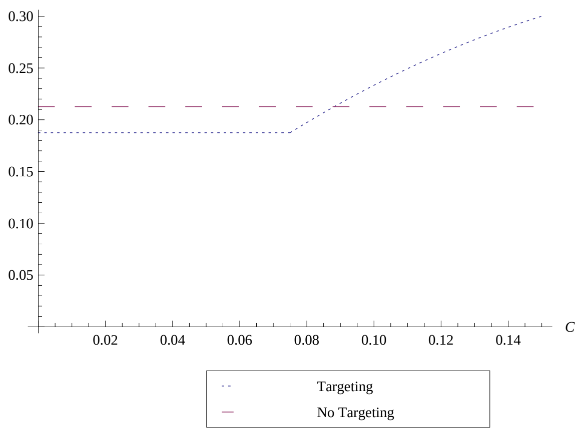

Figure 6: $W(C)$

A rise in firms' fixed costs C has opposite effects on welfare and the search engine's profit. For the search engine, increasing C implies a more binding constraint, both with and without targeting. For the welfare, it has no effect (in the case of *no targeting*) or a positive one (with targeting). The positive effect comes from the fact that an increase in C leads the search engine to lower its fee, which attracts more consumers.

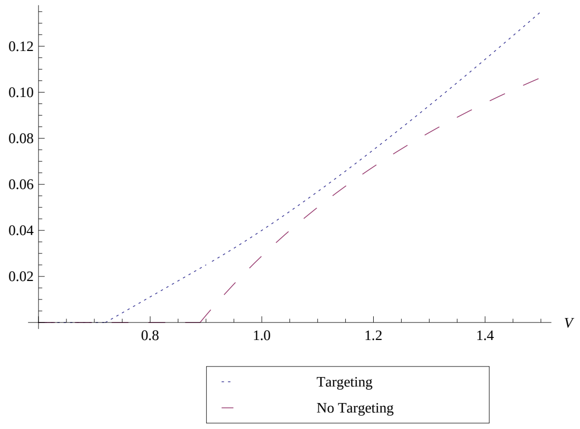

Figure 7: $\Pi^{SE}(V)$

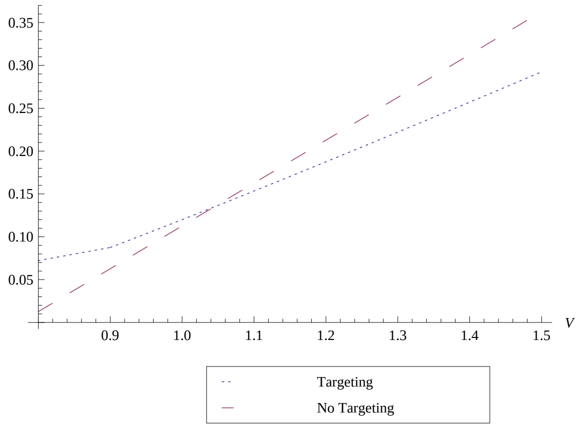

Figure 8: $W(V)$

Figures 7 indicates that higher values of V reinforce the incentives of the search engine to allow targeting, whereas firgure 8 shows that it would be socially optimal not to do it. A higher V leads the search engine to extract more surplus from consumers with its per-click fee, which is only possible with targeting.

A remark on competition Looking at the above figures, it seems that a monopolistic search engine generally tends to over-provide the targeting technology with respect to what is socially desirable. This suggests that competition between search engines, within the framework of section 6, may not be desirable. Indeed, if *Targeting* is not chosen by the monopolist, it is likely not socially efficient. If competition leads to the (*Targeting*, *Targeting*) equilibrium *and* increases the price, welfare is very likely to be reduced..

8 Finite number of firms

The assumption that there is a continuum of firms entails a loss in generality, since many industries are characterized by a small number of firms. However, in this section I show that the main insights about the interplay between targeting and search, namely that search costs are minimized and that price competition is more intense, are robust to a specification with only a small number of firms.

Using the Salop circle with only a small number of firms presents an unappealing feature. Indeed, suppose that firms are symmetrically located along the circle, with a distance $1/n$ between two firms. Then, after observing one firm's characteristics and prices, a consumer learns the exact distribution of positions, and computing his stopping rule becomes more cumbersome. Rather than trying to solve that model, I make the assumption that a consumer's willingness to pay for product i is independent of his willingness to pay for product j. This is the model developed in Wolinsky (1986) and further analyzed by Anderson and Renault (1999).

The utility of consumer l, if he buys product i, is

$$u_i^l = \epsilon_i^l - p_i - ks$$

where k is the total number of visits and s is the search cost. The ϵ_i^l are independently and identically distributed according to a cumulative distribution F, continuous and increasing on some interval $[a; b]$, with a density f. f is also assumed to be log-concave, which will ensure existence of the equilibrium.[15]

Prior to search, a consumer enters a keyword. As before, I assume that the keyword precisely describes what the consumer is looking for, so that conditional on a keyword, the vector of ϵ is deterministic. Stated differently, if consumer l enters a keyword ω, the utility that he will get if he buys product i is $\epsilon_i(\omega)$. Consumers do not know the functions $\epsilon_i(.)$ (and therefore still have to search), but I assume that firms do. This means that a firm knows how well it performs for a given keyword, but also how well its competitors perform. Such an assumption is more relevant for industries with a small number of firms, in which firms can devote resources to study their rivals' products in detail. Since the aim of this section is to examine how robust the previous results are to a set-up with few firms, the assumption seems appropriate.

The targeting mechanism is very similar to the "interval mechanism" used above. Each firm i can choose a set $\mathcal{K}_i \subset \mathcal{R}^n$. Whenever a consumer enters a keyword ω, if $(\epsilon_1(\omega), ..., \epsilon_n(\omega)) \in \mathcal{K}_i$, firm i is willing to appear on the screen. For example, a firm could decide to appear whenever the willingness to pay of a consumer is above a certain threshold. Alternatively, it could decide to appear only to consumers for whom it is the best match.

Similarly to the analysis above, and to the literature on consumer search, the optimal stopping rule for consumers is stationary as long as there is at least one firm left. If a consumer has visited all the firms, he buys from the one that gives him the highest utility. As in Anderson and Renault (1999) and Perloff and Salop (1985), consumers have no outside option, and so they have to buy one product. Results would be qualitatively similar under the more natural assumption that the match value of the product a consumer buys must be at least as large as the price, but ignoring this constraint allows one to find an explicit solution for the equilibrium price.

In equilibrium, the stopping rule is characterized by a reservation utility $\hat{x}$, which is the

[15] see Anderson and Renault (1999) for a discussion.

solution of

$$\int_{\hat{x}}^{b} (\epsilon - \hat{x}) f_T(\epsilon) d\epsilon = s$$

where f_T is the pdf of ϵ_i *conditional on the keyword being targeted by firm* i. If a consumer faces an offer at price p_i when the equilibrium price is p^*, he accepts the offer if and only if $\epsilon_i - p_i \geq \hat{x} - p^*$.

In order to focus on the effect of targeting when there is consumer search, and not on the role of the search engine, I assume that the per-click fee is set arbitrarily close to zero.[16] Thus if a firm is indifferent between targeting a keyword and not targeting it, I assume that it chooses not to target it.

With a finite number of firms, the game in which firms simultaneously choose their price and advertising strategy does not have an equilibrium in pure strategies. Instead of trying to find the mixed strategy equilibrium, I modify the timing of the game as follows: firms choose their price first, then observe everyone else's price. Only after observing the prices can firms decide which consumers to target. I discuss the implications of such a change at the end of the section.

The "no search" result The following proposition extends Proposition 3 to a different product space and to a situation with an arbitrary (and finite) number of firms.

Proposition 12 *In any symmetric equilibrium in pure strategies in the model with i.i.d match values, consumers never search more than once.*

Proof : The following is a proof of a stronger claim, namely that if all the firms except firm i have chosen a price p^* and if firm i has chosen a price p_i, then the outcome of the subgame in which firms target consumers is such that consumers will not search more than once. Let $\hat{x}(p_i) = \hat{x} - p^* + p_i$. If a consumer clicks on i's link, he immediately buys if and only if $\epsilon_i \geq \hat{x}(p_i)$. It is clear that firm i always wants to target consumers who would immediately buy. Now suppose that the consumer is such that $\epsilon_i < \hat{x}(p_i)$, and that firm i targets him nonetheless. There are two possibilities. Either firm i is the best offer among the firms that target him, or it is not. It is straightforward to see that firm i wants to target the consumer in the former case, because he will come back to i after having visited all the firms. In the latter case, the consumer won't come back, and so the firm prefers not to target him.

This means that there are only two possibilities in equilibrium for a given consumer. The first one is that he sees several links on his screen, which implies that all the firms provide him

[16] The discussion therefore neglects the price distortion.

with a larger utility than his reservation utility. In this case the consumer buys from the first firm he visits. [17] There is also the possibility that the consumer sees only one firm, which may or may not provide him with a larger utility than his reservation utility. Since he sees only one firm, he searches only once (if at all). □

By Proposition 12, one sees that the fact that targeting minimizes search costs does not rely on a large number of firms being present on the market. Rather, what matters here is that firms are able to target consumers as finely as they want. The result would not necessarily hold if firms could not condition their targeting strategy on the value of the match between a consumer and all the other firms. For instance, if we only allowed firms to set a threshold and target all the consumers such that the value of the match is larger than the threshold, there could be equilibria in which the targeting threshold is lower than the stopping rule threshold $\hat{x}$, thus resulting in the possibility of multiple visits in equilibrium.

Targeting endogenously reduces search costs. One of the main results of the paper is that targeting leads to an increase in the intensity of price competition. The random match framework, among other things, allows to exactly understand the force that leads to this result.

Consider the case of non-targeted search, extensively studied in the literature. In equilibrium , as long as there is at least one firm that has not been sampled by the consumer, the optimal search rule is to continue as long as the value of a match is less that $\hat{x}$, where $\hat{x}$ is given by

$$\int_{\hat{x}}^{b} (\epsilon - \hat{x})\, f(\epsilon) d\epsilon = s \tag{20}$$

The left-hand side of equation (20) measures the expected improvement upon an offer of quality $\hat{x}$, and the right-hand side is the search cost. In particular, the optimal threshold $\hat{x}$ is a decreasing function of the search cost.

Now, suppose that firms can target consumers. By Proposition 12, if a consumer receives several advertisements in equilibrium, it must be that all of the match values are higher than the optimal threshold $\tilde{x}$ (which differs from the threshold without targeting $\hat{x}$). The expected improvement from a match value z is given by $\int_{z}^{b} (\epsilon - z)\, f(\epsilon | \epsilon \geq \tilde{x}) d\epsilon = \int_{z}^{b} (\epsilon - z)\, \frac{f(\epsilon)}{1-F(\tilde{x})} d\epsilon$.

Therefore the optimal threshold $\tilde{x}$ is given by

$$\int_{\tilde{x}}^{b} (\epsilon - \tilde{x})\, f(\epsilon) d\epsilon = (1 - F(\tilde{x}))s \equiv \tilde{s} \tag{21}$$

[17] This corresponds to the equilibrium in section 3.

Therefore one may see the threshold $\tilde{x}$ as the optimal rule of non targeted search process in which consumers' search costs is $\tilde{s} \leq s$. As a result, we have the following:

Lemma 4 *The optimal threshold in the consumers' search rule is higher with targeting.*

By improving the quality of the pool of matches that a consumer can sample, targeting induces consumers to be more demanding. As we will see below, this effect is the one that drives the price down compared to the case of non-targeted search.

Equilibrium pricing. As we saw just above, one can rewrite the optimal stopping rule with targeting as a stopping rule without targeting but with a lower search cost. It turns out, perhaps surprisingly, that the profit function of a firm under targeting can be written just as in the case without targeting. Indeed, if firm i chooses a price p_i and other firms choose p^*, in the second stage firm i will target all the consumers such that $\epsilon_i \geq \tilde{x} - p^* + p_i$, as well as the consumers such that $\tilde{x} - p^* + p_i > \epsilon_i \geq \max_{j \neq i} \epsilon_j - p^* + p_i$.

The demand for i's product will be given by

$$D_i(p_i, p^*) = \sum_{k=0}^{n-1} \frac{1}{k+1} \binom{n-1}{k} [1-F(\tilde{x})]^k F(\tilde{x})^{n-k-1} [1-F(\tilde{x}+p_i-p^*)] + \int_a^{\tilde{x}+p_i-p^*} F(\epsilon+p_i-p^*)^{n-1} f(\epsilon) d\epsilon \tag{22}$$

The first term on the right-hand side corresponds to consumers who see several advertisements. A consumer who sees $k+1$ ads, including firm i's ad, is such that k firms among $n-1$ present a match value greater than $\tilde{x}$. On top of that, the match value for firm i must be greater than $\tilde{x} - p^* + p_i$. Such a consumer visits firm i with probability $1/(k+1)$. The second term corresponds to consumers who would continue to search if they could, but who will be targeted by firm i only.

After some straightforward manipulations, one gets

$$D_i(p_i, p^*) = \frac{1}{n} [1 - F(\tilde{x} + p_i - p^*)] \frac{1 - F(\tilde{x})^n}{1 - F(\tilde{x})} + \int_a^{\tilde{x}+p_i-p^*} F(\epsilon + p_i - p^*)^{n-1} f(\epsilon) d\epsilon \tag{23}$$

which corresponds to expression (5) in Anderson and Renault (1999)[18], the only difference being that the threshold $\tilde{x}$ is higher than the threshold without targeting $\hat{x}$. It follows from this observation that the equilibrium price is given by almost the same first order condition for the case with targeting and without targeting respectively:

[18] with $L = 1$ and $a = -\infty$, and $\mu = 1$

$$p_T^* = \frac{1}{\frac{1-F(\tilde{x})^n}{1-F(\tilde{x})} f(\tilde{x}) - n\int_a^{\tilde{x}} f'(\epsilon)F(\epsilon)^{n-1}d\epsilon} \tag{24}$$

$$p_{NT}^* = \frac{1}{\frac{1-F(\hat{x})^n}{1-F(\hat{x})} f(\hat{x}) - n\int_a^{\hat{x}} f'(\epsilon)F(\epsilon)^{n-1}d\epsilon} \tag{25}$$

Using Anderson and Renault (1999)'s Proposition 1, which states that the equilibrium price is a decreasing function of search costs, we can conclude that $p_T^* \leq p_{NT}^*$.

Proposition 13 *Allowing firms to target consumers reduces the equilibrium price.*

Two comments are in order at this point. First, one can easily see that in order to obtain the previous result, it is essential that consumers incur search costs. Indeed, if it was not the case, as in Perloff and Salop (1985), targeting would have no impact on the equilibrium price. Firms would choose their price, and then target all the consumers that would buy. Each consumer would only face one firm, but since the price is fixed that firm could not use its monopoly power. The second comment is related to this last point. The timing that I have used here, in which firms choose their price prior to their advertising strategy, prevents firms from exploiting the "captive consumers" who only see one advertisement.[19] Therefore, if one believes that this effect is important in practice, the conclusion that the price decreases with targeting is less likely to hold, especially if the number of firms is small. For a large number of firms, the probability that no firm gives a match value greater than $\tilde{x}$ becomes negligible, and the result would still hold.

9 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 is potentially very efficient, 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

[19] see Roy (2000), Iyer, Soberman, and Villas-Boas (2005) for examples in which the timing is different, but in which consumers do not actively search.

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.

There are several ways through which the model could be extended. The most natural one is to incorporate an auction stage that determines the order through which consumers visit firms. The version of the model presented in section 8 would probably be more amenable to such an approach, since it features a finite number of firms. This version would also be useful in order to study in more depth the question of entry by advertisers, and in particular whether the search engine induces excessive or insufficient entry.

The main advantage of the circular model is that it allows to model how keywords relate to products in a very convenient way. This approach could be exploited further. For example, one could think of a model in which some consumers search for generic keywords ("Luxury Hotel Paris", represented as an interval) while others search for brands ("Ritz Paris", represented as a point). A seemingly promising avenue for future research is to provide a model that would relate how consumers search (i.e what kind of keywords they enter) with the amount of information that they have been exposed to prior to search, through advertising or word-of-mouth.[20]

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.

[20]See Goldfarb and Tucker (2011) and Joo, Wilbur, and Zhu (2011) for empirical studies of the relationship between online search and offline advertising.

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.

Batelle, J. (2005): *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Nicholas Brealey Publishing.

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 (forthcoming): "Paid Placement: Advertising and Search on the Internet," *Economic Journal*.

De Corniere, A., and R. de Nijs (2011): "Online Advertising and Privacy," mimeo.

Dellarocas, C. (forthcoming): "Double Marginalization in Performance-Based Advertising: Implications and Solutions," *Management Science*.

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 (forthcoming): "A Simple Model of Search Engine Pricing," *Economic Journal*.

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.

Goldfarb, A., and C. Tucker (2011): "Substitution between Online and Offline Advertising Markets," *Journal of Competition Law and Economics*, 7(1), 37–44.

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. (2010): "Targeted Advertising and Advertising Avoidance," Working paper.

Joo, M., K. C. Wilbur, and Y. Zhu (2011): "Does Television Advertising Influence Online Search," working paper.

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.

Perloff, J. M., and S. C. Salop (1985): "Equilibrium with Product Differentiation," *Review of Economic Studies*, 52(1), 107–20.

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. (2010): "Search Quality and Revenue Cannibalisation by Competing Search Engines," working paper.

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*.

A Proofs

A.1 Proof of Proposition 2

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 4

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

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

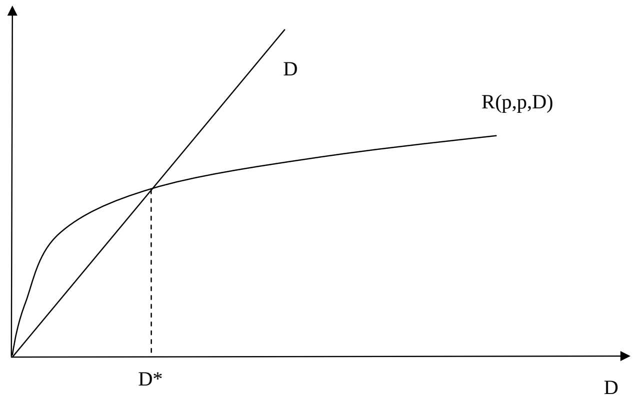

Figure 9: D versus $R(D)$

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, $r(1/2) \leq 1/2$ (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

$$r''(D) = -sr'(D)[\phi'(r(D)) + r(D)\phi''(r(D))][r(D)\phi'(r(D))]^{-2} \tag{26}$$

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 Proposition 1 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 $\pi(p,p^*) \equiv (p-a)R(p,p^*,D^*)$. 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).

[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.

There is thus a unique symmetric equilibrium. □

A.3 Proof of Proposition 9

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

$$\pi_\theta(p, p^*) = Dp \int_{v^*}^{\overline{v}} \int_0^{\overline{x}(v,p,p^*,D)} \frac{1}{D} f(v)dv = p \int_{v^*}^{\overline{v}} \overline{x}(v, p, p^*, D) f(v)dv \tag{27}$$

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

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

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, (27) rewrites

$$\pi_\theta(p, p^*) = p \int_{v^*}^{\overline{v}} \min\left(D, R(p, p^*, D)\right) f(v)dv = p \min\left(D, R(p, p^*, D)\right) [1 - F(v^*)] \tag{28}$$

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 $\min(D, R(p, p^*, D)) = R(p^*, p^*, D)$. 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,

$$\frac{\partial \pi^*_\theta(D)}{\partial D} = p^*(D)\frac{\partial R(p^*, p^*, D)}{\partial D}[1 - F(v^*(D))] - v^{*'}(D)f(v^*(D))p^*(D)R(p^*(D), p^*(D), D) \tag{29}$$

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

$$R(p(D), p(D), D) + p(D)\frac{\partial R}{\partial p}(p(D), p(D), D) = 0 \tag{30}$$

Given that $\frac{\partial R}{\partial p} = -\frac{\partial R}{\partial p(D)}$, totally differentiating (30) gives

$$\frac{dp(D)}{dD} = -\frac{\frac{\partial R}{\partial D}\left(1 + p(D)\phi''(R)(\phi'(R))^{-2}\right)}{\frac{\partial R}{\partial p}} \tag{31}$$

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

B The linear-uniform case: calculations

Under the linear-uniform specification, Assumption 1 rewrites $4s < t$. From Proposition 4, the symmetric equilibrium of the game is given by

$$D^* = R^* = \frac{2s}{t} \tag{32}$$

$$p^*(a) = a + 2s \tag{33}$$

which implies that the indifferent consumer is such that $v^*(a) = 4s + a$. Using equation 10 and the definition of a_θ, the optimal per-click fee for the search engine is

$$a^* = \min\{\frac{V - 4s}{2}, (1 - \frac{C}{2s})V - 4s\} \tag{34}$$

Without targeting, the equilibrium, as described in equations (2)[22] and (14) is given by

$$R_{NT} = \sqrt{\frac{2s}{t}} \tag{35}$$

$$p_{NT} = \sqrt{st} \tag{36}$$

so that the infdifferent consumer is $v_{NT} = 2\sqrt{st}$. From firms' participation constraint (15) which must be binding, the optimal fee for the search engine is

$$a_{NT} = 2s - \frac{VC}{V - 2\sqrt{st}}\sqrt{\frac{2s}{t}} \tag{37}$$

The search engine's profit is then straightforward to compute. Social welfare with and without targeting is given respectively by

$$W_T = \int_{v^*(a_T)}^{V} \frac{v - 2s}{V} dv$$

and

$$W_{NT} = \int_{v_{NT}}^{V} \frac{v - \sqrt{st}}{V} dv$$

Optimal design If we allow the search engine to choose D, the equilibrium is

$$R^*(D) = \sqrt{\frac{2sD}{t}} \tag{38}$$

$$p^*(D) = \sqrt{2stD} \tag{39}$$

and the indifferent consumer is $v^*(D) = 2\sqrt{2stD}$. From equation (27), a firm's profit is

$$\pi_\theta(D) = 2sD\left(\frac{V - 2\sqrt{2stD}}{V}\right) \tag{40}$$

and is maximized for $D = \frac{V^2}{18st}$ (provided this expression is larger than the *laisser-faire* targeting distance $D^* = \frac{2s}{t}$). Thus the optimal targeting distance for the search engine is

$$D^{SE} = \max\{\frac{V^2}{18st}, \frac{2s}{t}\} \tag{41}$$

[22] with $D = 1/2$