Targeting with Consumer Search: an Economic Analysis of Keyword Advertising

Alexandre de Cornière*

February 18, 2010

Abstract

This article investigates the role of a search engine as an intermediary between firms and consumers. Search engines enable firms to target consumers who have revealed some specific needs through their query. In a framework with horizontal product differentiation, imperfect product information and in which consumers incur search costs, I show that introducing a mechanism which enables firm to target consumers tends to reduce social inefficiencies and the equilibrium price. A profit maximizing search engine has incentives to design the mechanism so as to soften price competition between firms, in order to extract profit from them, and this may or not be socially efficient. Sufficient conditions for efficiency or inefficiency of the outcome are provided.

Keywords: search engine, targeted advertising, consumer search, product differentiation.

JEL Classification: D43, D83, L13,M37.

*Paris School of Economics, Email: adecorniere@gmail.com. I am grateful to Heski Bar-Isaac and Frederic Koessler for very useful suggestions. The paper has also benefited from presentations at the Paris School of Economics, Université de Cergy, Northwestern University (2nd Transatlantic Theory Workshop), and Stanford University (5th workshop on Ad Auctions).

1 Introduction

In 2007, online advertising expenses amount to 21 billion dollars in the United States, which is about 7% of total advertising expenses (Evans (2008)). The main actors in this industry are the internet search engines, such as Google or Yahoo. Indeed, 40 percent of online advertising is search-related. Moreover, search-related advertising expenses have been multiplied by seven between 2002 and 2006.

It turns out that advertising through a search engine is the cheapest way of attracting new consumers (see Batelle (2005)). One may wonder what are the ingredients that make it so profitable. Two aspects seem to be of particular importance, namely the facts that (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 to know what consumers are looking for. Typically, when a consumer enters keywords such as "ink jet photo printer" on a search engine, he or she reveals a need, and firms which can satisfy this need are able to target this consumer, instead of having to rely on less-relevant characteristics of the audience which would be used with more traditional advertising, such as TV or magazines.

The other ingredient, the "per click" pricing, is aimed at ensuring announcers that their investments are not wasted, i.e that the consumers for whom they pay are those who actually see the ad *and* were looking for it.

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*, and consumers do not have prior knowledge of firms' prices or positions on the circle. The search engine is an intermediary between firms and consumers: announcers choose which keywords they want to target, and consumers enter keywords and then search sequentially (and costly) at random through the links that appear. I do not study the format of the auction through which slots are allocated[1]. Rather, I shall explore the links between what information is revealed by the search engine and the resulting market outcomes.

In sections 2, 3 and 4, the search engine is "neutral", in the sense that it does not modify the messages which it receives from firms and consumers. I compare the outcome with a situation in which there is no intermediary through which firms can target consumers. Basically, I find that consumers benefit from firms' ability to target through three channels: better matches, smaller

[1]The per-click cost is determined through a *Generalized Second Price Auction* (See Edelman, Ostrovsky, and Schwarz (2007), Varian (2007), for the properties of this mechanism)

expenses in search costs and lower prices than without targeting. The fact that consumers find products more suited to their tastes is rather in line with the intuition that one may have before going into the details of the model. Indeed, since announcers target them, consumers no longer receive non-relevant advertisements and thus choose from a better pool of offers. The model also predicts that, with targeting, consumers do not visit more than one firm, and thus minimize their search costs. These two results combine to improve the efficiency of advertising: the social costs due to imperfect information (bad matches and high search costs) are significantly reduced and thus the presence of a search engine contributes to improving social welfare.

Consumers are the main beneficiaries of this welfare improvement, for they also benefit from a lower price of the final good. To grasp the intuition of this result, it is useful to emphasize that in the model consumers actively search for goods. This search process is sequential: after learning an offer, a consumer compares this offer to the expected offer that he is going to receive if he continues searching (his "outside option"). If the difference between the outside option and the current offer is larger than the search cost, then the consumer continues searching. Now, when firms can target consumers, the relative quality of the outside option increases, because consumers know that the offers they will get after rejecting the current one are targeted at them, and thus very likely to be good matches. Thus, since firms essentially compete against outside options, a rise in the quality of the latter implies less bargaining power for the firms and thus a lower price for the final good.

The "neutral" matching technology is an approximation of how search engines really proceed. For instance, Google sorts announcers 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. Such practices may be regarded as an attempt by the search engine to influence the accuracy of the information transmitted by firms, in one way or another. In section 5 I look at a situation in which the search engine can introduce an arbitrary level of noise (in a sense made precise below) in the information revealed to consumers. The analysis reveals that the equilibrium price is a U-shaped function of the level of noise. For low levels of noise, firms behave like monopolies, whereas the competitive pressure is higher for intermediate levels of noise and decreases thereafter. The reason for this switch from monopoly-like equilibrium to oligopoly-like equilibrium lies in the fact that when there is little noise, the outside-option constraint (consumers must be better-off if they buy than if they continue searching) is never binding and therefore firms are not constrained by their competitors. For higher levels of noise, the price is increasing function because increasing the noise makes consumers' outside option less attractive, and thus improves firms' bargaining power.

If a profit-maximizing search engine was to design a matching mechanism, it would try to implement as high a price as possible, while ensuring consumers' participation. The reason for that is that the search engine can extract firms' profit through advertising fees, but cannot extract consumers' surplus. Depending on the shape of consumers' preferences, maximizing the price implies either to minimize or maximize the level of noise. In the former case, the result is a socially optimal situation, while the latter case is inefficient. I show that when "transportation costs"' are linear, we are in the latter - inefficient - case.

Related literature

This paper is related to the literature on search models and advertising, as well as to more recent contributions which study internet search engines and two-sided markets.

The literature on search models on a product market has provided important insights. In a seminal paper, Diamond (1971) shows that as soon as there is a positive cost for consumers to learn the price of a homogenous good, the only equilibrium outcome is for all the firms to charge the monopoly price. This result is known as the "Diamond paradox". Stahl (1989) studies situations in which consumers have different search costs. This heterogeneity implies that some consumers will be better informed than others. As in Varian (1980), the fact that consumers differ in their level of information generates equilibrium price dispersion, because some firms want to compete for the informed consumers (i.e. with low search costs) whereas other firms charge high prices and sell only to the uninformed consumers (i.e. with high search costs).

When products are differentiated, the price is an increasing function of the search cost and entry is generally excessive with respect to the social optimum (Anderson and Renault (1999), Wolinsky (1984)). Bakos (1997) studies the impact of a drop in search costs due to the development of electronic markets. He highlights the importance of the nature of the information which is costly to get: prices go down when consumers have a cheaper access to price information, whereas prices rise when it is easier to get product information.

The relationship between advertising and consumer search is not a new topic: in Robert and

Stahl (1993), consumers may learn the price of homogenous goods either by receiving an ad or by searching actively. In a monopoly framework with uncertainty regarding the product's characteristics, Anderson and Renault (2006) study the optimal content of advertising. They highlight the differences between product information and price information, and show that the optimal advertising content varies with consumers' search costs.

The issue of targeting has received rather little attention in the economic literature. Esteban, Gil, and Hernandez (2001) show that in a monopoly framework, firms' ability to target consumers reduces both consumers' and total surplus. Iyer, Soberman, and Villas-Boas (2005) study targeting in a duopoly. Targeting induces endogenous differentiation of products, since firms advertise less to consumers who do not have "strong" preferences. The average price thus goes up. In their model, targeted advertising is more valuable to firms than targeted pricing. Also, interestingly, the effect of targeting on the optimal level of advertising depends on the initial cost of wasted advertising.

Van Zandt (2004) deals with the issue of information overload. He shows that, when firms can target consumers, a rise in the cost of advertising induces firms to send more accurate information to consumers, and this alleviates the effects of information overload.

Some recent papers study the interactions between firms and consumers on a search engine, but focus more on the ranking of ads than on the choice of relevant keywords. Athey and Ellison (2007) show that there exists an equilibrium in which efficient firms get the higher slots, and in which consumers search sequentially from top to bottom. They discuss mechanisms which could improve the efficiency of the generalized second-price auction.

Armstrong, Vickers, and Zhou (2009) study the impact of prominence on the market outcome. A prominent firm is sampled first by all consumers. Interestingly, they show that when firms are symmetric, prominence reduces welfare. On the other hand, when firms are vertically differentiated, firms with better quality would be willing to pay more to be made prominent, while consumers would sample these firms first even if they did not have to. Making the best firm prominent would improve welfare. This underlines the force that drives "better" firms to bid aggressively in order to secure the best slots, even when pricing is endogenous.

Finally, my paper is related to the growing literature on two-sided markets, seminal papers of which include Armstrong (2006), Caillaud and Jullien (2003), Rochet and Tirole (2006) orBaye and Morgan (2001) . My approach (see section 5) is a bit different from these papers, in the sense that I do not use a reduced-form way of modelling interactions between agents on the platform, in order to account for some important details. Hagiu and Jullien (2007) have a similar approach, although in a different set-up. Here I do not allow complete flexibility in terms of pricing, and I instead focus on the design of the matching process as a way to increase the platform's profit.

To the best of my knowledge, this paper is the first to explicitly model the transmission of information from firms to consumers through a search engine, and how this process may affect prices and welfare. It is also the first to study a model of consumer search with targeted advertising.

2 The model

Description of the market and of preferences

The framework is based on Wolinsky (1983). Consider a market where a continuum of mass μ_F of firms (or "announcers") 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. Thus a firm is characterized by the position of its product's keyword on the circle. Keywords' positions are denoted by $x \in [0;1]$.

There is a continuum of mass μ_C of consumers, each one having a favorite, or ideal, brand (or keyword), $y \in [0;1]$.

Consumers have use for at most one product, and the utility that a consumer y gets from consuming a good located in x, with $d(x,y) = d$, is

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

where p is the price of the good and v is decreasing and twice continuously differentiable.

Advertising technology

Consumers have imperfect information about firms' characteristics: they do not know firms' position on the circle nor their price, and thus have to search before buying. Interactions between firms and consumers are only possible through a search engine. The search engine plays the role of an intermediary: firms communicate the set of keywords that they want to target, and consumers communicate the keyword they are interested in. Consumers cannot enter several keywords at the same time. If a certain keyword is entered by a consumer, all the firms who want to target this keyword appear on the consumer's screen. If a consumer clicks on a firm's link, he incurs a search cost $s > 0$ and learns the price and position of this firm. The search cost corresponds to the time spent in order to find the relevant information on a website. On the other hand, when a consumer clicks on an announcer's link, the announcer pays a fee $a > 0$ to the search engine.

The assumption that consumers do not observe anything before clicking on a link seems appropriate in many contexts. Indeed, announcers 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 infer correctly that a firm which targets them is not farther than a certain distance.[2].

After a consumer has sampled a firm and learned its price and position, he can come back at no cost (recall is costless). It is the case if for instance consumers open a new window every time they click on a link.

Strategies and equilibrium concept

A strategy for a firm x consists in the simultaneous choice of *(i)* a price $p \in [0; v(0)]$, and *(ii)* a set of keywords $S = [x - D; x + D]$, with $D \in [0; 1/2]$.[3]

Consumers's strategy consists in choosing an optimal stopping rule, that is in setting a reservation distance R, such that the consumer is indifferent between buying a product at a distance R and continuing to search. R depends on the price that the consumer observes (p) as well as the strategy that he expects firms to play ($\overline{\sigma} \equiv (\overline{p}, \overline{D})$). Thus I will use the notation $R(p, \overline{\sigma})$ to describe the stopping rule. The optimality of such a strategy is discussed at length in Stahl (1989) and Anderson and Renault (1999). Basically, when recall is costless, as long as there is at least one firm left to visit, the problem faced by the consumer is stationary and he cannot do better than searching sequentially using a stopping rule.

The equilibrium concept used is the perfect bayesian equilibrium: every firm sets its price and

[2]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 situation, which certainly deserve a proper analysis

[3]One could imagine a richer strategy space regarding the set of keywords. As a matter of fact, a richer strategy space would not destroy the equilibrium of this simpler game, even though there might be other equilibria using more complex strategies

advertising policies so as to maximize its profit given the other firms' strategies and the stopping rule used by consumers. The stopping rule is itself a best-response to firms' strategies.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.

I will focus on symmetric equilibria in pure strategies.

3 Equilibrium analysis

Optimal stopping rule

In equilibrium, when a consumer located at y clicks on a link, the expected utility he gets from this click if he buys is

$$\int_{y-D^*}^{y+D^*} \frac{v(x)-p^*}{2D^*}dx = \int_0^{D^*} \frac{v(x)-p^*}{D^*}dx$$

Consumers regard each click as a random draw of a location x from a uniform distribution, whose support is $[y-D^*; y+D^*]$. Indeed a firm located at a distance greater than D^* from y 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 closer firm. For $R^* \equiv R(p^*, \sigma^*)$ to be a reservation distance it must be such that a consumer is indifferent between continuing to search and buying the product:

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

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

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

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.

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 $R(p, \sigma^*)$ is such that accepting a price p at a distance $R(p, \sigma^*)$ gives the same utility as accepting a price p^* at a distance R^*, i.e $v(R(p, \sigma^*)) - p = v(R^*) - p^*$. Thus we have the following proposition.

Proposition 1 *Given other firms' expected strategy $\sigma^* = (p^*, D^*)$, a consumer accepts to buy a good at price p if and only if the selling firm is located at a distance less than $R(p, \sigma^*)$, with $R(p, \sigma^*)$ such that*

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

where R^* is given by (2).

Moreover,

$$\frac{dR(p, \sigma^*)}{dp} = \frac{1}{v'(R(p, \sigma^*))} < 0$$

Optimal advertising and pricing strategies

Suppose that firm x sets a price p. Since it only has to pay for consumers who actually visit its link, firm x's optimal targeting strategy is to appear to every consumer y such that the expected profit made by x through a sale to y conditionally on y clicking on x's link is positive, i.e

$$Pr(y \text{ buys } x\text{'s product}|y \text{ clicks on } x\text{'s link}) \times (p - a) \geq 0 \tag{4}$$

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

The next lemmas will enable us to derive the only symmetric equilibrium. At this equilibrium, every firm chooses to advertise only to the consumers who buy the product as soon as they click on its link. Thus no consumer visits more than one firm.

The first lemma gives a necessary condition satisfied by any symmetric equilibrium.

Lemma 1 *Any symmetric profile of strategy $\sigma = (p, D)$ such that $D \neq R(p, \sigma)$ cannot be an equilibrium.*

Proof: This proof is in two stages: (1) if firms set $D < R(p,\sigma)$, then a firm can profitably deviate by targeting more consumers (2) if $D > R(p,\sigma)$, there is always at least one firm which can profitably deviate and lower its targeting distance.

1. The first stage is rather straightforward: suppose that all firms have a targeting distance D smaller than $R(p,\sigma)$. Take a consumer y and a firm x such that $D < d(x,y) < R(p,\sigma)$. If x were to deviate and choose to appear to consumer y, then it would sell the good with probability one if y clicked on its link. Thus it would be a profitable deviation.

2. Now suppose that all firms set $D > R(p,\sigma)$. Take a consumer y, and denote $\overline{x}$ the firm which is located farthest away from him. Since $d(\overline{x},y) > R(p,\sigma)$, the probability that y buys from $\overline{x}$ is zero. By reducing its reach, firm $\overline{x}$ can improve its profit. $\square$

Therefore, if a symmetric equilibrium exists, it must be the case that firms choose a targeting distance equal to consumers' equilibrium reservation distance. The next step in order to derive a symmetric equilibrium of the game is to study the best response of a firm when other firms play a symmetric strategy $\sigma^* = (p^*, D^*)$ with $D^* = R(p^*, \sigma^*)$.

Lemma 2 *Let x be the location of a given firm on the circle. If:*

- *all the other firms play the strategy $\sigma^* = (p^*, D^*)$ where $D^* = R(p^*, \sigma^*)$, and*

- *consumers expect all firms to play $\sigma^* = (p^*, D^*)$ and thus play $R(p, \sigma^*)$,*

then, whatever price p that firm x decides to set , the optimal advertising strategy is to set $D(p) = R(p, \sigma^*)$, i.e. a targeting distance equal to the reservation distance of consumers who face an "out of equilibrium" price.

This lemma states that if a firm wants to deviate from a situation where all firms set a targeting distance equal to the "equilibrium" reservation distance, the deviation implies to set a scope of relevance equal to the "out of equilibrium" reservation distance. Thus, the deviation does not change the number of clicks per consumer, since they find it optimal to buy from the first firm they visit. The proof is very similar to the previous lemma's one, and is omitted.

Now we can state an existence theorem and provide sufficient conditions for uniqueness. Notice first that there always exists a "trivial"' equilibrium, in which firms set $D^* = 0$ and $p^* = v(0)$, 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.

Proposition 2 *Suppose that $v(.)$ and s are such that $R(p, p, 1/2) \leq 1/2$, for any p. Then there exists a non trivial equilibrium of the game.*

Sufficient conditions for the equilibrium to be unique is that (C1): $x \mapsto v'(x) + xv''(x)$ satisfies the single crossing property and (C2): $\forall x \in [0; 1/2], xv''(x) + 2v'(x) \leq 0$.

The proof is provided in the appendix.

4 Comparative statics

In order to better capture some ongoing effects, in this section I focus a functional form for consumers' utility function $v(.)$:

$$v(d) = v - td^b \tag{5}$$

with $v, t > 0$, and $b \in (0; 1)$.[4] This specification gives the possibility to obtain closed-form expressions and to study the effect transportation costs t, while still allowing different curvatures of $v(.)$.

The proofs of the following propositions are in the appendix.

Proposition 3 *Let $\sigma^* = (p^*, D^*)$. When the utility of consumers is $u(d, p) = v - td^b - p$, the following results obtain:*

The optimal reservation distance is

$$R(p, \sigma^*) = \left(\frac{b+1}{b}\frac{sD^*}{t}\right)^{\frac{1}{b+1}} + \frac{p^* - p}{t}$$

The equilibrium advertising strategy, such that $D^* = R(p^*, p^*, D^*) = R(p^*, \sigma^*)$ is

$$D^* = \left(\frac{b+1}{b}\frac{s}{t}\right)^{\frac{1}{b}}$$

The equilibrium price is given by

$$p^* = \left(\frac{b+1}{b}\right)^{\frac{1}{b}} t^{\frac{b-1}{b}} s^{\frac{1}{b}} + a$$

The proof of the proposition is given in the appendix. Thanks to Proposition 3, we can obtain some straightforward comparative statics results.

[4] $b \in (0; 1)$ is sufficient for (C1) and (C2) to hold, and thus the non-trivial equilibrium is unique)

Proposition 4 *In equilibrium, the advertising strategy is an increasing function of the search costs s and a decreasing function of the transportation cost parameter t.*

To better understand this proposition, it may be useful to emphasize that firms choose the advertising strategy such that the marginal consumer (located at a distance D^*) is indifferent between buying and searching again. But as s increases, the marginal consumer is located further away, and therefore firms expand their scope. A rise in t has the opposite effect on the marginal consumer.

Proposition 5 *The equilibrium price is an increasing function of the search cost s. If $v(.)$ is convex (i.e. $b < 1$) the price is a decreasing function of t, whereas if $v(.)$ is concave ($b > 1$) the price is a decreasing function of t.*

The difference in the sign of the effect of t on the price comes from two opposite effects. On the one hand, *an increase in t raises the average valuation of targeted consumers.* Indeed, each firm targets fewer consumers, and they are closer than with a lower t. This exerts an upward pressure on the price. On the other hand, a rise in t leads every other firm to reduce its level of advertising, and thus consumers have *better outside options*, because they know that if they do not accept the current offer, the next offer is likely to be a similar product. This effect tends to reduce the price. The relative magnitude of these two opposite effects depends on the curvature of $v(.)$: the "average valuation effect" is stronger than the "outside option effect" when $v(.)$ is concave, and it is weaker when $v(.)$ is convex.

Proposition 6 *The equilibrium price is an increasing function of the advertising fee a, whereas the advertising strategy does not depend on it.*

As we will see below, the fact that the price depends on a comes from the fact that firms choose their advertising strategy, and thus internalize the cost of advertising. This may have implications in terms of tarification from the search engine point of view.

This specification of $v(.)$ makes it possible to measure the social value of the targeting technology. A useful benchmark is provided by Wolinsky (1983). In that model, firms cannot target consumers, but everything else is unchanged.

Proposition 7 *Giving firms the ability to target consumers unambiguously improves total welfare. Moreover, the final price is lower when firms target consumers as long as a is below some threshold $\overline{a}$.*

The welfare improving property of the targeting technology is due to a reduction of social inefficiencies caused by search costs and imperfect consumer-to-product matching. When firms target consumers in the model, consumers search only once, and the average distance between a seller and a buyer is smaller than without targeting.

5 Optimal design of the matching process

So far in the model the search engine is neutral, in the sense that it does not influence the amount of information which is revealed in equilibrium. But in practice there is evidence that search engines pay a lot of attention to the way advertisements are displayed. The ranking of ads through a "quality score" illustrates this concern, as well as the use of a "broad match" technology aimed at matching consumers to announcers when the keywords do not correspond exactly but are close enough according to some metrics. 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 additionnal noise.

In this section I will argue that a profit maximizing search engine has an incentive not to let firms target consumers as they wish. By introducing some additional level of noise in the process, the search engine may alleviate price competition between announcers, thus extracting more profit from them.

Noiseless matching leads to a Diamond paradox

In order to understand the consequences of changing the level of noise in the matching process, let us start with a natural setup, in which the search engine reveals firms' positions on the circle. I still assume that the search engine cannot observe the price set by firms, and therefore consumers ignore the price prior to clicking.

Proposition 8 *When the search engine reveals firms' positions, a variant of the Diamond paradox obtains: the only equilibrium is such that firms hold-up consumers and set a very high price.*

Proof: Suppose that consumers expect firms to set a price p^*. We need to find which firms a given consumer $y \in [0; 1]$ will visit , as well as his stopping rule. Since he anticipates that all firms set the same price, he strictly prefers to visit the firm which is the perfect match for him, i.e $x = y$.

Now, if firm x's price is $p \leq p^*$, he stops searching and buys. But if $p > p^*$, he faces a trade-off between buying at a high price (p) and paying a search cost in order to buy at a lower price (p^*) from a slightly less satisfying firm (from his point of view). Since there is a continuum of firms, the difference in positions between two firms can be made arbitrarily small, and thus the consumer buys the product at price $p > p^*$ if and only if $p \leq p^* + s$.

We recognize the classical hold-up problem (see Diamond (1971)): knowing how consumers behave, the only symmetric equilibrium is such that $p^* = v(0)$. Indeed, suppose that $p^* < v(0)$ is the price set by all firms. Then any firm can profitably deviate by setting a price equal to $p^* + s$, since at that price the consumers who visit the firm buy from it.□

This equilibrium is thus such that firms get all the surplus from trade. But then it is not individually rational for a consumer to start searching, because he will incur the search cost s and get zero surplus. Therefore the market collapses.

This proposition sheds light on a potential drawback of revealing too much non-price information to consumers, namely that firms could benefit from a hold-up situation *vis-à-vis* consumers and that trade could be hampered to some extent. This result is very similar to a result in Bakos (1997) and the intuition is also present in Anderson and Renault (2000), although in a different set-up.

Equilibrium as a function of the accuracy of the process.

The situation above corresponds to a case in which the search engine chooses to impose $D = 0$ to firms, that is a case in which all noise has been removed from the sampling process. Let us look at the equilibrium when the search engine is able to choose the level of noise in the sampling process, i.e to choose arbitrary values for D.

This technology might be regarded as an approximation of the "broad match" technology which is used by Google to match queries and advertisements. Basically, with broad match, the search engine will display an advertisement even if the keyword has not been selected by the announcer, provided it is regarded as relevant by the search engine. For instance, suppose that an announcer selects only one keyword, namely "web hosting". If a consumer enters the keyword "web hosting company" or "webhost", then the announcer'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 announcers: they no longer have to spend time and resources finding exactly 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.

Suppose that the search engine chooses the advertising distance D, everything else being unchanged. Now firms' strategy consists only in setting a price.

Proposition 9 *Suppose that $v(.)$ is concave, so that $pR(p, p^*, D)$ is a strictly concave function of p.*

Recall that D^* is such that $D^* = R(p, p, D^*)$ for any p. Then, when the search engine sets a broad match distance equal to $\overline{D}$:

- *If $D < D^*$, the only symmetric equilibrium is such that firms act like constrained monopolies. Marginal consumers are left with no surplus from the consumption of the good. The price is a decreasing function of $\overline{D}$. Consumers visit at most one firm.*
- *If $D > D^*$, the only symmetric equilibrium is such that firms are constrained by their competitors. Marginal consumers are indifferent between buying and searching again, but enjoy some surplus. The average number of firms visited by consumers is strictly larger than one. The price is an increasing function of $\overline{D}$.*
- *If $D = D^*$, there are at least two equilibria, corresponding to the limit of the two previous cases.*

The proof of this proposition is in the appendix.

Intuitively, when $\overline{D}$ is small, each firm knows that it is sufficiently close to the consumers who visit it that none of them will want to search again. Firms act like monopolies. For intermediary values of $\overline{D}$, this virtual isolation disappears: some consumers are now willing to switch to another firm if the price is too high. There is now an "outside option constraint" exerted by competitors, which leads the price to drop : firms can no longer act as if they were monopolies, and the price is thus the competitive price[5]. But as $\overline{D}$ further increases, this outside option constraint becomes less stringent, because the average distance between a consumer and the next firm is larger, leading to a rise in the price.

[5]The competitive price is still above marginal cost, since competition is imperfect because of information frictions.

Maximizing the search engine profit

This analysis suggests that the design of the matching technology is not neutral, and thus a natural question to ask is whether a search engine finds it profitable to change the "natural" level of noise in the matching process. Suppose that the search engine may choose both the advertising fee a and the targeting distance D. As we saw above, firms do not internalize a when they do not choose the targeting distance, and thus the search engine may capture all their profit, by setting $aD = \pi$. Because it cannot charge consumers here (for exogenous reasons), the search engine will distort the mechanism so as to induce a higher price while making sure that consumers participate. The highest price would obtain if the search engine was to set $D = 0$, but then consumers would choose not to enter the market. The search engine's profit is proportionnal to the equilibrium price as long as consumers' participation constraint is not violated, and we therefore have the following proposition:

Proposition 10 *For values of D such that consumers are willing to participate, the search engine profit is quasiconvex in D. The profit maximizing value of D is thus either the largest or the lowest value of D compatible with consumers' participation.*

In the former case, the situation is inefficient because it implies relatively bad matches and high search costs. On the contrary, the latter case is socially efficient because it maximizes the value of trade and induces small search expenses.

Proof: Because of Proposition 9, we know that the price is decreasing if $D \leq D^*$, and increasing if $D \geq D^*$. Thus the price is quasiconvex in D, and so is the search engine profit.□

The last proposition gives sufficient conditions for the matching accuracy to be improved or deteriorated when the utility is given by $u(d,p) = v - td^b - p$.

Let D_{min} and D_{max} denote respectively the minimal and maximal values of D compatible with consumers' participation.

Proposition 11 *Assume that a consumers' utility is given by $u(d,p) = v - td^b - p$, with $b \in [0;2]$ so that there is a unique price equilibrium for any value of D.*

Then

- *If $b > 1$ the search engine's optimal policy is to set $D = D_{max}$, i.e to choose a low-quality matching.*

- *If $b < 1$ and $v - 2S \geq (\frac{b+1}{b}\frac{s}{2}t^b)^{1/(b+1)}$, the search engine's optimal policy is to set $D = D_{min}$, i.e to choose a high-quality matching.*

Proof: If $b \geq 1$, it is straightforward to check that $D_{min} \geq D^*$, so that consumers do not participate if $D \leq D^*$. Therefore the maximal profit is obtained by setting $D = D_{max}$.

If $b < 1$ and $v - 2S \geq (\frac{b+1}{b}\frac{s}{2}t^b)^{1/(b+1)}$, we have $\pi(D_{min}) > \pi(1/2) \geq \pi(D_{max})$, which proves the result.□

Proposition 11 sheds some light on what may push the search engine to add or remove some noise to the matching process. The goal of the search engine is to raise the price of the good as much as possible under the constraint that consumers participate. This can be done either by revealing precise information about consumers to firms (when D is small) or by making consumers' outside options less attractive. Now consider a drop in s. If $b < 1$, such a drop will make it more likely that the search engine uses a high-quality matching technology, because it is more costly to degrade the outside option.

6 Concluding remarks

Search engines allow intent-related targeted advertising, and this paper illustrates the potential efficiency gains generated from firms' ability to target consumers. An interesting effect is the fact that targeting improves consumers' outside options, and thus leads to a lower price. A profit maximizing search engine wants to soften price competition between firms in order to extract their profit. In some cases this implies maximizing the value of trade, because firms are able to capture a large part of consumers' surplus. In other instances, maximizing the price implies degrading the quality of the matching process in order to improve firms' bargaining power (through a worsening of consumers' outside option).

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.

——— (2000): "Consumer Information and Firm Pricing: Negative Externalities from Improved Information," *International Economic Review*, 41(3), 721–42.

——— (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 (2007): "Position Auctions with Consumer Search," Levine's Bibliography 122247000000001633, UCLA Department of Economics.

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

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.

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

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

Edelman, B., M. Ostrovsky, and M. Schwarz (2007): "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1), 242–259.

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.

Hagiu, A., and B. Jullien (2007): "Designing a Two-Sided Platform: When To Increase Search Costs?," IDEI Working Papers 473, Institut d'Économie Industrielle (IDEI), Toulouse.

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

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.

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

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.

——— (2007): "Position auctions," *International Journal of Industrial Organization*, 25(6), 1163–1178.

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

——— (1984): "Product Differentiation with Imperfect Information," *Review of Economic Studies*, 51(1), 53–61.

Proof of Proposition 2

The equilibrium is obtained through the following steps:

- Compute the equilibrium reservation distance R^* (which equals the equilibrium targeting distance D^*, from lemma 1). This amounts to finding a fixed point $D^* > 0$ such that $D^* = R(p^*, \sigma^*)$ (such a fixed point is unique when (C1) is satisfied.).
- Compute $R(p, \sigma^*)$ as given by Proposition 1.
- Find a firm's best-response $argmax[(p-a)R(p,\sigma^*)]$ (which is a singleton if (C2) holds).
- The equilibrium price is such that $p^* \in argmax[(p-a)R(p,\sigma^*)]$.

The four steps in the proposition enable us to find an equilibrium when it exists. In order to show that an equilibrium exists it is sufficient to show that there exist fixed points D^* such that $D^* = R(p^*, \sigma^*)$ and $p^* \in argmax[(p-a)R(p,\sigma^*)]$.

For notational simplicity define $R(D) \equiv R(p^*, p^*, D)$. Using (2) and the implicit functions theorem on the open interval $(0; 1/2)$, we get $\frac{\partial R(D^*)}{\partial D} = -\frac{s}{R(D^*)v'(R(D^*))}$. As D^* decreases to zero, the right-derivative tends to $+\infty$, because $\lim_{D\to 0^+} R(D) = 0$ and $v'(.)$ is bounded and negative. Moreover, $R(1/2) \leq 1/2$ (by assumption), and therefore there must be a $D^* \in (0; 1/2)$ such that $D^* = R(D^*)$. Such a D^* is unique if $R(.)$ does not have more than one inflexion point. Differentiating $R(D)$ a second time with respect to D, one gets

$$\frac{\partial^2 R(D)}{\partial D^2} = sR'(D)[v'(R(D)) + R(D)v''(R(D))][R(D)v'(R(D))]^{-2} \tag{6}$$

Using the fact that $\lim_{D\to 0^+} \frac{\partial R(D)}{\partial D} = +\infty$, a sufficient condition for $R(.)$ to have at most one inflexion point is that $x \mapsto v'(x) + xv''(x)$ satisfies the single crossing property (i.e crosses the origin at most once, and from below). In that case, one can see that $R(D)$ is above D when $D < D^*$, and below D otherwise.

Once the existence of D^* has been proven, the existence of an equilibrium price results from Brouwer's fixed point theorem. The uniqueness of the equilibrium price obtains if we show that the profit is strictly quasi-concave in the firm's price. A sufficient condition for that is that $1/R(p, \sigma^*)$ is convex in p. For notational convenience let us drop the arguments in $R(p, \sigma^*)$. From Proposition 1 and the implicit functions theorem, one gets $\frac{\partial R}{\partial p} = \frac{1}{v'(R)}$. Straightforward computations show that $1/R(p, \sigma^*)$ is convex in p if and only if $-2v'(R) \geq Rv''(R)$. □

Proof of Proposition 3

Equation (2) writes

$$\int_0^{R^*} t\left(R^{*b} - x^b\right) dx = sD^*$$

that is $R^* = \left(\frac{b+1}{b}\frac{sD^*}{t}\right)^{1/(1+b)}$.

Using lemma 1, the equilibrium advertising distance is $D^* = R^* = \left(\frac{b+1}{b}\frac{s}{t}\right)^{1/b}$.

The optimal stopping rule is thus $R(p, p^*, D^*) = R^* + \frac{p^*-p}{t}$.

Maximizing $(p-a)R(p,p^*,D^*)$ with respect to p, and solving $p=p^*$ gives us the equilibrium price.□

Proof of Proposition 4

We have

$$\frac{\partial D^*}{\partial s} = (b+1)(\frac{(b+1)s}{bt})^{1/b-1} > 0$$

and

$$\frac{\partial D^*}{\partial t} = -\frac{1}{s(b+1)}(\frac{(b+1)s}{bt})^{1/b+1} < 0$$

Proof of Proposition 6

We have

$$\frac{\partial p^*}{\partial s} = \left(\frac{b+1}{b}\right)^{\frac{1}{b}} t^{(b-1)/b}\frac{1}{b}s^{(1-b)/b} > 0$$

and

$$\frac{\partial p^*}{\partial t} = \left(s\frac{b+1}{b}\right)^{\frac{1}{b}} \frac{b-1}{b}t^{(b-1)/b-1} > 0 \Longleftrightarrow b > 1$$

Proof of Proposition 7

Without targeting, finding the equilibrium amounts to solving the game with $D = 1/2$ for every firm. Note however that there is a difference in the profit function, which now writes $\pi(p) = pR(p,p^*,1/2) - \frac{1}{2}a\frac{\mu_c}{\mu_F}$. Indeed, because firms do not choose the consumers who receive the advertisement, advertising expenses are an exogenous fixed cost. Therefore it will not affect the equilibrium price, for a given mass of firms.

The reservation distance of consumers is now $R_1^* = \left(\frac{b+1}{b}\frac{s}{2t}\right)^{1/(1+b)}$, which is larger than R^*. The average distance between a seller and a buyer equals $R_1^*/2$, and is therefore larger than in the model with targeting. Moreover, because $R(p^*,p^*,1/2) \leq 1/2$, on average consumers search more than once, which creates some additionnal inefficiency with respect to the model with targeting. Finally the price is $p_1^* = \left(\frac{b+1}{b}\right)^{1/(b+1)} t^{b/(b+1)} s^{1/(b+1)} \left(\frac{1}{2}\right)^{1/(b+1)}$. It is straightforward to check that this price is higher than p^* when a tends to zero, and lower than p^* for high values of a.□

Proof of Proposition 9

First, let us chack that $\pi : p \mapsto pR(p,p^*,D)$ is indeed strictly concave when $v(.)$ is concave. By differentiating twice, and using Proposition 1, one gets

$$\pi''(p) = \frac{\partial v'(R(p,p^*,D)) - pv''(R(p,p^*,D))/v'(R(p,p^*,D)))}{v'(R(p,p^*,D))^2} + \frac{1}{v'(R(p,p^*,D))}$$

It is straightforward to check that $v''(.) \leq 0 \Longrightarrow \pi''(.) < 0$.

I will show that the nature of the symmetric equilibrium which is reached depends on the relative position of D and D^* (where D^* is the solution to $D^* = R(p,p,D^*)$, for all p).

- $D < D^*$

Suppose that D is such that $D < D^*$. As was shown in Proposition 2, $D \mapsto R(p,p,D)$ crosses $D \mapsto D$ only once, at $D = D^*$, and from above, this implies that $R(p,p,D) > D$ for every p. Thus if there is an equilibrium price p^*, it verifies $R(p^*,p^*,D) > D$, meaning all the consumers who visit a firm strictly prefer buying than searching again. Firms are somehow isolated from competition, and the unique potential equilibrium price is therefore

$$p^* = max\left(p^m, v(\overline{D})\right) \tag{7}$$

where p^m is the monopoly price.

To see why 7 is the only equilibrium price, suppose that there is only one firm on the market. The optimal price is the monopoly price p^m. Now, suppose that the firm cannot reach consumers farther away than D. If $v(D) \geq p^m$, the firm's optimal price is $v(D)$ (profit is strictly concave). If, on the other hand, $v(D) < p^m$, the optimal price is p^m.

Now let us check that this is indeed an equilibrium.

Suppose that all firms set a price p^* as defined in 7. If a firm decides to lower its price, this does not bring new consumers because D is fixed and all consumers closer than D would buy at price p^* (because $R(p^*,p^*,D) > D$). Thus it is an unprofitable deviation. On the other hand, if a firm decides to raise its price with respect to p^*, it is also an unprofitable deviation: because profit is strictly concave, it is decreasing for $p \geq p^m$ (which is the case: $p > p^* \geq p^m$).

From (7) one sees that the equilibrium price is a non-increasing function of D, and this completes the first part of the proposition.

- $D > D^*$

Now suppose that $D > D^*$. By a similar argument to previously, this implies that $R(p,p,D) < D$ for any p. If p^* is a symmetric equilibrium, it must maximize $pR(p,p^*,D)$ because $R(p,p^*,D)$ is the relevant demand. Indeed, start from p^* maximizing $pR(p,p^*,D)$. The only possibly profitable deviation would involve reducing the price and sell to every consumer on the segment of length $2\overline{D}$. But once the price is p such that $R(p,p^*,D) = D$, it is not profitable to further lower the price. Thus the profit which is maximized is $pR(p,p^*,D)$, and we know that there exists a symmetric equilibrium in this case, from proposition 2.

Let $\pi(p,p^*,D) = pR(p,p^*,D)$. Let us show that π is supermodular in (p,D) if v is concave.

$$\frac{\partial \pi(p,p^*,D)}{\partial p} = R(p,p^*,D) + p\frac{\partial R(p,p^*,D)}{\partial p} = R(p,p^*,D) + \frac{p}{v'(R(p,p^*,D))}$$

Then

$$\frac{\partial^2 \pi(p,p^*,D)}{\partial p \partial D} = \frac{R(p,p^*,D)}{\partial D} - p\frac{R(p,p^*,D)}{\partial D}\frac{v''(R(p,p^*,D))}{v'(R(p,p^*,D)^2}$$

Since $\frac{\partial R(p,p^*,D)}{\partial D} > 0$, we see that the profit is supermodular in (p,D) if v is concave.

Supermodularity of π in (p, D) implies that the equilibrium price is an increasing function of D, by Topkis's theorem.