

Advertising on a search engine: how to target consumers?

Alexandre de Cornière*

March 30, 2009

Abstract

Firms' ability to target consumers has increased thanks notably to the recent development of search engines. This paper studies the effects of targeted advertising through a search engine when products are differentiated, product information is imperfect and consumers incur search costs. In this framework, it is shown that targeting benefits consumers through three channels: better product-to-consumer matching, less search expenses and lower prices. The paper considers alternative advertising technologies, and identifies a potential source of market failure, when the search engine reveals too much non-price information.

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

JEL Classification: D43, D83, L13,M37.

*PSE (Paris-Jourdan Sciences Economiques / Paris School of Economics), 48 boulevard Jourdan, 75014 Paris, France. Email: alexandredecorniere@yahoo.fr. Phone: (+33) (0)1 43 13 63

1 Introduction

In 2007, online advertising expenses amount to 21 billion dollars in the United States, which is about 7% of total advertising expenses (see Evans (2008,[10])). 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 John Battelle (2005,[6])). 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 “cheap flights Australia” 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. As this example illustrates, intent-related advertising goes hand in hand with the well-known “long tail” phenomenon of the worldwide web, since it enables advertising for products whose consumer base is too small to make a non-targeted campaign.

The other ingredient, the “per click” pricing, is aimed at ensuring announcers that their investments are not wasted, i.e that the consumers for who they pay are those who actually see the ad *and* were looking for it¹.

In this paper I present a model of targeted advertising with differentiated products which includes the main features mentioned above. Firms are uniformly distributed around a circle, and consumers cannot observe prices nor positions on the circle. The search engine is an intermediary between firms and consumers: announcers choose which keywords they want to target, while consumers enter keywords and then search sequentially (and costly) at random through the links that appear. The per-click cost is exogenous.

Basically, I find that consumers benefit from firms’ ability to target through three

¹The per-click cost is determined through a *Generalized Second Price Auction* (See Edelman, Ostrovsky and Schwarz (2007,[8]), Varian (2007,[14]) for the properties of this mechanism)

channels: better matches, fewer 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 targeting 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, this raises the relative quality of the outside option, 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 raise in the quality of the latter implies less bargaining power for the firms and thus a lower price for the final good.

The matching technology is an approximation of how search engines really proceed. Google, for instance, does not display ads at random, but rather uses a weighted average of the firms’ bids and of their “quality score”, an index which is supposed to reflect the relevance and quality of an announcer. It seems interesting to study the effects of such a policy, and I therefore turn to this issue in the end of the paper. By focusing on a polar case, I find that if advertisements were always sorted by relevance to the query, the Diamond paradox would apply, i.e the resulting equilibrium price would be equal to the reservation value of consumers. This results underlines the importance of the question of the optimal level of information disclosed by the search engine.

- *Related literature*

This paper is related to the large literature on search models and advertising, as well as to more recent contributions which study internet search engines. The literature on search models on a product market has provided important insights. In a seminal paper, Diamond (1971)[7] 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.

When products are differentiated, the price is an increasing function of the search cost and entry is generally excessive (Anderson and Renault (1999,[1]), Wolinsky (1984,[16])).

The relationship between advertising and consumer search is not a new topic: in a framework with an homogenous good, Varian (1980,[13]) and Robert and Stahl (1993,[12]) see advertising as a substitute to costly information acquisition by consumers. Advertising generates a differentiation between informed and uninformed consumers, which leads firms to price discriminate. In a monopoly framework with uncertainty regarding the product's characteristics, Anderson and Renault (2006,[3]) 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,[9]) show that in a monopoly framework, firms' ability to target consumers reduces both consumers' and total surplus, and also reduces the price.

Iyer, Soberman and Villas-Boa (2005,[11]) 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,[17]) 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,[5]) 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 (2008,[4]) 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.

Section 2 presents the basic model, and section 3 shows the main results. These results are analyzed and compared to a benchmark without targeting in section 4. The revelation of products’ characteristics is analyzed in section 5.

2 The model

- *Description of the market and of preferences*

The model is a modified version of Wolinsky (1983,[15]).

Consider a market where a continuum 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. 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 one 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 gets

from consuming a good located at a distance d from his favorite brand is

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

where t is a transportation cost and p is the price of the good.

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.

- *Advertising technology*

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 (more precisely, they reveal their ideal brand). Consumers cannot enter several keywords at the same time². If a certain keyword is entered by a consumer, all the firms who want to target this keyword appear on the consumer's screen. Consumers do not observe neither the prices nor the positions on the circle of the firms before they visit them (they click on the link). If they click on a firm's link, they incur a search cost $s \in (0; t/4)$. 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. 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.³ After a consumer has sampled a firm and learned its price and position, he can come back at no cost (recall is costless).

- *Strategies and equilibrium concept*

²allowing this behavior might produce interesting results

³This advertising technology would maximize the profit of a monopolist for intermediate values of the search cost, for it exactly corresponds to the "threshold revelation" used by Anderson and Renault (2005).

A strategy for a firm x consists in the choice of a price p and in a set of keywords S . I only examine the case in which firms select closed symmetric intervals $S = [x - D; x + D]$. Nevertheless this is not a restriction on the strategy space : if all the other firms choose a closed symmetric interval, then, as we will see below, it is a best response for firm x to choose a closed symmetric interval.

As for consumers, their strategy consists in choosing an optimal stopping rule. This stopping rule consists in setting a reservation distance R , which depends on the (expected and observed) prices, such that the consumer is indifferent between buying a product at a distance R and continuing searching.

The equilibrium concept used is the Nash equilibrium: every firm sets its price and 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. In the remaining of the paper I focus on pure strategy symmetric equilibria.

3 Equilibrium analysis

Let (p^*, D^*, R^*) be the equilibrium strategies. Consumers' and firms's strategies need to be a best-response to this strategy profile.

- *Optimal stopping rule*

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

$$\int_{y-D^*}^{y+D^*} (v - td(x, y) - p^*)f(x)dx = 2 \int_0^{D^*} \frac{v - t|z| - p^*}{2D^*} dz$$

Each click is regarded by the consumer 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^* to be a reservation distance it must be such that a consumer is indifferent between continuing to search and

buying the product:

$$2 \int_0^{R^*} \frac{t(R^* - |z|)}{2D^*} dz = s \quad (2)$$

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. After a simple calculation one gets

$$R^* = \sqrt{\frac{2D^*s}{t}} \quad (3)$$

One may notice that the equilibrium reservation distance is independent from the equilibrium price. It 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 very precise targeting. Now, when a consumer samples a firm which has set an out-of-equilibrium price $p \neq p^*$, his optimal stopping rule $R(p)$ is such that accepting a price p at a distance $R(p)$ gives the same utility as accepting a price p^* at a distance R^* , i.e $v - tR(p) - p = v - tR^* - p^*$. Thus we have the following proposition.

Proposition 1 *A consumer accepts to buy the good at price p if and only if the firm is located at a distance less than $R(p)$, with $R(p)$ such that*

$$R(p) = R^* + \frac{p^* - p}{t} = \sqrt{\frac{2D^*s}{t}} + \frac{p^* - p}{t}$$

- *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 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

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

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

The next lemmas will enable us to derive an 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. We introduce the following notation : $R(p, D)$ is the reservation distance if consumers expect firms to set a price p and a targeting distance D .

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

Proof: This proof is in two stages: (1) if firms set $D < R(p, D)$, then a firm can profitably deviate targeting more consumers (2) if $D > R(p, D)$, 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 $R(p, D)$. Take a consumer y and a firm x such that $D < d(x, y) < R(p, D)$. 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 $D > R(p, D)$. Take a consumer y , and denote \bar{x} the firm which is located farthest away from him. Since $d(\bar{x}, y) > R(p, D)$, the probability that y buys from \bar{x} is zero. By reducing its reach, firm \bar{x} can improve its profit.

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 (p^*, D^*) with $D^* = R(p^*, D^*)$.

Lemma 2 *Suppose that a firm x anticipates that*

- *all the other firms play the strategy (p^*, D^*) where $D^* = R(p^*, D^*)$.*
- *consumers expect all firms to play (p^*, D^*) and thus play $R(p^*, D^*) = \sqrt{\frac{2sD^*}{t}}$.*

Then, whatever price p firm x decides to set, the optimal advertising strategy is to set $D(p) = R(p^, D^*) + \frac{p^* - p}{t}$, 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 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 left to the reader.

Corollary 1 *In equilibrium, consumers always buy to the first firm they visit.*

Thanks to Lemma 2, it is straightforward to compute the optimal strategy of a firm. Given that the other firms play $D^* = \frac{2s}{t}$ (which is obtained by solving $D^* = R(p^*, D^*)$), and given that $D = R(p^*, D^*) + \frac{p^* - p}{t} = \frac{2s + p^* - p}{t}$, firm x ’s profit is proportional to $2D(p)(p - a)$ that is

$$\pi(p) \propto 2(p - a) \frac{2s + p^* - p}{t}$$

Notice here that a plays the role of a marginal cost: since consumers buy at their first visit, each firm pays a exactly the same number of times as it sells the product. Firm x ’s best response to the equilibrium strategy is therefore $p^{BR}(p^*) = \frac{2s + p^* + a}{2}$. For p^* to be an equilibrium, it must be the case that $p^* = p^{BR}(p^*)$, i.e

$$p^* = 2s + a \tag{5}$$

The equilibrium strategies are summarized below:

Proposition 2 *There exists a unique symmetric equilibrium in pure-strategy.*

- *Firms set a price equal to $p^* = 2s + a$*

- They target all the keywords located at a distance less than or equal to $D^* = \frac{2s}{t}$
- Consumers buy whenever they find a firm at a distance less than or equal to $R^* = \frac{2s}{t}$

4 Comments on the equilibrium

- *Some comparative statics*

Some conclusions of the preceding analysis deserve particular attention. Regarding the level of advertising, as measured by the equilibrium reach of advertising D^* , we see that it is an increasing function of the search cost s and a decreasing function of the transportation cost t . This is in line with the intuitive signification that one may give to these parameters. Indeed, s and t are both a source of market power for the firms, but of a different nature. The search cost s strenghten a firm's bargaining position with respect to *all the consumers who have just clicked on its link*. Such consumers have an outside option, which is to continue searching, but it costs them s to use it. A rise in s *uniformly* makes the outside option more expensive and thus improve the firm's position. *Uniformly* means that this phenomenon applies to every consumer who has clicked, now matter how far he is from the firm.

On the other hand, a rise in the transportation cost t does not affect all the firm-consumer relationships the same way. Intuitively, if a consumer is close from the firm which makes him the offer, a rise in t implies that the consumer pays more attention to the distance between him and the firm, and thus he is more likely to accept, other things being equal, *in particular the value of the outside option*. But if the distance between the consumer and the firm is greater, a rise in t makes the consumer more reluctant to buy, other things being equal. Thus we see that a rise in t improves the firms' bargaining power *vis-à-vis* close consumers.

Having said that, it is easy to see why s and t have opposite effects on the equilibrium advertising level. A rise in s makes distant consumer more willing to buy, and thus the firm wants to target them, and inversely for a rise in t .

The above reasoning does not explain why the transportation cost t does not have any effect on the price level $p^* = 2s + a$.

Basically, a rise in s makes the offer more attractive to consumers who click on it, and, since announcers expand their reach, makes the outside option less valuable. Both effects improve firms' bargaining power, and thus lead to higher prices.

A rise in t has a more ambiguous effect: it improves the bargaining power of the firm vis-à-vis close consumers, deteriorates bargaining power vis-à-vis far consumers, and the firm advertises less. This effect tends to be positive. But there is another, more subtle, effect: a rise in t causes other firms to reduce their reach, and therefore *improves* the outside option of *all* the consumers, because they expect that the next firm they visit is at a distance smaller than if t was lower. In the model with linear transportation costs, these two effects offset each other and therefore firms do not benefit from more differentiation.

Regarding the reservation distance $R^* = \frac{2s}{t}$, the effects are roughly the same as above. The reservation distance raises with s , as the outside option is less valuable. It is a decreasing function of t . To understand clearly this point, it is useful to write R^* as a function of D^* the equilibrium targeting strategy. We have $R(D^*) = \sqrt{\frac{2sD^*}{t}}$, that is, even if the targeting strategy of the other firms was independent of t , R^* would be a decreasing function of t . But other firms' behavior depends on t ($D = \frac{2s}{t}$) and this phenomenon makes the reservation distance a less convex function of t .

- *A useful benchmark*

In order to correctly assess the impact of targeted advertising on market outcomes, it is useful to compare the results obtained above with results which would obtain if firms did not have the ability to target consumers.

This may be done by using Wolinsky 1983's model. The only difference ⁴ between that model and the model with targeting is that, in Wolinsky's model, each consumer receives all the ads. That model is therefore a benchmark which tends to underestimate the positive effects of targeting for consumers. Indeed, if one prevents targeting, the best thing for consumers is to receive all the ads and search sequentially.

⁴actually Wolinsky deals with a finite but large number of firms and uses approximations, but this does not change the results

In the linear version of Wolinsky's model, consumers' reservation distance writes $R_W = \sqrt{\frac{s}{t}}$, and the equilibrium price is $p_W = \sqrt{st}$. The average number of visits per consumer is $1/R_W = \sqrt{\frac{t}{s}} \geq 2$. The average distance between a buyer and a seller is $R_W/2$, and thus the average consumer utility is

$$u_W = v - t \times \frac{\sqrt{\frac{s}{t}}}{2} - s \times \sqrt{\frac{t}{s}} - \sqrt{st} = v - \frac{5}{2}\sqrt{st}$$

How do these findings compare to the model with targeting? To facilitate the comparison, let the advertising cost a tend to zero.

First, the price with targeting is lower, since $2s \leq \sqrt{st} \iff s/t \leq 1/4$, which is true (otherwise $\frac{2s}{t}$ would be larger than $1/2$, and this would not make much sense in the model with targeting). This result is in stark contradiction to the results obtained by Iyer, Soberman and Villas-Boa (2005) and Esteban, Gil and Hernandez (2001). In Iyer, Soberman and Villas-Boa (2005)'s model, targeted advertising enables firms to differentiate: consumers with strong preferences for one product are not targeted by the other firm, and therefore firms are in local monopoly. In Esteban, Gil and Hernandez (2001), the monopolist faces a less elastic demand with targeting and is therefore able to raise its price. In this paper, these effects are offset by an improvement of the outside option of the consumer, and therefore the elasticity of demand is raised, which leads to a lower price.

Consumers' reservation distance is higher without targeting, due to the low value of the outside option: if a consumer refuses an offer, the next offer he receives is a random draw uniformly distributed around the circle, instead of a random draw from an interval around his position. A direct consequence is that the average distance (which equals $R/2$ in both models) is also lower with targeting, which implies that targeting improves efficiency on the ground of better matching.

Targeting also reduces the number of visits before a purchase. Indeed, although the reservation distance is higher without targeting, it is still smaller than $1/2$, which implies that some consumers will receive offers that they do not accept in equilibrium.

One may also see that the differentiation parameter t has a positive effect on the price: $p_W = \sqrt{st}$. In light of the previous comments on the effects of t , the reason is simple: a rise in t does not affect the value of the outside option, because the

offers are drawn from the same distribution. Thus the positive differentiation effect on the mark-up is not offset by the “outside option effect”.

5 Should ads be sorted by relevance?

In this section I turn to the question of the amount of information revealed by the search engine. In the basic model, no “hard” information is revealed to consumer regarding firms’ positions on the circle. In equilibrium consumers anticipate correctly that firms are somehow close to them, but they have no other information.

The actual system is a bit different, in the sense that ads are sorted on the screen of a consumer. The sorting of ads is done by the search engine, on the basis of the announcers’ bids (which is ignored in this paper, see Edelman, Ostrovsky and Schwarz 2007 , Varian 2007 or Athey and Ellison 2008 for more on this) and of a so-called “quality score”. The way the quality score of an announcer is computed is unclear. Google, for instance, only gives some of the factors that are used to compute it⁵:historical clickthrough-rate (which measures the number of clicks generated by ads from a given announcer), relevance to the query, quality of the landing site, among other factors.

How would the conclusions of the model be affected by the introduction of a quality score? To see this this, I deal with a very simple proxy for the quality score, namely the position on the circle: the search engine reveals firms’ positions on the circle, so that consumers can choose which one to visit. Nevertheless I still assume that the search engine cannot observe the price set by firms. To simplify matters even more, I restrict firms’ strategy space to the set of prices: firms do not choose which keywords to target (this does not change the symmetric equilibrium outcome which is studied below). On the other hand, consumers’ strategy space now includes the *choice* of firms to visit.

In this slightly modified framework, there is an equilibrium in which firms hold-up consumers and set a very high price. This equilibrium is the only symmetric equilibrium in which firms use pure-strategies.

To see this, suppose that consumers expect firms to set a price p^* . We need to find

⁵This information may be found at <http://adwords.google.com/support/bin/answer.py?answer=10215>

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 $p^* = v$. Indeed, suppose that $p^* < v$ 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, as the reader may have anticipated, this 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!

Although a bit extreme, this conclusion sheds light on a potential difficulty, namely that firms could benefit from a hold-up situation *vis-à-vis* consumers and that trade could be hampered to some extent. Revealing too much information to consumers can be damaging as long as this information is price-irrelevant. This intuition is also present in Anderson and Renault (2000,[2]), although in a different set-up.

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 affects consumers' outside options, and thus leads to a lower price. Nevertheless, one has to be cautious about the amount of information which is revealed, for firms may be able to hold-up consumers, resulting in a market collapse.

The next step would be to model the search engine as a strategic player of the game, or to adopt a mechanism design approach in order to study the optimality of the advertising technology.

References

- [1] Simon P. Anderson and Regis Renault. Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *RAND Journal of Economics*, 30(4):719–735, Winter 1999.
- [2] Simon P. Anderson and Regis Renault. Consumer information and firm pricing: Negative externalities from improved information. *International Economic Review*, 41(3):721–42, August 2000.
- [3] Simon P. Anderson and Regis Renault. Advertising content. *American Economic Review*, 96(1):93–113, March 2006.
- [4] Mark Armstrong, John Vickers, and Jidong Zhou. Prominence and consumer search. Technical report, 2008.
- [5] Susan Athey and Glenn Ellison. Position auctions with consumer search. Levine’s Bibliography 122247000000001633, UCLA Department of Economics, October 2007.
- [6] John Batelle. *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Portfolio, 2005.
- [7] Peter A. Diamond. A model of price adjustment. *Journal of Economic Theory*, 3(2):156–168, June 1971.
- [8] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.
- [9] Lola Esteban, Agustin Gil, and Jose M Hernandez. Informative advertising and optimal targeting in a monopoly. *Journal of Industrial Economics*, 49(2):161–80, June 2001.

- [10] David S. Evans. The economics of the online advertising industry. *Review of Network Economics*, 7(3):359–391, September 2008.
- [11] Ganesh Iyer, David Soberman, and J. Miguel Villas Boas. The targeting of advertising. *Marketing Science*, 24(3):461–476, Summer 2005.
- [12] Jacques Robert and Dale O. Stahl. Informative price advertising in a sequential search model. *Econometrica*, 61(3):657–86, May 1993.
- [13] Hal R. Varian. A model of sales. *American Economic Review*, 70(4):651–59, September 1980.
- [14] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.
- [15] Asher Wolinsky. Retail trade concentration due to consumers’ imperfect information. *Bell Journal of Economics*, 14(1):275–282, Spring 1983.
- [16] Asher Wolinsky. Product differentiation with imperfect information. *Review of Economic Studies*, 51(1):53–61, January 1984.
- [17] Timothy Van Zandt. Information overload in a network of targeted communication. *RAND Journal of Economics*, 35(3):542–560, Autumn 2004.