Search Advertising

Alexandre de Cornière *

November 3, 2015

Abstract

Search engines enable advertisers to target consumers based on the query they have entered. In a framework in which consumers search sequentially after having entered a query, I show that such targeting reduces search costs, improves matches and intensifies price competition. However, a profit-maximizing monopolistic search engine imposes a distortion by charging too high an advertising fee, which may negate the benefits of targeting. The search engine also has incentives to provide a suboptimal quality of sponsored links. Competition among search engines can increase or decrease welfare, depending on the extent of multi-homing by advertisers.

Keywords: search engine, targeted advertising, consumer search.

JEL Classification: D43, D83, L13, M37.

1 Introduction

Search advertising designates the display of "sponsored links" on a search engine results page, alongside "organic links". Whereas organic links are free, sponsored links are the main source of revenue for search engines. The standard pricing scheme in the industry is *per-click pricing*: search engines collect fees from advertisers every time a consumer clicks on their link.

Unlike more traditional advertising formats where ads are displayed alongside content (TV, newspapers), search advertising reaches consumers at a point at which they are actively looking for information about products, and is therefore less of a nuisance. This is all the more so that, as advertisers

*Toulouse School of Economics. Email: adecorniere@gmail.com. I thank the editor and three anonymous referees for helping me improve the paper. I am also grateful to Simon Anderson, Mark Armstrong, Heski Bar-Isaac, Luc Bridet, Bernard Caillaud, Olivier Compte, Jacques Crémer, Gianni De Fraja, Gabrielle Demange, Renaud Foucart, Gilles Grandjean, Bruno Jullien, Frederic Koessler, Dina Mayzlin, Romain de Nijs, Marco Ottaviani, Régis Renault, Joana Resende, Andy Skrzypacz, Greg Taylor and Xavier Wauthy for very useful suggestions. Previous versions of this paper have circulated under the title "Targeting with consumer search: an economic analysis of keyword advertising".

can select precise keywords to target, sponsored links are generally relevant to the queries and thus valuable to consumers.

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

To answer these questions, I present (sections 2 and 3) a model of targeted advertising through a search engine, with differentiated products, which includes the main features mentioned above. Firms are horizontally differentiated *à la* Salop (1979), and consumers do not have prior knowledge of firms' prices or products' characteristics. The search engine is an intermediary between firms and consumers: firms choose which keywords they want to target, and consumers enter keywords and then search sequentially at random through the links that appear. Firms incur a fixed cost to be registered on the search engine, and they pay the latter on a per-click basis. In this model, I do not consider organic links (see for instance de Cornière and Taylor (2014) for a model with both organic and sponsored links).

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 *composition* effect[1]). 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 about the quality of the match between a firm and a keyword, they will most likely try to use it so as to optimally design the matching mechanism. For instance, Google sorts firms using a weighted average of the firms' bids and of a "quality score" index. Consumers are also sometimes provided with additional information on the results page, such as a map showing the locations of different vendors. On the other hand, the "broad match" technology enables search engines to expand the set of keywords corresponding to a given advertisement. In section 4, I study a situation in which the search engine can more finely design

[1] I thank a referee for suggesting this terminology.

the matching mechanism. 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, when the matching is too accurate, consumers view firms as almost homogenous. In that case the only equilibrium is such that consumers expect high prices and do not search (the Diamond paradox). It can even be optimal for the search engine to implement a matching that is less accurate than the laissez-faire outcome (in which accuracy is the result of equilibrium behavior by firms and consumers). The reason is that offering a noisy matching mechanism makes consumers more willing to accept high-prices on the product market, because it is now more costly to refuse an offer and to search again, as the next firm is less likely to be a good match (another instance of composition effect). 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.

In section 5, I build upon my baseline model to incorporate the issue of competition between search engines. I show that there may exist equilibria in which competition is socially harmful, but also equilibria in which competition is desirable. A key factor for competition to be desirable is the extent to which advertisers multi-home, which in turns depends on the magnitude of economies of scale in advertising. Full multi-homing makes advertising fees irrelevant as a competitive tool, and competing search engine thus behave like monopolists facing a low elasticity of demand, which lowers welfare compared to the monopoly case. When advertisers single-home, competition leads to lower advertising prices, and therefore improves welfare.

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, in particular in relation to its impact on product market pricing. Most of the models in this literature discuss mechanisms through which targeting tends to increase prices. In Roy (2000), Galeotti and Moraga-Gonzalez (2008) or Iyer, Soberman, and Villas-Boas (2005), targeting enables firms to segment the market, thereby relaxing price competition. In Esteban, Gil, and Hernandez (2001), competition between advertisers is shut down, and targeting, by making it optimal to advertise only on specialized outlets that cater to high willingness-to-pay consumers, leads to a price increase. In a framework with several advertisers who compete in an auction for an advertising slot on a platform, de Cornière and de Nijs (Forthcoming) show that targeted advertising (i.e allowing firms to condition their bids on consumers' characteristics) changes the expected composition of demand for each firm, with more weight on consumers with less elastic demands. Firms then have an incentive to charge higher prices.

In contrast, as discussed above, the present paper shows that targeted advertising, when coupled with consumer search, intensifies price competition. Grossman and Shapiro (1984) also find that targeting can lower prices, although through a reduction in advertising costs.

Other recent works on targeted advertising include Van Zandt (2004) who shows that targeted advertising can lead to information overload, Johnson (2013), who examines ad avoidance behavior, or Bergemann and Bonatti (2011) and Athey and Gans (2010), who study competition between medias with different targeting technologies. These papers ignore the issue of product market pricing.

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)'s circular city model. Wolinsky (1986) and Anderson and Renault (1999) also deal with heterogenous preferences, modeling match values as i.i.d shocks.[2]

Some models of consumer search are more directly relevant to the search engine industry. Athey and Ellison (2011) focus on the design of the auction to allocate advertisement slots, given that consumers search strategically through the slots. However their analysis does not include competition between firms on the product market. Armstrong, Vickers, and Zhou (2009) deal with price competition between firms, in a model in which one firm is made prominent, meaning that although consumers search strategically, they always visit the prominent firm first. Chen and He (2011) and Haan and Moraga-Gonzalez (2011) endogenize prominence by including an advertising stage prior to firms' pricing decision and consumer search. In Chen and He (2011) this advertising stage is an auction in which the more relevant firms submit higher bids, making it rational for consumers to sample them first. Haan and Moraga-Gonzalez (2011) assume that consumers are boundedly rational, in the sense that the probability that a consumer remembers a firm is proportional to that firm's advertising expenses. Yang (2013) looks at the impact of improvements in the search technology on product design decisions. None of these papers study the strategic choice of keywords by advertisers, nor the role of the search engine.

The model captures the complementarity between search and advertising that is inherent to the search advertising technology, in the sense that the firm must target a keyword searched for by a consumer for a match to be possible (both have to be active). This is unlike Robert and

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

Stahl (1993) for instance, where receiving an advertisement dispenses the consumer from searching, but reminiscent of Anderson and Renault (2006), whose equilibrium shares some features with this paper.[3]

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. Other papers have a similar approach: Baye and Morgan (2001) model an intermediary who acts as an information gatekeeper on a homogenous product market, and look at the optimal two-sided pricing, taking into account subsequent price setting by firms and consumer search. Hagiu and Jullien (2011) focus on the design of a platform in terms of search diversion, and highlight several reasons why an intermediary does not want to provide the highest quality matching, even when the technology is costless. Eliaz and Spiegler (2011), in a related paper, also show that a search engine wants to implement a matching with a suboptimal quality. White (2013) and Taylor (2013) 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).Gomes (2014) characterizes the optimal mechanism to sell an advertising slot when consumers and advertisers are heterogenous.

2 The model

2.1 Description of the market and of preferences

The framework is based on Wolinsky (1983). Consider a market in which there is a continuum of products uniformly distributed along a circle whose perimeter is normalized to one. Each product can be described by a keyword. For each product, there is a continuum of firms that are potential entrants.[4] When a firm enters the market, its type, i.e the keyword that perfectly describes its product, is denoted $\theta \in [0, 1]$. θ is private information.[5]

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 in their willingness to pay for their favorite product. More specifically, for each product ω, there is a continuum of mass 1 of consumers whose willingness to pay v is distributed on $[0, \overline{v}]$ according to a continuous and increasing cumulative distribution function F , with a log concave density f.[6]

[3]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).

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

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

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

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 the distance between them $d(\theta, \omega) = d$, is

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

where p is the price of the good and ϕ is a mismatch cost. I assume that ϕ is increasing, and convex, which implies that consumers are risk-averse with respect to the quality of the match. $\phi(d)$ is often referred to as a transportation cost in traditional models of spatial competition. Here, I 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 launches an online advertising campaign using the search engine incurs a fixed cost C. This cost corresponds to the marketing or monitoring expenses that accompany the advertising campaign, and is not a payment to the search engine.

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

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

2.3 Strategies and equilibrium concept

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

will ensure that the search engine's profit is quasi-concave in the advertising fee (see Caplin and Nalebuff (1991)). This property is satisfied by many usual distributions (see Bagnoli and Bergstrom (2005)).

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

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. The mass of active firms is μ. Entrants incur the fixed cost C. A firm θ that decides to use the search engine chooses a price p_θ and an advertising strategy D_θ. Consumers do not observe firms' strategies.

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 involve cutoff rules. First, let $EU(v)$ be the expected utility of a consumer of type v if he uses the search engine. If he does not search, his utility is normalized to zero. Let $v^*(a)$ (sometimes noted v^*) be such that $EU(v^*(a)) = 0$. Consumers with $v \geq v^*(a)$ use the search engine, while consumers with $v < v^*(a)$ 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 consumer continues to search if and only if $v - \phi(\hat{d}) - \hat{p} < U_R$. The strategy of a consumer thus consists in the choice of the reservation utility U_R, or, alternatively, in the choice of a reservation distance $R \equiv \phi^{-1}(v - \hat{p} - U_R)$. R depends on the expected future prices and locations if the consumer keeps on searching. Because a consumer who starts searching eventually buys a product (and because of separability), v does not affect R (but it does affect the decision to participate). Figure 1 illustrates how the market works.

The equilibrium concept used is perfect Bayesian equilibrium with passive beliefs. 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 number of entrants is such that there is no profit for advertisers in equilibrium.

The stopping rule R^* is a best-response to firms' strategies. I focus on symmetric equilibria in pure strategies $(a^*, R^*, v^*, p^*, D^*, \mu^*)$. The reservation distance R^* depends on the price p that the

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

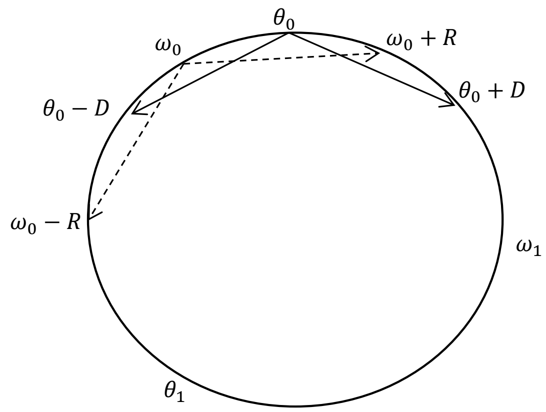

Figure 1: Targeting strategy and stopping rule. A firm located in θ_0 targets all the keywords in $[\theta_0 - D, \theta_0 + D]$ (clockwise), and a consumer located in ω_0 stops searching as soon as he samples a firm in $[\omega_0 - R, \omega_0 + R]$. Here, the ω_0-consumer may buy from the θ_0-firm, but would not accept the offer by the θ_1-firm. The ω_1-consumer cannot see an ad by the θ_0-firm, because he is not in its targeted set of keywords.

consumer is facing, but also on the price and targeting distance he expects other firms to set. I use the notation $R^*(p, p^*, D^*)$ where (p^*, D^*) refer to what consumers expect other firms to play.

The following assumption ensures existence of a symmetric 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, $D = 1/2$) in a symmetric equilibrium, some consumers search more than once before buying. In particular, this assumption requires search costs not to be too large. It is 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).

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, advertising and entry. Finally, given the equilibrium of the subgame, we can find the search engine's optimal per click fee a.

3.1 Consumer search

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

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

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

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

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

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

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

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

Proof: From (2), $R(p, p, D)$ is the solution to $\int_0^R \frac{\phi(R) - \phi(x)}{D} dx = s$, and hence 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 (and thus the firm's demand) $R(p, p^*, D^*)$ is such that accepting a price p at a distance $R(p, p^*, D^*)$ gives the same utility as accepting a price p^* at a distance R^*, i.e $v - \phi(R(p, p^*, D^*)) - p = v - \phi(R^*) - p^*$. Thus we have the following result:

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

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

where R^* is given by (2).

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

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

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

3.2 Advertisers' strategy

Advertising Now that we know consumers' search behavior, it is possible to characterize firms' optimal targeting strategy. It turns out that this optimal strategy is surprisingly simple: a firm should target a consumer if and only if the distance between the two is smaller than the reservation distance. Indeed, suppose that firm θ sets a price p. Since it only has to pay for consumers who actually visit its link, firm θ's optimal targeting strategy is to target 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. With a continuum of firms, consumers' stopping rule is stationary, and a consumer never comes back to a firm he previously visited. The conditional probability is then either 0 (when $d(\omega, \theta) > R(p, p^*, D^*)$) or 1 (when $d(\omega, \theta) \leq R(p, p^*, D^*)$). Thus we have the following result, the proof of which is in the appendix:

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

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

Pricing Thanks to Lemma 3, it is straightforward to find the per-(search engine)-user profit function of a firm if the other firms and consumers play their respective equilibrium strategies.[9] Indeed, if that firm wants to set a price p different from the candidate equilibrium price p^*, it must also change the set of consumers that it targets. By the same argument as in Lemma 3, the optimal advertising strategy is to target consumers if and only if they are located at a distance smaller than the new reservation distance $R(p, p^*, D^*)$. Since every consumer within this reservation distance is targeted by a mass $2R(p^*, p^*, D^*)\mu^*$ of firms,[10] the per-user demand for the firm's product is $\frac{R(p,p^*,D^*)}{R(p^*,p^*,D^*)\mu^*}$. Conditional on visiting the firm, all consumers buy without searching further, and this implies that a is formally equivalent to the firm's marginal cost of production. Therefore, if all the other players (firms and consumers) follow the equilibrium strategy profile, a firm's per-user profit function is

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

The previous reasoning 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 do not participate 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.

Entry Recall that $v^*(a)$ is the lowest value of v such that a consumer participates. Given the profit function (6), the free-entry condition writes:[11]

$$(p^*(a) - a)\frac{1 - F(v^*(a))}{\mu^*(a)} = C \tag{7}$$

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

$$s = \int_0^{R^*} \frac{\phi(R^*) - \phi(x)}{D^*} dx \tag{8}$$

$$R^* = D^* \tag{9}$$

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

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

[11] I adopt the convention that if a mass μ of symmetric firms serve a mass λ of consumers, each firm sells to λ/μ consumers.

$$p^*(a) - a = \phi'(R^*)R^* \tag{10}$$

$$v^*(a) = a + \phi'(R^*)R^* + \phi(R^*) \tag{11}$$

$$\mu^*(a) = \phi'(R^*)R^* \frac{(1 - F(v^*(a)))}{C} \tag{12}$$

Proof: The proof of the existence and uniqueness is provided in the appendix. Equation (8) is simply a rewriting of equation (2), while (9) comes directly from Lemma 3. Equation (10) 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. To obtain (11), note that the expected surplus of a consumer is $v - p^*(a) - s - E[\phi(d)|d \leq R^*]$. Now, using (8) and (9), one can show that $s + E[\phi(d)|d \leq R^*] = \phi(R^*)$. Given (10), the indifferent consumer is thus such that $v - \phi'(R^*)R^* - a - \phi(R^*) = 0$. Finally, (12) is simply a rewriting of the free-entry condition.□

Equation (10) gives the mark-up in equilibrium. By convexity of ϕ 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 that has been selected by a firm, instead of a per-click basis. Indeed, in that case the per-user profit function of a firm would be $\pi(p, p^*, a) = \frac{1}{\mu}\left(p\frac{R(p,p^*,D^*)}{R(p^*,p^*,D^*)} - aR(p, p^*, D^*)\right)$, and one would just need to replace a by $aR(p^*, p^*, D^*)$ in the expression of the equilibrium price (10).[12]

3.3 Search engine 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.[13] Given that a higher fee a leads to a higher product price (see 10), it also leads to fewer consumers using the platform, as shown by (11).

Since in equilibrium every consumer who uses the search engine clicks only once, the search engine's profit is

$$\Pi^{SE}(a) = a\,(1 - F(v^*(a))) \tag{13}$$

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

An interior solution to the search engine's program is given in the next proposition:

[12] If firms paid a to the search engine on a per-sale basis, (8), (9) and (10) would still constitute an equilibrium, but there would also be an equilibrium without targeting (such that $D^* = 1/2$). See Taylor (2011) for a comparison of these payment schemes.

[13] In Appendix B I show that it would be equivalent for the search engine to choose a quantity of available slots. The price-setting model is more convenient to use.

Proposition 2 *The optimal fee for the search engine is*

$$a^* = \frac{1 - F(v^*(a^*))}{f(v^*(a^*))} \tag{14}$$

Proof: The first-order condition to (13) is $a^* = \frac{1-F(v^*(a^*))}{v^{*\prime}(a^*)f(v^*(a^*))}$. But according to (11), $v^{*\prime}(a^*) = 1$.□

The optimal fee for the search engine is greater than the socially optimal fee, which is zero. Indeed, looking at (8), (9) and (10), one sees that a has no impact on the quality of the matching in equilibrium, while a high a implies a higher price paid by consumers.[14] However, the search engine would not make any profit if this was the case. In order to increase its profit, the search engine imposes a distortion, because the higher fee results in a higher equilibrium price.

3.4 The effects of targeting

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

The main results are the following:

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

1. *reduces the expected number of clicks;*

2. *reduces the mismatch frictions;*

3. *has an ambiguous effect on the price of the final good.*

Proof: In the case of no-targeting, the equilibrium reservation distance for consumers, R_{NT}, is given by (8) with $D_{NT} = 1/2$. The first point of the proposition is a direct corollary of Lemma 3 and of Assumption 1.

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

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

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

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

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

$$p_{NT} = \phi'(R_{NT})R_{NT} \tag{15}$$

Comparing this latter expression with (10), one gets

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

Indeed, we have $R_T < R_{NT}$ and, because ϕ is convex, the function $x \mapsto \phi'(x)x$ is non-decreasing, which explains why the second term in the right-hand side is negative. □

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

The intuition for the composition effect is that targeting, by affecting the composition of the pool of firms from which consumers sample, increases the continuation value of search for consumers and thus the (semi-) elasticity of demand. This in turn puts more pressure on firms to reduce their mark-up.

Example. Suppose that $\phi(d) = td$, and that $F(v) = 1 - e^{-\eta v}$. Then we get $p_T = 2s + a^*$, $p_{NT} = \sqrt{st}$ and $a^* = \frac{1}{\eta}$. Targeting leads to a price reduction if and only if $\frac{1}{\eta} + (2s - \sqrt{st}) < 0$. Numerical results show that it is possible for welfare to be higher or lower with targeting. Welfare is more likely to be higher with targeting when t is high, s is intermediate and η is high. (See Appendix C for details.)

3.5 Robustness

For the sake of tractability, I have made several simplifying assumptions, the two main being (i) that consumers have identical search costs, and (ii) that they search randomly accross the links they see. In this subsection I show that the main insights of the model do not depend critically on these two assumptions.

Heterogenous search costs The assumption of homogenous search costs is the driving force behind Lemma 3, and relaxing it implies that different consumers have a different reservation distance. It follows that the result that consumers buy from the first firm they visit no longer holds.

More importantly, introducing heterogenous search costs makes firms' advertising strategies more responsive to the advertising fee a, which is certainly a more realsitic feature.

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

To study the equilibrium in a setup with heterogenous search cost, assume that consumers have the same valuation v for their ideal product,[16] that their utility is linear in the mismatch cost ($\phi(d) = td$) and that participation is exogenous. Let the search cost s be uniformly distributed over $[0, 1]$.

When the equilibrium price is p^* and when firms target up to a distance D^*, a consumer with search cost s, facing a price p has a reservation distance $R(s, p, p^*, D^*) = \sqrt{\frac{2sD^*}{t}} + \frac{p^*-p}{t}$. Therefore, if a consumer visits a firm at a distance x, he buys with probability $Pr[x \leq R(s, p, p^*, D^*)] = Pr[s \geq \frac{t}{2D^*}(x + \frac{p-p^*}{t})^2] = 1 - \frac{t}{2D^*}(x + \frac{p-p^*}{t})^2$.

A firm charging a price p is willing to target all the consumers located at a distance x such that $pPr[x \leq R(s, p, p^*, D^*)] \geq a$. The targeting distance is the value of x such that the previous inequality is binding: $D(a, p, p^*, D^*) = \frac{\sqrt{2D^*t}\sqrt{p(p-a)}}{p} + \frac{p^*-p}{t}$.

The profit of a firm is then

$$\pi(p, p^*, D^*) = \int_0^{D(a,p,p^*,D^*)} \left(pPr[s \geq \frac{t}{2D^*}(x + \frac{p-p^*}{t})^2] - a \right) dx$$

The symmetric equilibrium is the solution to $\frac{\partial \pi(p^*,p^*,D^*)}{\partial p} = 0$ and $D(p^*, p^*, D^*) = D^*$. The solution is $p^* = \frac{2+\sqrt{4+6a}}{3}$ and $D^* = \frac{4-\sqrt{4+6a}}{t}$.[17]

Two points merit comments. First, the equilibrium targeting distance is now a decreasing function of the fee a. Therefore, charging a higher fee would allow the search engine to implement a more accurate matching through self-selection, thereby increasing consumers' utility.

On the other hand, even though a drop in D^* following an increase in a exerts downward pressure on the equilibrium price through the composition effect, this effect is not enough to offset the direct pass-through effect of an increase in the fee: the equilibrium price is an increasing function of a.

The heterogenous search cost version of the model thus generates the same prediction as Eliaz and Spiegler (2011) regarding the effect of advertising fees on the quality of the matching, but the opposite prediction regarding its impact on equlibrium price.

Non-uniform sampling While in the model I assume that the links that consumers see are identical and therefore that consumers click randomly, in practice firms can convey some information regarding the products they offer, so that consumers may not be ex ante indifferent between links.

Formally, one can capture this effect by assuming that consumers are more likely to click on a firm corresponding to a better match. Below I show that the results of Proposition 3 continue to hold under such a specification.

Consider first the case without targeting (indexed by NT). Consumers are exposed to ads from firms located all over the unit circle, but are more likely to click a link the closer to their position

[16] Heterogeneity in v does not imply heterogenous behavior for consumers who decide to use the search engine.
[17] The differentiation parameter t has to be large enough so that $D^* \leq 1/2$.

the corresponding firm is: the probability that the consumer clicks on a firm located at a distance between x and $x+dx$ is $g(x)dx$, with g non-increasing.[18] Let $G(x)$ be the probability that the consumer clicks on a firm located at a distance closer than x. G is increasing and concave, with $G(1/2)=1$.

In equilibrium, consumers' reservation distance is given by

$$\int_0^{R_{NT}} g(x)(\phi(R_{NT})-\phi(x))dx = s$$

As before, if consumers face a price $p \neq p_{NT}$, they adjust their reservation distance to $R_{NT}(p)$ such that they are indifferent between buying at price p_{NT} and distance R_{NT} and buying at price p at s distance $R_{NT}(p)$

A firm's profit is then proportional to $\int_0^{R_{NT}(p)} pg(x)dx - aK$, where K is the expected number of clicks (and does not depend of a firm's price). The first-order condition is

$$p_{NT} = \phi'(R_{NT})\frac{G(R_{NT})}{g(R_{NT})}$$

Suppose now that firms can choose to target up to a distance D. Consumers are still more likely to click on a close firm, but, given that firms farther than D decide not to target a consumer, the probability with which he clicks on a firm closer than x is $\min\{\frac{G(x)}{G(D)},1\}$.

Because payment is on a per-click basis, Lemma 3 (and thus equation (9)) still holds, i.e. in equilibrium $D=R_T$ and Proposition 3 (1) holds.

The reservation distance is given by

$$\int_0^{R_T} \frac{g(x)}{G(R_T)}(\phi(R_T)-\phi(x))dx = s$$

As in (4), we have $\frac{\partial R_T(p)}{\partial p} = -\frac{1}{\phi'(R_T(p))}$.

The profit of a firm is proportional to $\int_0^{R_T(p)}(p-a)\frac{g(x)}{G(R_T)}dx$. The first-order condition for a symmetric equilibrium is thus

$$p_T - a = \phi'(R_T)\frac{G(R_T)}{g(R_T)}$$

Let us now compare the two situations. First, let $h(R) \equiv \int_0^R g(x)(\phi(R)-\phi(x))dx$. R_T is such that $h(R_T) = sG(R_T)$, while R_{NT} is such that $h(R_{NT}) = s$. Because h is increasing, we have $R_T < R_{NT}$, i.e. targeting reduces the expected mismatch, as in Proposition 3 (2).

To show that Proposition 3 (3) also holds, note first that the click fee is only passed through to consumers with targeting, for the same reason as before. Moreover, because $R \mapsto \phi'(R)\frac{G(R)}{g(R)}$ is increasing by convexity of ϕ and concavity of G, we have $\phi'(R_{NT})\frac{G(R_{NT})}{g(R_{NT})} > \phi'(R_T)\frac{G(R_T)}{g(R_T)}$, i.e. the

[18] When g is constant, sampling is uniform.

composition effect is still present.

Alternative ad pricing mechanism Another simplification I use in the model is that the search engine directly sets the per-click fee, rather than using an auction. In appendix B, I look at a mechanism where the search engine chooses a number of slots to display (technically, a mass of firms μ) , and where the per-click fee a is determined through an ascending auction. I then show that the equilibrium described in Proposition 1 is an equilibrium of this modified game, the only difference being that the "free" parameter is μ instead of a.

4 Platform design

The assumption that the search engine does not behave strategically with respect to information revelation leaves aside interesting theoretical as well as practical issues. 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, which is the default option on Google, the search engine might 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, a company that selects the keyword "hat" may appear following a query for "caps". 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 assume that the search engine can influence the relevance of ads by choosing the value of D, on top of chosing a per-click fee a. To simplify the exposition I assume in this section that firms' participation is exogenous, and I normalize μ to 1.

We saw in the previous section that the accuracy of targeting affects firms' market power (via their mark-up, equal to $R^*\phi'(R^*)$). The search engine faces a trade-off between giving firms enough market power and ensuring sufficient consumer participation. The main result of this section is that the optimal value of D^* is *always* at least as large as the equilibrium value obtained in section 3. The following lemma will be useful in proving that result.

Lemma 4 *If the search engine has the possibility to choose the accuracy of the matching, then (i) the equilibrium price no longer depends on a, and (ii) the search engine can entirely extract firms' profit.*

Proof: Let $v^*(D)$ be the consumer who is indifferent between using the search engine and his outside option 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 search only once, whereas otherwise they search on average $\frac{D}{R(p^*,p^*,D)}$ times. It is straightforward to see that the level of a does not affect which price a firm should charge. In equilibrium, by setting $a = p^*)/\max\{\frac{D}{R(p^*,p^*,D)}, 1\}$, the search engine extracts all the profit. □

Recall that D^* is the equilibrium distance in the game in which firms choose their targeting strategy.

Proposition 4 *The optimal matching accuracy, from the search engine's point of view, is* $D^{SE} \geq D^*$.

The complete proof of this proposition is in the appendix, but its logic is the following. If the search engine chooses $D < D^*$, there cannot be an equilibrium with consumer participation. Indeed, suppose that consumer participation is positive, that firms charge a price p, and let $v(p)$ be the willingness to pay of the marginal consumer (indifferent between using the search engine and staying out of the market). As I show in Lemma 7 in the appendix, $D < D^*$ implies that $D < R(p, p, D)$, i.e. that the consumers who use the search engine strictly prefer to buy from the first firm they visit rather than to search. This implies that firms' demand is inelastic around p, so that charging $p + \epsilon$ such that $R(p, p, D) > R(p + \epsilon, p, D) > D$ is a profitable deviation. Therefore, for p to be an equilibrium, it must be that no consumer participates (i.e. $v(p) > \overline{v}$). This is a variant of the well-known Diamond paradox (Diamond (1971)). A corollary of this observation is that D^* is the targeting accuracy that would be chosen by a benevolent planner unable to affect firms pricing.

When $D > D^*$, we have $R(p, p, D) < D$, so that some users find it optimal to visit several firms. This ensures that there exists an equilibrium price compatible with a positive participation. There are then two effects of increasing D: on the one hand, by the composition effect (discussed in section 3.4), the price-elasticity of each user's demand goes down, so that firms generate higher per-user profit. Such profit is then captured by the search engine. On the other hand, a higher D reduces the expected utility of users (through higher prices, less accurate matches, and higher expected search costs), such that participation declines. The search engine chooses the optimal D^{SE} so as to balance these two effects.

5 Competing search engines

In this section I come back to a setup of decentralized targeting, that is in which firms are free to choose their targeting strategy. Recall from Proposition 2 that a monopolistic search engine imposes a distortion on the economy through a per-click fee that is higher than the socially optimal fee (here, zero). The purpose of this section is to determine under which conditions, if any, can competition between search engines improve welfare.

To do so, I study a stylized game of competition in which a second search engine operates on the market. Consumers can use at most one search engine, whereas advertisers can be present on both platforms (multi-home).[19] The cost structure of firms is the following: it costs C_H to register on one search engine, and $C_H + C_L \in [C_H, 2C_H]$ to register on both search engines (to multi-home). That $C_H \geq C_L$ means that there may exist economies of scale. For instance, whereas registering on a search engine for the first time implies developing a website and devising an advertising strategy, many of these expenses need not be incurred when registering on another search engine. However, if monitoring the performance on a search engine is the main expense, it may be that economies of scale are not very important.

The timing is the following:

1. Both search engines choose their per-click fees a_1 and a_2.

2. Advertisers observe a_1 and a_2, make their participation decision, and choose their targeting (D) and pricing (p) strategies. Advertisers can have different targeting strategies across search engines, but are constrained to charge a uniform price.[20]

3. Consumers observe observe a_1 and a_2, choose a search engine (or none) and start a sequential search. If consumers are indifferent between the two search engines, search engine 1 receives a market share $n_1 \geq n_2$.

Because of the importance of coordination, multiple equilibria arise in this setup. Rather than characterising the whole set of equilibria, I focus on two kinds of equilibria: equilibria with full multi-homing and equilibria with full single-homing.[21] The reason for this shortcut is that, with partial multi-homing, all the firms on a given search engine may no longer be symmetric (some only use search engine i, while others multi-home), and the analysis of the model is much more delicate. Focusing on full multi-homing and full single-homing restores symmetry. Hopefully this

[19]See Ambrus, Calvano, and Reisinger (2014) or Athey, Calvano, and Gans (2013) for models with consumer multi-homing.

[20]The latter assumption is consistent with casual empiricism. If firms could price-discriminate between search engines, search engines would then compete by lowering their fees, à la Bertrand, which would lead to an efficient outcome.

[21]The former exists if $C_L = 0$, the latter if $C_L = C_H$ and there are enough potential firms. Propositions 5 and 6 focus on these polar cases.

will be enough to convey the point that the extent of multi-homing is a key driver of the desirability of competition between search engines.

Indeed, when advertisers multi-home, the Bertrand logic of price competition between search engines cannot apply: a decrease in the fee a_1 results in a decrease in the final price of the good on both search engines, and thus such a strategy does not increase the market share of search engine 1. On the other hand, when advertisers single home on search engine 1, reducing a_1 allows to attract consumers. Advertising fees are then driven down to zero.

The following intermediary result will prove useful in the subsequent analysis of equilibrium under competition. Consider a situation in which a share $\alpha_i \in \{0, n_i, 1\}$ of search engine users utilize search engine i, and in which there is no partial multi-homing.[22] By the same logic as that of Lemma 3 , we have:

Lemma 5 *If firms decide to advertise on search engine i, the equilibrium targeting strategy must satisfy $D_i = R(p_i, p_i, D_i)$, where p_i is the price charged by advertisers who use search engine i.*

Indeed, if we had $D_i > R(p_i, p_i, D_i)$, a firm could deviate by choosing a smaller D_i, whereas for $D_i < R(p_i, p_i, D_i)$ a profitable deviation would consist in targeting a larger set of keywords.

Multihoming equilibrium Intuitively, multi-homing is more likely to occur in equilibrium if there are large enough economies of scale. To see this, let's assume that $C_L = 0$ so that registering on a second search engine is costless.

Proposition 5 *If $C_L = 0$, there exists an equilibrium in which all active firms multi-home. In this equilibrium, the expected per-click fee $n_1 a_1^M + n_2 a_2^M$ is higher than the monopoly per-click fee a^*. Therefore welfare is lower than under monopoly.*

Proof: First, note that given that $C_L = 0$, single-homing is dominated by multi-homing. Now consider a situation in which all active firms advertise on both search engines. Lemma 5 implies that $D_1 = D_2 = D^* = R^* = R_1^* = R_2^*$, as given by (8) and (9). Given that firms charge the same price irrespective of the search engine used by consumers, consumers are then indifferent between the two search engines and market shares are thus n_1 and n_2. It should be clear that the situation is exactly the same as if there was a unique search engine charging a per-click fee of $n_1 a_1 + n_2 a_2$. Thus, as in Proposition 1, there is a unique equilibrium in the subgame. We have

$$p(a_1, a_2) = R^* \phi'(R^*) + n_1 a_1 + n_2 a_2 \tag{16}$$

The consumer type who is indifferent between using a search engine and his outside option is

$$v(a_1, a_2) = R^* \phi'(R^*) + n_1 a_1 + n_2 a_2 + \phi(R^*) \tag{17}$$

[22]This means that if either all firms multi-home or none does.

The mass of active firms, $\mu(a_1, a_2)$, is then given by

$$(p(a_1,a_2) - n_1 a_1 - n_2 a_2)(1 - F(v^*(n_1 a_1 + n_2 a_2))) = C_H \mu(a_1, a_2) \tag{18}$$

Given the above analysis, search engine i's profit maximization program is

$$\max_{a_i} a_i n_i \left(1 - F(v^*(n_1 a_1 + n_2 a_2))\right)$$

The first-order condition is

$$n_i a_i = \frac{1 - F(v^*(n_1 a_1 + n_2 a_2))}{f(v^*(n_1 a_1 + n_2 a_2))} \tag{19}$$

Let us now show that the system of first-order conditions has a unique solution, such that $n_1 a_1^M + n_2 a_2^M > a^*$. First, from assumption 3, f is log-concave. Theorem 2 in Bagnoli and Bergstrom (2005) thus ensures that $1 - F$ is also log-concave, which implies that $\frac{1-F}{f}$ is decreasing. To prove uniqueness, notice that (19) leads to $a_2 = \frac{n_1}{n_2} a_1$. Thus we can rewrite $v^*(n_1 a_1 + n_2 a_2) = v^*(2 n_1 a_1) = 2 n_1 a_1 + R^* \phi'(R^*) + \phi(R^*)$. Equation (19) for $i = 1$ thus rewrites $n_1 a_1 = \frac{1 - F(v^*(2 n_1 a_1))}{f(v^*(2 n_1 a_1))}$. The left-hand side is an increasing function of a_1 that covers the range of positive reals, and the right-hand side is a decreasing (and positive) function of a_1. Therefore there exists a unique a_1^M that satisfies the equation (and thus a unique a_2^M).

Note that since $n_1 a_1^M = n_2 a_2^M$, we also have $a_1^M \leq a_2^M$. Finally, in order to show that $n_1 a_1^M + n_2 a_2^M \geq a^*$, rewrite (19) as

$$\frac{n_1 a_1^M + n_2 a_2^M}{2} = \frac{1 - F(v^*(n_1 a_1^M + n_2 a_2^M)))}{f(v^*(n_1 a_1^M + n_2 a_2^M))}$$

and compare with equation (14):

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

The solution to the first equation ($n_1 a_1^M + n_2 a_2^M$) must be larger than the solution to the second (a^*). $\square$

With multi-homing, and under the assumption that firms cannot price-discriminate, each search engine behaves like a monopolist. However, the relevant demand for search engine i is $n_i(1 - F(v^*(n_1 a_1 + n_2 a_2^M)))$, and its elasticity is lower than that of the monopoly demand $1 - F(v^*(a))$, because an increase in a_i is passed through to consumers at a rate $n_i < 1$.[23]

Single-homing equilibrium The previous result relies on the fact that with multi-homing, search engines do not benefit relative to their competitors from having firms lower their prices. When firms

[23]This intuition is also present in Wright (2002), although in a different setup.

single-home, this logic no longer applies, as I show now. Suppose that there are no economies of scale, i.e that $C_L = C_H$,[24] and that the mass of potential firms is very large. Then we have the following result:

Proposition 6 *When $C_L = C_H$ and there are many potential entrants, there exists an equilibrium in which all entrants single-home. In this equilibrium, the advertising fees are $a_1^S = a_2^S = 0$, so that welfare is higher than with a monopolistic search engine.*

Proof: Consider the following strategy profile:

1. Search engines charge zero advertising fees: $a_1^S = a_2^S = 0$;

2. Targeting strategies and reservation distances are given by (8) and (9) on both search engines;

3. Firms on both search engines charge the same price $p_1 = p_2 = R^*\phi'(R^*)$;

4. The number of consumers on search engine i is $n_i(1 - F(v^*(0)))$;

5. All entrants single-home, and the mass of firms who advertise on search engine i is given by

$$2\mu_i C_H = n_i R^* \phi'(R^*)(1 - F(v^*(0)))$$

6. If $a_i > a_j$, consumers and firms behave as if search engine j was the only one on the market.

Given $a_1^S = a_2^S = 0$, points 2 to 5 clearly form an equilibrium. By 6, no search engine has an incentive to charge a higher fee.□

Discussion. This model of competition between search engines has several drawbacks. Given the potential multiplicity of equilibria, it is delicate to derive unambiguous results while comparing monopoly and duopoly. Moreover, the use of infinitely elastic market shares leads to very stark results that seem at odds with what one observes in practice. However, this simplistic model delivers an original insight, namely that the desirability of competition on the search engine market depends on the extent of multi-homing. Even though equilibria with partial multi-homing are difficult to study, it seems reasonnable to conjecture that, as economies of scale in advertising increase, more advertisers will multi-home, which relaxes competition between search engines.

An intriguing consequence of such a result is that efforts to foster advertisers multi-homing, such as the ones made by the European Commission[25] in its recent investigation on Google, might have adverse consequences by increasing the advertising fees. It may well be that these effects are of

[24]For $C_L \in (0, C_H)$ pure single-homing cannot be part of the equilibrium, since a firm who single-homes on search engine i would make a positive profit by also registering on j. As mentionned above, partial multi-homing is difficult to study, but pure single-homing delivers some insights that should, to a certain extent, carry over to the partial multi-homing case.

[25]See europa.eu, IP/10/1624

second order compared to the risks of exclusion of rival search engines through exclusivity clauses, but, given that the theory is still far from being established, these elements probably deserve further investigation.

On a different note, the results here differ from the standard results on multi-homing (Armstrong (2006)) due to the limited instruments that can be used. Indeed, if search engines can only use the fee (or the quantity of sponsored links), they are not able to extract profit from advertisers (the multi-homing side) while at the same time providing high surplus to consumers (the single-homing side), as would be the case with a richer set of instruments.

6 Concluding remarks

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

Although the model provides insights regarding the links between the design of the platform and market outcomes, it ignores some dimensions that are potentially important, such as the presence of organic links or and the issue of own-content bias, which was at the center of the EU's recent investigation against Google (see de Cornière and Taylor (2014) for instance). Future work will hopefully further improve our understanding of how these aspects interact with each other.

References

Ambrus, A., E. Calvano, and M. Reisinger (2014): "Either or Both Competition: A "Two-Sided" Theory of Advertising with Overlapping Viewership," *mimeo.*

Anderson, S. P., and R. Renault (1999): "Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model," *The 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," *The RAND Journal of Economics*, 37(3), 668–691.

Armstrong, M., J. Vickers, and J. Zhou (2009): "Prominence and consumer search," *The RAND Journal of Economics*, 40(2), 209–233.

Athey, S., E. Calvano, and J. Gans (2013): "The impact of the internet on advertising markets for news media," Discussion paper, National Bureau of Economic Research.

Athey, S., and G. Ellison (2011): "Position Auctions with Consumer Search," *Quarterly Journal of Economics*, 126.

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

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

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

Bar-Isaac, H., G. Caruana, and V. Cunat (2010): "Information Gathering and Marketing," *Journal of Economics & Management Strategy*, 19.

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

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

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

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

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

de Cornière, A., and R. de Nijs (Forthcoming): "Online Advertising and Privacy," *The RAND Journal of Economics*.

de Cornière, A., and G. Taylor (2014): "Integration and Search Engine Bias," *The RAND Journal of Economics*, 45, 576–597.

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

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

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

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

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

Gomes, R. (2014): "Optimal Auction Design in Two-Sided Markets," *The RAND Journal of Economics*, 45, 248–272.

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?," *The 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. (2013): "Targeted Advertising and Advertising Avoidance," *The RAND Journal of Economics*, 44(1), 128–144.

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

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

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

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

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

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

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

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

——— (2013): "Search Quality and Revenue Cannibalisation by Competing Search Engines," *Journal of Economics & Management Strategy*, 22(3), 445–467.

Van Zandt, T. (2004): "Information Overload in a Network of Targeted Communication," *The 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. (2013): "Search engines: Left side quality versus right side profits," *International Journal of Industrial Organization*, 31(6), 690–701.

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

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

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

Yang, H. (2013): "Targeted Search and the Long Tail Effect," *The RAND Journal of Economics*, 44, 733–756.

A Proofs

A.1 Proof of Lemma 3

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

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

Lemma 6 *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 Lemma 3. 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 6, and since $d(\omega,\theta) < R^*(p^*,p^*,D^*)$, we know that $P[v \geq p^* + \phi(d(\theta,\omega)) | v \geq v^*] = P[\delta(v,p^*) \geq d(\theta,\omega) | v \geq v^*] = 1$. Thus it would be a profitable deviation.

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

A.2 Proof of Proposition 1

The equilibrium is obtained through the following steps:

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

   Lemma 7 *Under assumption 1, and for any price p, the function $r : D \mapsto R(p,p,D)$ has two fixed points: 0 and $D^* \in (0,1/2)$.*

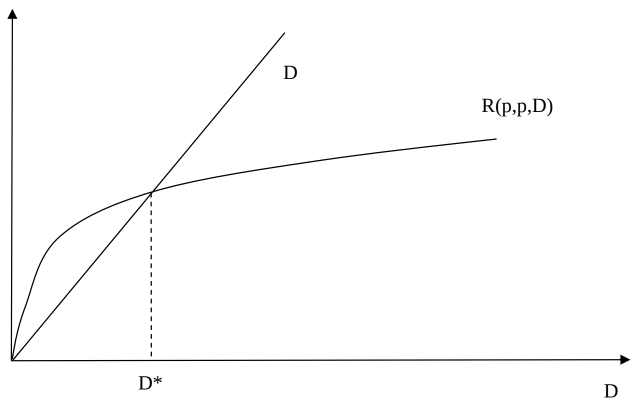

Figure 2: D versus $R(D)$

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

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

Using the implicit functions theorem on the open interval $(0, 1/2)$, we get $r'(D) = \frac{s}{r(D)\phi'(r(D))}$. As D goes to zero, $r'(D)$ tends to $+\infty$, because $\lim_{D\to 0} r(D) = 0$ and $\phi'(.)$ is bounded and positive.[26] 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

$$r''(D) = -sr'(D)[\phi'(r(D)) + r(D)\phi''(r(D))][r(D)\phi'(r(D))]^{-2} \tag{20}$$

By convexity of ϕ, 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 $(p - a)R(p, p^*, D^*) \times \frac{1-F(v^*(a))}{\mu^*(a)R(p^*,p^*,D^*)}$ if other firms play (p^*, D^*).

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

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

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

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

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

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

A.3 Proof of Proposition 4

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

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

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

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

[26]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.

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

$$\pi_\theta(p,p^*) = Dp\int_{v^*}^{\overline{v}}\int_0^{\overline{x}(v,p,p^*,D)} \frac{1}{D}f(v)dv = p\int_{v^*}^{\overline{v}} \overline{x}(v,p,p^*,D)f(v)dv \tag{21}$$

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 8 *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, (21) rewrites

$$\pi_\theta(p,p^*) = p\int_{v^*}^{\overline{v}} \min\left(D, R(p,p^*,D)\right) f(v)dv = p\min\left(D, R(p,p^*,D)\right)[1-F(v^*)] \tag{22}$$

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 9 *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 7) Therefore, at any symmetric strategy profile p, demand is inelastic around p. Each firm has an incentive to raise the price by ϵ, since such a deviation is not enough to trigger an additional search by consumers. □

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

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

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

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

By the envelope theorem,

$$\frac{\partial \pi_\theta^*(D)}{\partial D} = p^*(D)\frac{\partial R(p^*,p^*,D)}{\partial D}[1-F(v^*(D))] - v^{*'}(D)f(v^*(D))p^*(D)R(p^*(D),p^*(D),D) \tag{23}$$

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 10 *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

$$R(p(D), p(D), D) + p(D)\frac{\partial R}{\partial p}(p(D), p(D), D) = 0 \tag{24}$$

Given that $\frac{\partial R}{\partial p} = -\frac{\partial R}{\partial p(D)}$, totally differentiating (24) gives

$$\frac{dp(D)}{dD} = -\frac{\frac{\partial R}{\partial D}\left(1 + p(D)\phi''(R)(\phi'(R))^{-2}\right)}{\frac{\partial R}{\partial p}} \tag{25}$$

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

B Alternative pricing mechanism

In the main text I focus on the simplest mechanism possible, i.e in which the search engine selects a per-click fee. In order to check the robustness of the results, let us look at an auction-like mechanism which is still tractable.

Suppose that there is a mass μ of ad slots available on the search engine. The per-click fee a is determined through an ascending uniform auction. The timing of this modified game is the following

1. For each keyword θ, the fee a_θ starts at zero, and is continuously increased until a mass μ of firms remain. Let a^*_θ be the clearing fee.

2. Each firm who has won a slot for at least one keyword chooses a price p for its product. Active firms incur a cost C of monitoring the ad campaign on the search engine.

3. Consumers observe the fees and decide whether to start a sequential search with uniform sampling.

Given the importance of coordination by firms, the previous game may have many, potentially asymmetric, equilibria. Below I show that one equilibrium is closely connected to the equilibrium given in Propositions 1 and 2.

For each position θ on the circle, let us index firms by $i \in [0, 1]$.

When C is high enough, the following strategy profile is an equilibrium. For each θ each firm located in θ stays in the auction for all keywords $x \in [\theta - R^*, \theta + R^*]$ as long as the per-click fee a_x is lower than a^μ given by

$$R^*\phi'(R^*)\frac{(1 - F(v^*(a^\mu)))}{K} = C.\text{[27]}$$

When a_x reaches a^μ, all firms with $i > \mu$ drop out while those with $i \leq \mu$ remain. Firms do not bid for keywords further away than R^*.

Such a strategy profile leads to a per-click fee of a^μ. Given this per-click fee, the equilibrium price is given by equation (10): $p^*(a^\mu) = R^*\phi'(R^*) + a^\mu$, all consumers with $v \geq v^*(a^\mu)$ use the search engine, and each advertiser makes zero profit.

[27] This is the free entry condition (7).

Let us check that no deviation by advertisers is profitable in the auction stage. Suppose that a firm in θ decides to remain active at $a_x = a^\mu$ for $x > \theta + R^*$ (or $x < \theta - R^*$). In order to sell to these extra consumers, the firm has to lower its price (otherwise the consumers would continue searching after clicking on its link). But we saw in Proposition 1 that setting $p = p^*(a^\mu)$ is optimal when all other firms do the same. So the deviation cannot be profitable.

Moreover a^μ is the unique symmetric equilibrium per-click fee consistent with a mass μ of advertisers remaining active.[28] Indeed, suppose that the auction leads to a symmetric price $a < a^\mu$. Then some firms with $i > \mu$ could make a strictly positive profit by staying longer in the auction. With $a > a^\mu$, there would be too few users on the search engine to cover the fixed cost C with a mass μ of firms.

The main difference with the analysis in the main text is that the search engine can only affect the advertising price indirectly, by changing the number of slots μ. For instance, a way for the search engine to increase the advertising price is to reduce the number of firms allowed on the platform. Such a move would in turn lead to an increase in the product price. Note that the negative correlation between the mass of active firms and the equilibrium price of the goods is not the result of stronger competition on the product market, but rather of softer competition at the auction stage, leading to a lower per-click fee.

C Welfare effects of targeting

In order to assess whether targeting increases welfare or not, I use the following specification: $\phi(d) = td$, and $F(v) = 1 - e^{-\eta v}$.

Using results from section 3, I find that the equilibrium with targeting is given by :

$$a_T = \frac{1}{\eta}, \quad p_T = 2s + a_T, \quad R_T = \frac{2s}{t}, \quad v_T = 4s + \frac{1}{\eta}$$

And the equilibrium without targeting is:

$$a_{NT} = p_{NT}, \quad p_{NT} = \sqrt{st}, \quad R_{NT} = \sqrt{\frac{2s}{t}}, \quad v_{NT} = 2\sqrt{st}$$

Welfare with targeting is then given by

$$W_T = \int_{v_T}^{\infty} (v - v_T + a_T) dF(v) = \int_{v_T}^{\infty} (v - 4s) dF(v)$$

and welfare without targeting is given by

$$W_{NT} = \int_{v_{NT}}^{\infty} (v - v_{NT} + a_{NT}) dF(v) = \int_{v_{NT}}^{\infty} (v - \sqrt{st}) dF(v)$$

[28] Symmetric equilibrium meaning that all keywords sell for the same fee.

The following figures depict the welfare gain from targeting as a function of s, t and η . The default values in the figures are $s = 0.15, t = 2$ and $\eta = 2$.

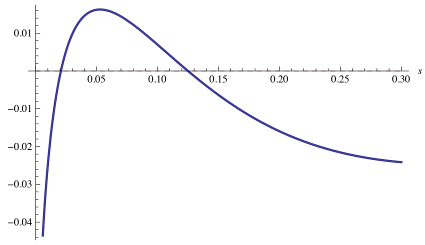

Figure 3: Welfare gain from targeting as a function of s

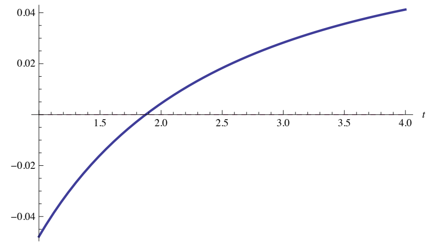

Figure 4: Welfare gain from targeting as a function of t

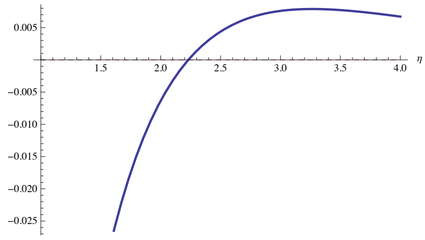

Figure 5: Welfare gain from targeting as a function of η

As expected, targeting is most valuable for high values of transportation costs (Figure 4) and of the elasticity of participation (Figure 5). Figure 3 shows that welfare gains from targeting are higher for intermediate values of the search costs. This is perhaps surprizing, but one should keep in mind that, although search costs are minimized thanks to targeting, firms' mark-up with targeting is increasing with s at a linear rate ($p_T = 2s + a_T$), whereas the rate is lower for high values of s without targeting ($p_{NT} = \sqrt{st}$).