A Simple Model of Search Engine Pricing[*]

Kfir Eliaz[†] and Ran Spiegler[‡]

December 5, 2009

Abstract

We present a simple model of how a monopolistic search engine optimally determines the average quality of firms in its search pool. In our model, there is a continuum of consumers, who use the search engine's pool, and there is a continuum of firms, whose entry to the pool is restricted by a price set by the search engine. We show that a monopolistic search engine may have an incentive to set a relatively low price that encouarges low-relevance advertisers to enter the search pool. This conclusion is independent of whether the search engine charges a price per click or a fixed access fee.

KEYWORDS: search engines, internet, two-sided markets, sequential search

1 Introduction

A search engine is a platform that serves a two-sided market. It is based on a technology that potentially improves the quality of consumer search. Before the advent of internet search engines, yellow pages were the closest example of a search engine. Firms pay in order to be included in the yellow pages, with various degrees of prominence. The yellow pages organize the set of firms according to some categorization system. In internet environments, consumers use search engines by submitting a query in a language dictated by the search engine. The objects that the query elicits depend on the search engine's method. In particular, a "sponsored links" system assigns objects to queries according to a mechanism in which firms pay the search engine for (prominent) appearance on the list of query results.

[†]Brown University. E-mail: Kfir_Eliaz@Brown.edu. URL: http://www.econ.brown.edu/fac/Kfir_Eliaz.
[‡]Tel Aviv University and University College London. E-mail: r.spiegler@ucl.ac.uk. URL: http://www.homepages.ucl.ac.uk/~uctprsp/

1

In the current age of Google, there is a near-monopoly in the industry of internet search engines. Our objective in this short paper is to present a simple, tractable model of sponsored-link pricing by a monopolistic search engine. Our model builds on a model of sequential consumer search due to Wolinsky (1986). We enrich Wolinsky's model by allowing for heterogeneity in the firms' degree of "relevance" for consumers, and by introducing a search engine that controls the search pool via its pricing decision.

Our main result is that the search engine may find it optimal to degrade the quality of the search pool by setting a low price-per-click that encourages low-relevance firms to enter. This leads to higher search costs and higher prices in the search pool. While it may come as no surprise that monopoly can generate an inefficient outcome, the distorting effects of monopoly in the case of search engine pricing are novel and therefore, worthy of separate enquiry. Here, a better pool of firms has a negative effect on the monopolist's profits because it leads to more competition among the firms, which, in turn, leads to lower prices and shorter searches (i.e., fewer "clicks").

Because we assume a large population of firms, our model allows us to abstract from auction-theoretic aspects and considerations of prominence (see our discussion in Section 5), in order to focus in the simplest manner possible on implications of search engine pricing for consumer search costs and product prices. As such, our paper complements existing theoretical work on search engines.

The closest paper to ours is Chen and He (2006), which develops a model of price competition with sequential consumer search, in which a finite number firms with different degrees of relevance bid for prominence. They show that the search engine may sometimes be better off with a lower quality pool in the sense that its revenue from the position auction has an inverted U-shape with respect to the highest level of relevance. However, unlike our model, the search engine in their framework cannot control the average quality (relevance) in its search pool. In addition, we assume a large population of firms, which enables us to get more mileage in the analysis of the search engine's problem.

Ellison and Athey (2008) combine a model of sequential consumer search with a position auction design by the search engine, without incorporating price setting by firms in the search pool. Armstrong, Vickers and Zhou (2009) analyze price competition with sequential consumer search, where one firm appears first in the consumers' search list. Finally, the main result in this paper is also reminiscent of Hagiu and Jullien (2009), who make a similar point - namely, that platforms in two-sided markets may have an incentive to put obstacles on consumer search - in the context of a very different two-sided-market model.

2 A Model

Let us begin with a market search model without a search engine, which extends a model due to Wolinsky (1986) by introducing heterogeneity among firms in a way that broadly follows Chen and He (2006). The market consists of a continuum of consumers and a continuum of firms. A firm's type is a number q, which is distributed over $[0, 1]$. When a consumer is matched with a firm of type q, the match has positive value for the consumer with probability q. Conditional on a positive-value match, the consumer's willingness to pay for the firm's product is randomly drawn (independently across all matches) from $U[0, 1]$. The firms' cost of providing their products is normalized to zero.

We interpret q as a measure of the firm's "relevance" for the consumers. For example, think of firms as websites providing holiday packages. A firm with a higher q corresponds to a website with a wider range of destinations and hotel types, such that the consumer's need is more likely to be met. Note that all the heterogeneity among firms is summarized by the probability of a positive-value match, but there is no heterogeneity conditional on this event. This modelling strategy greatly simplifies the analysis.

The market interaction proceeds as follows. Each firm simultaneously chooses a price for its product. Consumers form a belief of the distribution of prices in the market, and follow a conventional sequential-search process with a search cost of s per round. When a consumer samples a firm, he learns the value of the match and the firm's price, and optimally decides whether or not to continue searching (i.e., drawing a new sample from the population of firms). A stopping rule is a function that specifies the realized match values and prices for which the consumer stops searching.

For analytical convenience, we focus on market outcomes in which all firms charge the same price. A *uniform-price market equilibrium* is a price p^* and a stopping rule for consumers, which satisfy the following properties: (i) given that all firms charge p^*, the consumers' stopping rule is optimal; (ii) given the consumers' stopping rule and the belief that all firms charge p^*, no firm has an incentive to deviate to a different price.

Let us now introduce a monopolistic search engine into the model. Before the above market interaction takes place, the search engine limits firms' entry into the search pool. Specifically, the search engine posts a "price-per-click" r. This is a payment from the firm to the search engine each time a consumer visits the firm. Note that the payment is independent of whether the firm eventually transacts with the consumer. Only firms

that accept the posted price-per-click are admitted into the search pool. In the ensuing market equilibrium, consumers base their behavior on a correct expectation of the set of firms that entered the search pool. The search engine chooses r to maximize its revenue, which is r multiplied by the expected number of "clicks" - i.e., the expected number of samples that consumers draw in the market equilibrium induced by r.

Our assumption that firms are charged per click is motivated by the observation that this is how real-life search engines operate. We will depart from this assumption later in this paper. Following the same motivation, we assume that consumers are not charged for accessing the search engine's pool of firms. However, this assumption is also partly justified if there exists a "universal" pool where all firms belong (including those that are left outside the search engine's pool), where consumers can search for free. This pool can be interpreted as offline search. Since firms in the search engine's pool are on average more relevant than firms in the universal pool, consumers will tend to prefer searching in the former. However, if the search engine employs an access fee to extract consumers' surplus with access fee, this may impel them to switch to the universal pool.

3 Analysis

Our analysis proceeds in two steps. First, we take the set of firms that enter the search pool as given and characterize uniform-price equilibrium. Second, we incorporate this characterization into the search engine's problem and determine the optimal price-per-click.

3.1 Equilibrium Characterization for a Given Search Pool

Let us begin with a characterization of uniform-price market equilibria, taking the set of firms that entered the search pool as given. As in many other sequential-search models, our market model has a trivial equilibrium in which all firms post a price equal to the highest willingness to pay, $p = 1$, and consumers choose not to search at all. This is the equilibrium characterized by Diamond (1971) and known since then as the "Diamond Paradox". However, if search costs are sufficiently low (see below), there is also a uniform-price market equilibrium with active search, and we will focus on this equilibrium.

Proposition 1 *In a uniform-price market equilibrium with active search, consumers stop if and only if the value of a match with the current firm is at least*

$$v^* = 1 - \sqrt{\frac{2s}{E(q)}} \tag{1}$$

where $E(q)$ is the expectation with respect to the population of firms in the search pool. Firms charge the uniform price

$$p^* = 1 - v^* = \sqrt{\frac{2s}{E(q)}} \tag{2}$$

Proof. Our proof is a minor extension of a derivation by Wolinsky (1986). Let us begin with the consumers' stopping rule. Because all firms charge p^*, consumers face a stationary environment. Therefore, their stopping decision obeys a cutoff rule. That is, there exists $v^* \in [0, 1]$, given by

$$E(q) \cdot \int_{v^*}^1 (v - v^*)dv = s$$

such that in equilibrium, consumers stop if and only if the current match value is $v \geq v^*$. The L.H.S represents the incremental expected benefit from one more search, while the R.H.S represents the cost of one more search. The proof is standard and therefore omitted. Solving this equation yields the expression for v^*.

Now consider the pricing decision of a firm of type q. If the firm deviates from the equilibrium price p^* to another price p, a consumer who samples the firm and learns that the match value is $v > 0$ will buy the firm's product if

$$v - p > v^* - p^*$$

because the R.H.S of this inequality represents the consumer's reservation surplus conditional on a positive-value match. Thus, the probability that the consumer will buy at p is $1 - p - v^* + p^*$, by the assumption that v is drawn from $U[0, 1]$ (ignoring the possibility of corner solutions, which is easy to dismiss). Therefore, the firm will choose p to maximize

$$p \cdot (1 - p - v^* + p^*)$$

In equilibrium, the solution to this maximization problem coincides with p^*, yielding

$$p^* = 1 - v^*$$

∎

To see how our model relates to Wolinsky (1986), think of the consumer's "effective search cost" of a consumer as the total expected cost that a consumer incurs before reaching a positive-value (i.e., relevant) match. This is precisely $s/E(q)$. The model due to Wolinsky (1986) is a special case in which $q = 1$, hence the effective search cost coincides with s.

The gross profit-per-click that a firm of type q' earns in active-search equilibrium (i.e., excluding the transfer to the search engine) is

$$q' \cdot p^* \cdot (1 - v^*) = q' \cdot \frac{2s}{E(q)}$$

The conversion rate - namely, the expected stopping probability - is

$$E[q \cdot (1 - v^*)] = E[q \cdot \sqrt{\frac{2s}{E(q)}}] = \sqrt{2s \cdot E(q)}$$

Note that the inverse of this probability is the expected duration of search.

Turning to consumer welfare, note that consumers find it optimal to enter the market and face the uniform-price equilibrium only if their ex-ante expected surplus from searching in the pool is non-negative. The ex-ante expected surplus is equal to the expected value of the item that will ultimately be purchased, minus its equilibrium price minus the expected search costs. This amount is given by

$$E(v \mid v \geq v^*) - p^* - \frac{s}{\sqrt{2s \cdot E(q)}} = v^* - p^*$$

By Proposition 1, this reduces to

$$1 - \sqrt{\frac{8s}{E(q)}} \tag{3}$$

This means that the uniform-price equilibrium with active search exists if and only if $E(q) > 8s$.

3.2   The Optimal Price-Per-Click

In this sub-section, we assume that the uniform-price equilibrium with active search is played (whenever it exists) in the search pool induced by any given price-per-click. Let us characterize this search pool. Given that the search engine's price-per-click is r, a firm of type q' chooses to enter a pool if and only if

$$q' \cdot \frac{2s}{E(q)} \geq r$$

The expectation in $E(q)$ is taken with respect to the set of firms that choose to enter the search pool. If a firm of type q' prefers to enter, then any type $q'' > q'$ strictly prefer to enter. It follows that given r, the set of firm types that choose to enter is $[q^*, 1]$, where q^* is defined as follows:

$$q^* \cdot \frac{2s}{E_{q \geq q^*}(q)} = r \tag{4}$$

This equation may have multiple solutions. We will assume that in this case, the search engine is free to select its most desirable solution.

Recall that the search engine's expected revenue is the price-per-click r multiplied by the expected number of clicks in the induced equilibrium, which is the inverse of the conversion rate. This leads to our first main result.

Proposition 2 *The search engine's problem can be reformulated as follows: choose the critical type $q^* \in [0, 1)$ to maximize*

$$\frac{q^*}{[E_{q \geq q^*}(q)]^{\frac{3}{2}}} \tag{5}$$

subject to the constraint that a uniform-price market equilibrium with active search exists in the search pool induced by q^* - i.e.

$$E_{q \geq q^*}(q) \geq 8s$$

This maximization problem involves a subtle trade-off. When the search engine sets r in a way that effectively increases q^*, this has several implications. First, the change in the profit-per-click is ambiguous. On one hand, the equilibrium product price goes

7

down because a higher-quality search pool creates a more competitive environment. On the other hand, the conversion rate goes up. Both are equilibrium implications of the increase in the quality of the search pool, and they have opposite effects on the firms' profit-per-click. Second, the increase in q^* lowers the expected number of clicks (because it is the inverse of the conversion rate). This lowers the search engine's total profit.

Note that the domain in the search engine's maximization problem is $[0, 1)$. The reason $q^* = 1$ is not feasible is technical - we want to ensure that for any q^* that the firm may set, there is a strictly positive measure of firms that enter the search pool. When the maximization problem has no solution because of this open-set feature, the firm wants to set the cutoff q^* arbitrarily close to 1. Our interpretation is that it is optimal for the search engine to admit only the highest quality firms.

However, the search engine may find it optimal to set r such that $q^* < 1$, which is tantamount to contaminating the search pool with firms of relatively low relevance. The following simple example illustrates this motive.

Example 1 *Suppose that q is distributed as follows: with probability α, $q = 1$, and with probability $1 - \alpha$, $q = L < 1$. If the firm sets r such that $q^* = 1$, its normalized total profit is 1. If the search engine sets r such that $q^* = L$, its normalized total profit is*

$$\frac{L}{[\alpha + (1 - \alpha)L]^{\frac{3}{2}}}$$

The search engine will strictly prefer to induce $q^* = L$ by setting

$$r = \frac{2sL}{\alpha + (1 - \alpha)L}$$

whenever

$$\alpha < \frac{L^{\frac{2}{3}}(1 - L^{\frac{1}{3}})}{1 - L}$$

For example, when $L = \frac{1}{2}$, the search engine will degrade the quality of the search pool whenever $\alpha \lesssim \frac{1}{4}$.

It is interesting to note that even if the search engine charged consumers for accessing its pool, it may still find it optimal to degrade the quality of this pool. To see this, suppose that the search engine could charge consumers an access fee that extracts their entire ex-ante surplus from using its search pool. Then by (3) and (5), the search

engine's problem would consist of choosing q^* to maximize the sum,

$$1 + \frac{q^*}{[E_{q \geq q^*}(q)]^{\frac{3}{2}}} - \sqrt{\frac{8s}{E_{q \geq q^*}(q)}}$$

This means that the search engine's incentive to set a low q^* is somewhat curbed by the negative effect this has on consumer's surplus. However, the search engine may still find it optimal to set $q^* < 1$. For instance, if in Example 1, $s = \frac{1}{64}$, $\alpha = 0.18$ and $L = 0.6$, setting $q^* = L$ maximizes the search engine's profit.

3.3   Equivalence between Access Fee and Price-Per-Click

An alternative way to restrict firms' entry into the search engine's pool is to charge firms a lump-sum access fee a. This method is characteristic of yellow pages, because unlike internet-age search engines, yellow pages are unable to monitor the consumers' search activities. As with price-per-clicks, if a firm of type q decided to pay a, so would every $q' > q$. As in the previous subsection, let q^* denote the marginal firm type, which is indifferent between paying a and staying out of the pool. The access fee a would then be equal to the expected profit of firm q^*. To compute this profit, let ρ denote the ratio of consumers to the firms in the general population. The expected total number of clicks that each firm receives is

$$\frac{1}{1 - \{1 - E_{q \geq q^*}[q \cdot (1 - v^*)]\}} \cdot \frac{\rho}{1 - F(q^*)}$$

The marginal firm's profit-per-click is, as before, $q^* \cdot (1 - v^*) \cdot p^*$, where p^* is the equilibrium price in the search pool. Therefore:

$$a = \frac{\rho}{1 - F(q^*)} \cdot \frac{q^*}{E_{q \geq q^*}(q)} \cdot p^*$$

By (2), the search engine's decision problem is to choose $q^* \in [0, 1)$ that maximizes the expression,

$$\frac{\rho\sqrt{2s}}{1 - F(q^*)} \cdot \frac{q^*}{[E_{q \geq q^*}(q)]^{\frac{3}{2}}} \cdot [1 - F(q^*)]$$

This leads to our second main result.

Proposition 3 *The search pool that maximizes the search engine's profit is independent of whether it charges a fixed access fee or a price-per-click.*

Thus, our conclusion that the search engine may find it optimal to degrade the quality of its search pool extends to the case in which it charges a fixed access fee.

4 Competition between Search Engines

We conclude this paper with a brief discussion of a competitive version of our model of search engine pricing. Two search engines simultaneously choose a price-per-click. Given a profile (r_1, r_2) of prices-per-click, each firm chooses whether to enter each of the two search pools operated by the two search engines. Note that firms can enter both pools. Simultaneously, consumers choose to enter at most one of the two search pools. After firms and consumers make their entry decisions, they proceed to play an active-search equilibrium (if one exists) in the search pool.

Coordination failures introduce uninteresting equilibria where all consumers enter only one pool. We, therefore, consider the following equilibrium refinement: given (r_1, r_2), each firm enters any pool which would be profitable if a positive measure of consumers entered that pool. This is a reasonable refinement, since a firm is indifferent between entering and staying out of a search pool when no consumers enter it. The refinement may be viewed as capturing firms' response to some form of "trembles" on the part of consumers: with small probability a consumer enters a pool even if the other pool leads to strictly higher profits.

Thus, a firm of type q will enter the pool operated by search engine i if and only if $q \geq q_i^*$, where q_i^* solves equation (4) (we ignore the possibility of multiple solutions). Suppose that $q_1^* > q_2^*$. Then, consumers necessarily choose to enter the search pool operated by search engine 1. The reason is that this search pool is characterized by a lower equilibrium price and a higher stopping probability than the search pool operated by search engine 2. It follows that in market equilibrium, both firms will set $r_1 = r_2 = 2s$, such that $q_1^* = q_2^* = 1$. That is, competitive forces maximize search quality. As a result, they maximize the price-per-click that search engines post and the conversion rate, and they minimize the product price that consumers confront in the search process. In this sense, search engine pricing responds to competition in a way that benefits consumers.

5 Concluding Remarks

Throughout our analysis, we assumed that the search engine has only one means for controlling the quality of the search pool - namely, manipulating the price that it

charges from firms. Suppose that the search engine could also directly manipulate the search cost s (e.g., by slowing down its server, or by deliberately throwing irrelevant links into query results). Then, it is straightforward to see that the search engine would set a price-per-click $r^* = \frac{1}{4}$, coupled with a search cost $s^* = \frac{1}{8}$. This induces an outcome in which only firms of type $q = 1$ enter the search pool, while consumers earn a zero surplus. In reality, reputational concerns may prevent search engines from deliberately degrading their search pool in this direct manner.

Our analysis has also abstracted from issues of prominence (which are typically addresses by a search engine through the use of position auction). One justification for this is our assumption that there is a large supply of firms (approximated by infinity). To see why, consider the case of two firm types, $q = 1$ and $q = L < 1$, that bid in a position auction. In reality, search engines do not position the query results in exact accordance with the ordering of the bids, but "blur" it to some extent. Furthermore, the precise blurring algorithm is kept secret, so that the firms do not know exactly how their bidding behavior maps into position.

Suppose that the search engine does not blur the ordering of bids at all. Then, in the auction process, firms with $q = 1$ will outbid firms with $q = L$, and as a result, the search pool will effectively consist of $q = 1$ types only. In equilibrium, these firms will pay the amount that gives them zero profits, which is given by our model for $q^* = 1$. Now suppose that the search engine completely blurs the ordering of bids. Then, firms have no reason to bid for prominence. Therefore, in equilibrium firms will never bid above the reserve price-per-click set by the search engine. The optimal reserve price is given by our model. And as we saw, it may be optimal for the search engine to set the reserve price such that all firms will choose to enter. What our analysis in this paper does not capture is partial blurring of the ordering of bids in a position auction. However, our model at least shows that in a position auction with a large number of firms, the search engine may prefer total blurring to no blurring at all.

References

[1] Armstrong, M., J. Vickers and J. Zhou (2009), "Prominence and Consumer Search", RAND Journal of Economics 40, 209-233.

[2] Athey, S. and G. Ellison (2008), "Position Auctions with Consumer Search," mimeo.

[3] Chen, Y. and C. He (2006), "Paid Placement: Advertising and Search on the Internet," mimeo.

[4] Diamond, P. (1971): "A Model of Price Adjustment," Journal of Economic Theory, 3, 156-168.

[5] Hagiu, A. and B. Jullien (2009), "Why do Intermediaries Divert Search?", mimeo.

[6] Wolinsky, A (1986), "True Monopolistic Competition as a Result of Imperfect Information," Quarterly Journal of Economics 101, 493-511.