AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

A Simple Model of Search Engine Pricing*

Kfir Eliaz[†] and Ran Spiegler[‡]

September 29, 2010

Abstract

We present a simple model of how a monopolistic search engine optimally determines the average relevance 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-per-click set by the search engine. We show that a monopolistic search engine may have an incentive to set a relatively low price-per-click that encourages low-relevance advertisers to enter the search pool. In general, the ratio between the marginal and average relevance in the search pool induced by the search engine's policy is equal to the ratio between the search engine's profit per consumer and the equilibrium product price. These conclusions do not change if the search engine charges fixed access fee rather than a price-per-click.

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

---

*Financial support from ERC grant no. 230251, NSF grant SES #0802789, ESRC (UK) and the Sapir Center is gratefully acknowledged. We thank Mark Armstrong for helpful comments.

[†]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/>

queries according to a mechanism in which firms pay the search engine for (prominent) appearance on the list of query results.

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 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 according to a *cdf* G , the support of which is contained in $[0, 1]$. 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 a continuous density function f defined over the support $[0, 1]$. Let F denote the *cdf* induced by f . We assume that f satisfies the usual increasing hazard rate property. 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 about 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.

First, note that since the search engine charges per click, this rate does not affect the firms’ pricing decisions given that they are in the search pool. The reason is that once a firm encounters a consumer, it regards the amount it pays the search engine as a sunk cost. Accordingly, in what follows we entirely ignore the amount that firms pay the search engine.

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^* , and firms charge a price p^* , where v^* and p^* are uniquely given by the pair of equations*

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

$$p^* = \frac{1 - F(v^*)}{f(v^*)} \quad (2)$$

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

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 (1), such that in equilibrium, consumers stop if and only if the current match value is $v \geq v^*$. The L.H.S of (1) 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.

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 $1 - F(v^* + p - p^*)$. Therefore, the firm will choose p to maximize

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

The first-order condition is

$$1 - p \cdot f(v^* + p - p^*) - F(v^* + p - p^*) = 0$$

In equilibrium, the solution to this equation is $p = p^*$, yielding (2). ■

To see how our model relates to Wolinsky (1986), think of the consumer's "effective search cost" as the total expected cost he 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 equilibrium probability that a consumer who clicks on a firm of type q will buy from that firm - namely, the conversion rate that characterizes such a firm - is

$$q \cdot (1 - F(v^*))$$

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

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

The expected conversion rate in equilibrium is

$$E(q) \cdot (1 - F(v^*)) \tag{3}$$

Note that the inverse of the expected conversion rate is the equilibrium 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

$$E(v \mid v \geq v^*) - p^* - \frac{s}{E(q) \cdot (1 - F(v^*))}$$

By (1), this expression is equal to $v^* - p^*$. By (2) and the increasing hazard rate property, p^* is strictly decreasing in v^* ; and by (1), v^* is strictly decreasing in s . If s is sufficiently large, the expression $v^* - p^*$ is negative, in contradiction to the requirement that consumers choose to search in equilibrium. Thus, when s is sufficiently large, there exists no uniform-price equilibrium with active search.

An example: a uniform valuation distribution

When f is the uniform distribution over $[0, 1]$, equations (1) and (2) have a closed-form solution:

$$v^* = 1 - \sqrt{\frac{2s}{E(q)}} \quad (4)$$

$$p^* = \sqrt{\frac{2s}{E(q)}} \quad (5)$$

The equilibrium conversion rate is

$$\sqrt{2s \cdot E(q)}$$

Finally, a uniform-price equilibrium with active search exists if and only if $E(q) \geq 8s$.

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. We assume that the search engine incurs no costs. Fixing the distribution of firm types in the search pool, and 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

$$q' \cdot \frac{(1 - F(v^*))^2}{f(v^*)} \geq r$$

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:

$$q^* \cdot \frac{(1 - F(v^*))^2}{f(v^*)} = r \quad (6)$$

It follows that the term $E(q)$ in expression (1) that implicitly defines v^* - as well as any other expression that contains this term - should be written, more precisely, as

$E_G(q \mid q \geq q^*)$. Equation (6) 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 given by (3). 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*

$$\frac{q^*}{E_G(q \mid q \geq q^*)} \cdot \frac{1 - F(v^*)}{f(v^*)} \quad (7)$$

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

This maximization problem involves subtle trade-offs. When the search engine sets r in a way that effectively raises the search pool's marginal relevance q^* , this has several implications. Recall that the search engine's profit is equal to the price-per-click multiplied by the expected number of clicks, and that the price-per-click is equal to the marginal firm's gross profit per click. The effect of raising q^* on this gross profit is ambiguous. First, note that although the marginal relevance in the search pool goes up, the average relevance $E_G(q \mid q \geq q^*)$ can go either way, depending on the shape of G . However, for the sake of the argument, suppose that $E_G(q \mid q \geq q^*)$ increases with q^* . On one hand, the equilibrium product price goes down because a higher-quality search pool creates a more competitive environment. On the other hand, the net effect on the expected conversion rate is ambiguous, because an increase in the average relevance of the search pool causes v^* to go up, such that $1 - F(v^*)$ decreases. Thus, the marginal firm's gross profit per click can go either way. Finally, the expected number of clicks, which is the inverse of the conversion rate, can go either way because as we saw, the effect of raising q^* on the conversion rate is ambiguous. In order to obtain clear-cut predictions, we need to impose particular assumptions on F and G . We shall do so below.

Nevertheless, certain relations turn out to be independent of distributional assumptions. The following is another way of expressing Proposition 2. Expression (7) represents the expected profit that the search engine earns per consumer. It follows that the ratio between this profit per consumer, denoted π^* , and the equilibrium product

price that firms charge is given by the formula

$$\frac{\pi^*}{p^*} = \frac{q^*}{E_G(q \mid q \geq q^*)} \quad (8)$$

This formula is that it is independent of the distribution of consumer valuations; it is expressed entirely in terms of the distribution of firm types. Another nice feature is that the ratio π^*/p^* can be interpreted as a unit-free measure of the “intermediation fee” that the search engine extracts. Formula (8) connects this quantity to the equilibrium ratio between the relevance of the marginal and average firms in the search pool. When the search engine does not degrade the search pool, this ratio is equal to one.

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.

When v is uniformly distributed over $[0, 1]$, we can plug the closed expression for v^* derived in the previous sub-section into (7), and obtain that the critical type q^* maximizes

$$\frac{q^*}{[E_G(q \mid q \geq q^*)]^{\frac{3}{2}}}$$

It can be seen that depending on the shape of G , 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 effect.

Example 1 *Suppose that v is uniformly distributed over $[0, 1]$ and 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 $\alpha \lesssim \frac{1}{4}$.

Comment: How does invasion of relevant firms affect the search engine's profit?

Suppose that we add firms of type $q = 1$ to the general population of firms. Then, for every possible critical type q^* , this invasion of $q = 1$ types causes $E_G(q \mid q \geq q^*)$ to go up. This has two effects on (7). First, the effective search cost decreases, such that v^* goes up and p^* goes down. Second, the ratio between the marginal and average relevance in the search pool goes down. Both effects hurt the search engine. Since this holds for any possible q^* , it follows that the search engine's maximal profits decrease as a result of the invasion of firms of type $q = 1$ to the general population. On the other hand, invasion of firms of types close to $q = 0$ to the general pool has no effect on the search engine's profit, because these firms will be excluded from the search pool that the search engine creates.

Comment: charging consumers

Throughout the paper, we assume that the search engine charges fees from firms only. 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, consider the case in which v is uniformly distributed over $[0, 1]$, and suppose that the search engine could charge consumers an access fee that extracts their entire ex-ante surplus from using its search pool. Then, the search engine's problem would consist of choosing q^* to maximize

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

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, firms regard the payment to the search engine as a sunk cost at the time they set their prices. Thus, for a given search pool, the uniform-price market equilibrium is exactly the same as in the price-per-click case. Moreover, 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, normalize (without loss of generality) the ratio of consumers to the firms in the general population to one. Therefore, the ratio between the number of consumers and the number of firms in the search pool induced by the search engine's policy is

$$\frac{1}{1 - G(q^*)}$$

Denote this quantity by σ , and denote the equilibrium conversion rate by γ . Then, the total number of clicks that an individual firm receives consists of σ clicks by consumers for whom the firm is the first they visit, $\sigma(1 - \gamma)$ clicks by consumers for whom the firm is the second they visit, $\sigma(1 - \gamma)^2$ clicks by consumers for whom the firm is the third they visit, and so forth. It follows that the expected total number of clicks that each firm receives is

$$\frac{1}{E_G(q \mid q \geq q^*) \cdot (1 - F(v^*))} \cdot \frac{1}{1 - G(q^*)}$$

The marginal firm's gross profit per click is, as before, $q^* \cdot (1 - F(v^*)) \cdot p^*$, where p^* and v^* are given by (2) and (1). Therefore,

$$a = \frac{1}{E_G(q \mid q \geq q^*) \cdot (1 - F(v^*))} \cdot \frac{1}{1 - G(q^*)} \cdot q^* \cdot (1 - F(v^*)) \cdot p^*$$

The search engine's decision problem is to choose $q^* \in [0, 1)$ that maximizes

$$a \cdot (1 - G(q^*))$$

This immediately leads to the following 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.*

In particular, 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. The key formula (8) continues to hold as well.

4 Competition between Search Engines

Our results in the previous sections raise the question of whether competition between search engines would benefit consumers. To address this question we analyze the following extensive-form game with a pair of search engines and a continuum of firms. In the first stage, the two search engines simultaneously choose a price-per-click. Given the pair (r_1, r_2) of prices-per-click selected in the first stage, in the second stage 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. We assume that for each consumer, there is an arbitrarily small probability ε that he would “tremble” and enter a pool even if the other pool leads to higher expected payoffs. After firms and consumers make their entry decisions, they proceed to play an active-search equilibrium (if one exists) in the search pool. As in the previous section, we focus on a uniform-price equilibrium, which is characterized by a cutoff v^* for the consumers and a price p^* for the firms.

Before we present our result, we note that the role of consumers’ trembles is to avoid coordination failures, which introduce uninteresting equilibria where all consumers enter only one pool. Our assumption implies that 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.

Proposition 4 *In any equilibrium satisfying our refinement, each search pool consists of only the highest quality firms.*

Proof. Consider a firm of type q . This firm type will enter the pool operated by search engine i if and only if $q \geq q_i^*$, where q_i^* solves equation (6) (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, firms will set a price-per-click that will induce $q_1^* = q_2^* = 1$. ■

The above results shows that 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 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, in the case where consumers' valuations are uniformly distributed, it is straightforward to show 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 addressed by means of position auctions). One justification for this is our assumption that there is a large supply of firms. 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.