Search Engines: Left Side Quality versus Right Side Profits*

Alexander White†

May 13, 2008

Abstract

This paper models a monopoly search engine, focusing on the fact that a search results page typically involves both paid links and unpaid links. It argues that the presence of *both* commercial and seemingly uncommercial components is the key ingredient to search engines' profitability. It analyzes the search engine's decision of how many advertisements to place and what level of search quality to provide. Doing so, it gives an endogenous explanation for why a search engine gives "free advertising" to non-paying websites but also shows that the search engine has incentive not to offer maximal search quality, even when it would be technically costless to do so.

1 Introduction

This paper tries to understand and explain search engines' recipe for commercial success. In the time since the arrival of the Internet, the search engine has become an essential tool for doing many things that seem totally uncommercial. People use search engines for myriad different purposes, including locating specific websites, investigating the meaning of words, finding out information about historical events, just to name a few. Concomitantly, search engines have been at the center of a rapidly growing market for online advertising and a seemingly ubiquitous discussion thereabout. Picking up a major newspaper at random, these days, it would be very unusual *not* to find an article on this subject. Indeed, a quick search on the *New York Times* website for articles containing the word "Google", published in the 30 days prior to writing, returns 100 results.

A driving force for this attention is search engines' both current and potential profitability, an example of which comes from Microsoft's attempt to acquire Yahoo!. The latter, widely seen to be struggling, nevertheless rejected Microsoft's 2008 offer to pay $45 billion for it, claiming in that this sum "substantially undervalues" the company. Meanwhile, search behemoth Google,

*Many thanks to Klênio Barbosa, Bruno Jullien, Paul Seabright, Glen Weyl, a seminar audience in Toulouse and especially to Jacques Crémer for very helpful discussions. This is a preliminary draft which doubtless contains many errors, for which none of the aforementioned bear any responsibility.

†Toulouse School of Economics, alexander.white@univ-tlse1.fr

whose 2007 net income reached \$4.2 billion, is currently valued by the market at nearly \$140 billion.

A central reason for this profitability is the fact that search engines have turned out to be very good at facilitating commercial transactions. People who are potentially interested in buying a particular good or service often find it useful to type the name of this item into the "query bar" of a search engine and to look at the "results". Having done so, they also often find it convenient to make their purchase directly from a website whose link appears on the results page. Thus, a search engine seems, in some ways, to be a typical intermediary or platform that connects buyers and sellers.

At the same time, different types of platforms have different characteristics, and search engines have two such characteristics that, together, make them particularly interesting. First, since users type in a "search term" *before* being shown a results page, the search engine is in a position to exert very precise control over which websites a user interacts with. This contrasts, for instance, with a credit card, which also serves as an intermediary, but which does not receive such detailed information or play such an active role in linking the two sides of the market. Thus, it would seem that a search engine is more similar to a directory service, such as a traditional copy of the Yellow Pages, in which a reader selects a topic based on a need and then finds a list of corresponding advertisements.

Second – and unlike a traditional directory service – a search engine's results page consists of a *combination* of two types of results: paid advertisements or "sponsored" links, typically found on the right side of the screen, as well as unpaid results or "organic" links, typically found on the left side.[1] For a link to a given website be found among the former, the owner of this site must have actively sought (and paid) to advertise on the search engine's results page for a particular query. On the other hand, the organic results are chosen by the search engine's computer algorithm and do not depend, explicitly, on any type of payment. Under the typical arrangement, each time someone clicks on a paid link, a transfer is made from the advertiser to the search engine. In contrast, when someone clicks on an unpaid link, no money changes hands.

Therefore, to put things one way, if a user were definitely going to click on the link to a particular website, the search engine would prefer that this link be paid rather than an organic. However, an almost defining characteristic of the search engine is that it provides relevant organic

[1]This spacial arrangement for placing ads is, in fact, often violated, particularly by the presence of sponsored links at the top of the left-hand side of the page. In this paper, I nevertheless refer interchangeably to "organic", "unpaid" and "left side" links on the one hand, and to "sponsored", "paid" and "right side" links on the other.

results. This, it would appear, is both what distinguishes search engines from directory services or shopping websites and what has been central to their commercial success and overwhelming rise to prominence.

In this paper, I attempt to explain how search engines make money by focusing on the fact that a search engine is clearly *more* than just a directory service and, by asking more precisely, just what else is it? To do this, I develop, sequentially, a model of a monopoly search engine designed to address these issues. The simplest, baseline, version of the model represents the problem facing a traditional directory service. The two subsequent extensions modify the original framework in order to represent, with a finer brush, the unique characteristics of a search engine.

To do this, I introduce, in the baseline model, the notion of user "query costs", and in the subsequent extensions, a variable representing the engine's "search quality". In the game, the search engine endogenously chooses a search quality which affects user costs and thus affects the number of people who choose to use the search engine. In the first extension, I suppose that the placement of merchants between the sponsored links and the organic links is exogenous but that search quality is costly to produce. I then show that there is a non-trivial tradeoff between the choice of search quality and the choice of how many sponsored links to display.

In the second extension, I endogenous the layout of the search results page. Here, I assume that search quality entails no explicit costs to produce, but instead that it comes with the implicit cost to the search engine of driving user attention away from the paid links to the organic ones. One insight provided by this section is the fact that we should expect the "organic" results shown on a search page in fact to be motivated by commercial concerns even though they are not directly paid for. From a less sanctimonious perspective, this section and this entire paper suggest that an important reason for search engines' success is fact that they are a subtly-designed bottleneck. They are so effective precisely because they allow challengers to appeal directly to consumers instead of forcing all suiters to pay for access.

Related Literature. This paper represents the first attempt that I know of in the literature to explain both the organic and sponsored elements of search engines. Recently, there has been much attention paid to the auction mechanisms search engines use to sell their paid links, including Edelman, Ostrovsky & Scharz (2007), Varian (2007), and Borgers, Cox, Pesendorfer & Petricek (2007). Among this literature, the paper that perhaps is closest to this one is Athey & Ellison (2007), in the limited sense that explicitly models internet users and not just advertisers. This paper's initial setup is similar to that of Baye & Morgan (2001) but its focus on the specific characteristics of search engines is very different from the focus of that paper. The approach of this paper is broadly inspired by the two-sided markets literature and, in particular, by the results presented in Rochet & Tirole (2007). Analysis of how high a two-sided platform should set a different but related form of "search cost" can be found in Jullien & Hagiu (2007). Finally, this paper's model is emblematic of the room for ambiguity between a traditional "middleman" buying upstream, selling downstream and a two-sided platform, an issue examined more closely in Hagiu (2007).

Roadmap. The paper is organized as follows. Section 2 presents the baseline model of a "directory service", illustrating a basic tradeoff that is present throughout the paper. Section 3 contains the first extension designed to capture some distinguishing features of a search engine but abstracts from direct consideration of organic versus paid results. Section 4 presents the model that focuses on this issue.

2 Baseline Case: A Directory Service

I begin by considering the simplest case in order to illustrate the overarching trade-off. On the one hand, the search engine must attract users to its site, but on the other hand, it must create conditions that are profitable for advertising merchants. In this baseline version of the model, each user faces an exogenously given cost of using the search engine. This cost can be thought of as the personal disutility incurred by the user both to access the site *and* to learn her valuation for the (homogeneous) good in question.

To motivate the assumption that the user does not know, *ex ante*, her valuation for the good, despite the fact that it is homogeneous, it is helpful to ask why the user choses to visit the search site in the first place. Given the multitude of consumer goods that are sold in today's economy, it is quite obvious that no one knows "off hand" precisely how much he would be willing to pay in, a given circumstance, for every single item on sale. A natural reason for using a search engine is thus not only to find merchants who sell a given good but also to reduce one's ignorance about the good itself. One way to model this ignorance would be to consider heterogeneous goods with characteristics that the user does not observe.

For three reasons, I avoid using this modeling strategy. First, it would imply heterogeneity among competing merchants and would thus require modeling to the auctions that the search engine would hold to extract this private information. This is the subject of a quickly-growing body of literature, but it is not what I am focusing on in this paper. Second, it would suggest an unrealistically high degree of consumer sophistocation, since it would imply that the sole obstacle preventing users from knowing their "true" willingness to pay for any given good is the lack of access to information that *is* known by merchants. Finally, it would make the model much more complicated.

Instead, I assume simply that the good is homogenous but that the user is uninformed of her valuation before she performs the search. One might imagine that the user is employing the search engine to find out more precisely what a particular product is. Alternatively, it is also reasonable to picture a consumer with a particular "need", that she is able to describe using a search query. The results of the query would then identify a good that is potentially capable of satisfying this need. With access to these results, the user would then be privately able to determine to what extent (or with what probability) the good would fulfill her need, given her personal circumstances.

2.1 The Model

I now formally describe the model. There are three types of players in the game: a "search engine", an arbitrarily large set of "merchants", and a continuous mass of "users". To summarize the game, first the search engine sets an advertising fee. Second, merchants decide whether to pay this fee and become "advertisers". At the same time, users decide whether or not to search for the good. Finally, advertisers compete with one another to sell the good to users who have chosen to search.

Preferences. Each user has a privately known "query cost", θ_i, which is drawn from a continuous distribution, with cdf $F(\cdot)$ and pdf $f(\cdot)$, with positive support on the interval $[\underline{\theta}, \bar{\theta}]$, where $0 \leq \underline{\theta} < \bar{\theta}$. A user who doesn't search gets a payoff of zero (and, by assumption, is not able to access any merchants). By searching, a user learns her valuation for the good, v_i, which is drawn from a continuous distribution, with cdf $G(\cdot)$ and pdf $g(\cdot)$, with positive support on the interval $[\underline{v}, \bar{v}]$, where $0 \leq \underline{v} < \bar{v}$. I assume that the two distributions are independent of one another.

Letting u_i denote user i's (risk-neutral) utility function, we can write:

$$u_i = \begin{cases} v_i - p - \theta_i, & \text{if searches and buys good at price } p \\ -\theta_i, & \text{if searches and does not buy good} \\ 0, & \text{if does not search} \end{cases}$$

Merchant profits are straightforward. Each merchant either pays a fee, A, to the search engine in order to advertise and access users, or he doesn't and gets a payoff of zero. Merchants are assumed to produce the good at a common marginal cost, c. So, for a merchant, j, who advertises, profits are given by π_j, where:

$$\pi_j = (p_j - c)q_j - A$$

In this expression, p_j stands for the price j charges for the good, and q_j is the quantity he sells. I assume that in the final stage of the game, the advertising merchants engage in Cournot competition, and as is standard, that demand is such that (for a given mass of users) total merchant profits strictly decrease as the number of adversing merchants increases, tending to zero as the number of merchants grows large.

Finally, search engine profits come exclusively from advertising revenue. While in principle, it might be possible for the search engine also to charge users, I assume that it cannot. This seems like the most reasonable assumption to make, since, in practice, doing so would require users to go through a time-consuming registration process and would thus add significant transaction costs. As a result, search engine profits, Π^{SE}, are given by:

$$\Pi^{SE} = n \times A$$

Here n is the number of merchants who sign up to advertise, each paying the fee, A.

Timing. The stages of the game are as follows:

1. Nature assigns a private "query cost", θ_i, to each user, from distribution $F(\cdot)$.
2. The search engine sets the advertising fee, A.
3. Merchants choose whether to advertise, and users choose whether to search.
4. Each user who searches learns her valuation for the good, v_i, from distribution $G(\cdot)$.
5. The n merchants who advertise compete *à la* Cournot, to sell to the users who searched.
6. Users who searched either buy or don't buy the good from an advertising merchant.

2.2 Equilibrium Results

I look for pure-strategy Perfect Bayesian Equilibria. Thus, in the final stage, we can restrict our attention to the symmetric[2] Cournot equilibirum with n merchants facing a demand function $D(\cdot) \equiv 1 - G(\cdot)$. The true demand facing these merchants is given by $m \times D(\cdot)$, where m denotes the mass of users who decided to search; however, because of the independence of the

[2] Asymmetric equilibria in the last stage do not survive the "perfection" refinement, since all merchants are charged a uniform advertising fee that extracts all of their rent.

distribution of users' query costs and users' valuation for the good, the number of users who search does not directly affect the equilibrium price of good. It does, on the other hand, affect the price of the good *indirectly*, by influencing merchants' decision whether or not to advertise and thereby influencing the equilibrium number of merchants in the market.

The key feature of the model is the search engine's complete control over the equilibrium price in the final-stage Cournot competition. In setting the advertising fee, A, the search engine determines both the number of advertisers, n, which, for convenience, we treat as a continuous variable, and the mass of users, m. Consequently, we can view the search engine as if were solving a monopoly pricing problem. Proposition 1 establishes this.

Proposition 1. *The problem facing the search engine when choosing its advertising fee, A, can be rewritten as a problem of choosing the equilibrium price, p, of the good. This problem can be written as:*

$$\max_{p} \left[\Pi^{SE} = m(p)D(p)[p-c]\right] \tag{1}$$

where $m(\cdot)$ is the mass of users who search as a function of the price. Moreover, the expression for the mass of users, $m(p)$, is given by:

$$m(p) = F\left(\int_{p}^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right) \tag{2}$$

Proof: See Appendix.

Setting the problem up this way allows us to isolate the key issues facing search engine, illustrated by proposition 2.

Proposition 2. *The equilibrium price of the good, p^*, induced search engine, must satisfy the following dual inverse-elasticity rule:*

$$\frac{p^*-c}{p^*} = \frac{1}{\varepsilon_m + \varepsilon_D} \tag{3}$$

where ε_m is the elasticity of the mass of users who search with respect to price, $\varepsilon_m = \frac{dm(p^*)}{dp^*}\frac{p^*}{m(p^*)}$, and where ε_D is the final-stage elasticity of demand with respect to price, $\varepsilon_D = \frac{dD(p^*)}{dp^*}\frac{p^*}{D(p^*)}$.

The expression for ε_m is given by:

$$\varepsilon_m = \frac{f\left(\int_{p^*}^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right)}{F\left(\int_{p^*}^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right)} D(p^*)p^* \tag{4}$$

Proof: See Appendix.

Expression (3) gets to the heart of the issue facing the search engine. In maximizing profits, it must attract users to perform queries. Since this is costly for users, they do so only if the market price is low enough to compensate for this cost. Consequently, the search engine sets an advertising fee that is low enough to attract multiple advertisers. The following corollary gives the expression for the number of advertisers as a function of the aforementioned elasticities.

Corollary to Proposition 2. *The search engine chooses the number of advertisers, n^*, to satisfy:*

$$n^* = \frac{\varepsilon_m}{\varepsilon_D} + 1 \tag{5}$$

Proof: See Appendix.

Note that if the number of users were set first, *and afterwards* the search engine could choose its advertising fee, then the search engine would induce standard monopoly pricing, satisfying the classical Lerner formula. Here, however, the mass of demand that is present in the market depends on the expected price. Therefore, the search engine induces a lower price by allowing for competition among advertisers. The intensity of competition it induces (i.e. the number of merchants it chooses to let advertise) grows with the sensitivity of users to price *before they search relative to their sensitivity afterwards.*

Using equation the formula given in (4) for ε_m as well as the definition of ε_D gives the following expression for the ratio of these elasticities. Omitting the notationally cumbersome arguments of the functions, we have:

$$\frac{\varepsilon_m}{\varepsilon_D} = -\frac{f}{F}\frac{D^2}{D'} \tag{6}$$

This formula helps convey the intuition for how the ratio of elasticities varies as the distributions of the users' private valuations change. This result, on the one hand, is commensurate with the straightforward intuition that, from the search engine's perspective, the optimal level of competition is higher when new many new users are readily attracted by a small decrease in prices. Of particular note, however, is the quadratic term in the numerator. This term reflects the fact that the final stage demand, $D(\cdot)$, plays two roles "in support" of letting in more advertisers.

The first role follows from the standard result that in Cournot competition, if demand is less elastic a given number of competing firms will set higher prices. The second role $D(\cdot)$ plays is in the marginal user's decision whether or not to query. Since her expected utility from searching is given by the expression,

$$[E[v|\text{buy}] - p] \times \Pr[\text{buy}] = \int_{p}^{\bar{v}} D(\tilde{p})\, d\tilde{p} \tag{7}$$

ceteris paribus, a higher $D(\cdot)$, evaluated at the candidate price, means that this marginal user's expected utility from searching *increases faster* as competition increases.

2.3 Example with Uniform Distributions

I now give a simple example which allows us to examine closed-form solutions to the game. Assume that both the users' query costs and their valuations for the good are uniformly (and still independently) distributed, each over some non-negative interval: $\theta_i \sim \mathrm{U}\left[\underline{\theta}, \bar{\theta}\right]$ and $v_i \sim \mathrm{U}\left[\underline{v}, \bar{v}\right]$. (For notational ease, define $\Delta\theta \equiv \bar{\theta} - \underline{\theta}$ and $\Delta v \equiv \bar{v} - \underline{v}$.) With these assumptions, the game can easily be solved.

In the final stage, there are n advertisers competing *à la* Cournot for total demand given by:

$$m \cdot D(p) = m \cdot \frac{\bar{v} - p}{\Delta v}$$

which implies the inverse-demand function, $p(\cdot)$, given by:

$$p\left(q_j + \sum_{k \neq j} q_k\right) = \bar{v} - \frac{\Delta v \left[q_j + \sum_{k \neq j} q_k\right]}{m}$$

advertisers then solve the problem:

$$\max_{q_j} \left\{ q_j \times [p(q_j + \sum_{k \neq j} q_k) - c] \right\} \tag{8}$$

This leads to an equilibrium price among advertisers, $p^*_{ad}(n)$, given, as a function of the number of competitors, by:

$$p^*_{ad}(n) = \frac{\bar{v} + nc}{n + 1} \tag{9}$$

Note that, for a given number of advertisers, n, this price is independent of m, the mass of users who decide to search. Moreover, note that, as discussed above, this independence does not depend on the specific demand function used in this example.

At the "entry stage" during which merchants decide whether or not to advertise and users decide whether or not to search, user i will search if and only if:

$$[E[v|\text{buy}] - p] \times \Pr[\text{buy}] = \frac{[\bar{v} - p]^2}{2\Delta v} \geq \theta_i$$

The mass of users who search is then given by:

$$m(p) = F\left(\frac{[\bar{v} - p]^2}{2\Delta v}\right) = \frac{[\bar{v} - p]^2}{2\Delta v \Delta\theta} \tag{10}$$

Rather than performing the remaining calculations demanded by conventional "backward induction", we can jump directly to the search engine's pricing problem – to maximize the function $\Pi^{SE} = m(p)D(p)(p - c)$ with respect to price. Plugging in the specific functions $m(p)$ and $D(p)$, the search engine's problem writes as:

$$\max_p \left\{\frac{1}{2(\Delta v)^2 \Delta\theta} \times [\bar{v} - p]^3 p\right\} \tag{11}$$

The solution to this, p^*_{SE}, is:

$$p^*_{SE} = \frac{\bar{v} + 3c}{4} \tag{12}$$

Search engine profits are:

$$\Pi^{SE}(p^*_{SE}) = \frac{27/512}{(\Delta v)^2 \Delta\theta}[\bar{v} - c]^4 \tag{13}$$

Setting $p^*_{SE} = p^*_{ad}(n)$, we can then back out the optimal number of advertisers for the search engine, n^*, which is:

$$n^* = 3 \tag{14}$$

Thus we see that for with uniformly distributed valuations and query costs, regardless of the supported intervals, the optimal number of advertisers is constant. This is of course consistent with the fact that, with uniform distributions, it is always the case that $\varepsilon_m/\varepsilon_D = 2$.

3 Endogenous Search Quality

So far, we have considered the users' query costs to exogenous. This is unrealistic, as the level of disutility one experiences when searching for a good is clearly affected by various characteristics of the search engine. In the framework of this model, it is convenient to consider a one-dimensional variable called "search quality", which, as it increases, lowers the query cost of each user.

Seen this way, it would, at first glance, always appear beneficial to the search engine to increase its quality, since doing so would simply increase the mass of users who search. In fact, this is not so obvious, since providing different quality of service may entail different levels of costs. For example, in order to ensure that search results consistently appear very quickly, the search engine may have to invest more in "computing clouds" in many different locations around the world.

In this section I add, to the baseline model, a quality decision to be made by the search engine. In doing so, I assume that the trade-off is straightforward. On the one hand increasing quality reduces the query cost each user incurs if he searches. On the other hand, in order to increase the search engine must undertake a costly investment.

The key lesson to be derived from this section is the following. Even with this simple trade-off, the relationship between quality and pricing is quite subtle. To see this, imagine that a there is an exogenous increase in the search engine's quality, i.e. that there is a decrease in each user's query cost. How should we expect the search engine to respond in its pricing decision?

One intuition says that since this decrease in users' query costs leads to an increase in overall demand for the good, the search engine should let in fewer advertisers and thus induce a higher price. However, another intuition could lead one to believe the opposite. The decrease in query costs could make users *more price-sensitive* in their decision whether or not to search; consequently, the search engine should adjust by letting in more advertisers and thereby bringing down the price of the good.

Which one of these intuitions is correct? It turns out that either one can be, depending, in particular, on how strong the effect of a change in search quality is on users of different types. The main result of this section, given in proposition 2, pins down the precise condition on which this issue hinges.

3.1 The Model

Technology. The search engine can increase its quality, denoted by s, at a cost, $h(s,\alpha)$, where α is a cost parameter. I assume $\frac{\partial h}{\partial s} > 0$, $\frac{\partial^2 h}{\partial s^2} \geq 0$ and $\frac{\partial^2 h}{\partial \alpha \partial s} > 0$.

Preferences. Since increased quality reduces each user's query cost, we can now express user i's utility as:

$$u_i = \begin{cases} v_i - p - \varphi(\theta_i, s), & \text{if searches and buys good at price } p \\ -\varphi(\theta_i, s), & \text{if searches and does not buy good} \\ 0, & \text{if does not search} \end{cases}$$

Under this specification, user i's type is still drawn from an interval $[\underline{\theta}, \bar{\theta}]$ with cdf $F(\cdot)$. Now, however, her query cost also depends negatively on search quality. So we make the following assumptions on the partial derivatives of φ:

$$\frac{\partial\varphi}{\partial\theta_i} > 0 > \frac{\partial\varphi}{\partial s}$$

Timing. The timing of the game is essentially the same as in the baseline model. The one important difference is that here, in the first stage, the search engine selects quality, s, in addition to the advertising fee. Thus we have:

1. Nature assigns a type, θ_i, to each user, from distribution $F(\cdot)$.
2. The search engine sets the advertising fee, A and the quality, s.
3. Merchants choose whether to advertise, and users choose whether to search.
4. Each user who searches learns her valuation for the good, v_i, from distribution $D(\cdot)$.
5. The n merchants who advertise compete *à la* Cournot, to sell to the users who searched.
6. Users who searched either buy or don't buy the good from an advertising merchant.

3.2 Equilibrium Results

As in the previous section, the model can most easily be understood by thinking of the search engine as directly setting the price of the good. Here, in addition to choosing a price, the search engine must choose an optimal level of search quality. Its profit maximization problem can thus be written:

$$\max_{s,p} \left[\Pi^{SE} = m(s,p)D(p)[p-c] - h(s,\alpha)\right] \tag{15}$$

This differs from the previous section's expression for search engine profits (1) in two ways. First, the mass of users who search, m, is both a decreasing function of the equilibrium price of the good and an increasing function of the chosen search quality.[3] Second, the search engine incurs higher costs, given by h, the higher the search quality it sets.

Initially disregarding the choice of quality, it is apparent that the first-order condition of this problem with respect to price yields a dual inverse-elasticity pricing rule analogous to the one given by proposition 1:

$$\frac{p^* - c}{p^*} = \frac{1}{\varepsilon_m + \varepsilon_D}$$

[3]Note that $\frac{\partial m(s,p)}{\partial p} > 0$ is implied by the assumption that $\frac{\partial \varphi(\theta_i,s)}{\partial s} < 0$.

where, now, the price-elasticity of the mass of users involves a partial derivative:

$$\varepsilon_m \equiv -\frac{\partial m(s^*, p^*)}{\partial p^*} \frac{p^*}{m(s^*, p^*)}$$

Here, as in the previous section, users take into account the expected price of the good when they decide whether or not to search. As a result, the search engine brings in competing advertisers and thus induces price that is below the "industry optimum" for a given mass of demand. The crucial issue that arises when the search engine must also choose a level of search quality is the fact that *this quality choice can affect users' sensitivity to price when deciding whether or not to search.* This feature of the model is illustrated by proposition 3a and in the discussion thereafter.

Proposition 3a. *As technology improves and the search engine's optimal choice of quality, s^*, increases, then the optimal choice of price, p^*, varies according to the following condition:*

$$p^* \text{ increases if } \varepsilon_m > \varepsilon_{\frac{\partial m}{\partial s^*}} \text{ and decreases if } \varepsilon_{\frac{\partial m}{\partial s^*}} > \varepsilon_m$$

where $\varepsilon_{\frac{\partial m}{\partial s^*}} \equiv -\frac{\partial}{\partial p^*}\left\{\frac{\partial m(s^*,p^*)}{\partial s^*}\right\}\frac{p^*}{\partial m(s^*,p^*)/\partial s^*}$.

Proof: See Appendix.

3.3 Example with Price Independent of Quality.

In light of this result, we first examine an example that falls into the simplest possible case, where these two elasticities are equal to one another: $\varepsilon_m = \varepsilon_{\frac{\partial m}{\partial s^*}}$. Assume that users' private values are distributed according to $\theta_i \sim \mathrm{U}[0,1]$ and $v \sim \mathrm{U}[0,1]$. Moreover, let $c = 0$ and let the cost of producing quality be given by:

$$h = \alpha s^2/2$$

Finally, let the users' query cost function, φ, be given by:

$$\varphi(\theta_i, s) = \frac{\theta_i}{s^\gamma}$$

where $\gamma \in (0,1)$ is an exogenous parameter reflecting the rate at which an increase in quality reduces users' query costs. Under these assumptions, the mass of users, m, takes the form:

$$m(s,p) = \frac{s^\gamma (1-p)^2}{2} \tag{16}$$

which gives:[4]

$$\varepsilon_m = \varepsilon_{\frac{\partial m}{\partial s^*}} = \frac{2p}{1-p} \tag{17}$$

Turning to the search engine's profit maximization problem, it can be written as:

$$\max_{s,p} \left[\Pi^{SE} = s^\gamma (1-p)^3 p - \alpha \frac{s^2}{2} \right] \tag{18}$$

From this we can clearly see that the tradeoff regarding the pricing decision is unaffected by concerns of quality and is in fact identical to the one we have seen in the previous section. Thus for the search engine, the optimal price to induce, p^*, and the optimal number of advertisers to let in, n^*, are:

$$p^* = \frac{1}{4} \quad \text{and} \quad n^* = 3 \tag{19}$$

Meanwhile, the quality decision is a bit less straightforward. The optimal quality level, s^* is given, as a function of the parameters, by:

$$s^*(\gamma, \alpha) = \left[\frac{\gamma}{\alpha} \frac{27}{256} \right]^{\frac{1}{2-\gamma}} \tag{20}$$

As one could easily have guessed, s^* is decreasing in α, the cost of producing quality. The behavior of s^* in response to changes in γ is somewhat more complicated – for small enough values of α, s^* is increasing in γ, while for larger values of α, s^*, plotted as a function of γ takes an inverted u-shape. Since I do not know of an economically meaningful explanation for this behavior of s^* in response to γ, I will refrain from further discussion. In the next section, where non-advertising merchants on the "left side" of the page compete with advertisers on the "right side", I give an example that can naturally be compared to this one.

3.4 Interpretation in Terms of Query Costs

Here, to further develop the intuition for the relationship between optimal price and quality, I derive a result similar to that of proposition 3a, in terms of users' query costs. Recall that a given user i chooses to search if and only if the expected net utility from potentially buying the good is at least as great as the user's query cost:

$$\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p} \geq \varphi(\theta_i, s)$$

[4]Note that the fact that $\varepsilon_m = \varepsilon_{\frac{\partial m}{\partial s^*}}$ does not depend on the specific form of the demand function, $D(p)$. Indeed, generally, with $\varphi(\theta_i, s) = \theta_i / s^\gamma$, we have $\varepsilon_m = \varepsilon_{\frac{\partial m}{\partial s^*}} = \varepsilon_{\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}}$.

Let us now define the function $\varphi^{-1}(\cdot,\cdot)$ such that:

$$\varphi^{-1}\left(\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}, s\right) = \theta^*$$

where θ^* is the type of the "threshold user" who is indifferent between searching and not searching, for a price-quality pair. Note that the assumptions on the query cost function, $\varphi(\theta_i, s)$, imply that $\varphi^{-1}(\cdot,\cdot)$ is decreasing in price[5] and increasing in quality.

Using this notation, we can write the mass of users, $m(s,p)$ as:

$$m(s,p) = F\left(\varphi^{-1}\left(\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}, s\right)\right)$$

For the rest of this subsection, let us assume that the types are uniformly distributed over some positive interval: $\theta_i \sim \mathrm{U}\left[\underline{\theta}, \overline{\theta}\right]$. We can then derive the result given by proposition 3b.

Proposition 3b. *When user types, θ_i, are uniformly distributed, as technology improves and the search engine's optimal choice of quality, s^*, increases, then the optimal choice of price, p^*, varies according to the following condition:*

$$p^* \text{ increases if } \varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial\varphi}{\partial\theta}} < 0 \text{ and decreases if } \varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial\varphi}{\partial\theta}} > 0$$

where $\varepsilon_{\varphi^{-1}} \equiv -\frac{\partial\varphi^{-1}}{\partial s^*}\frac{s^*}{\varphi^{-1}}$ *and where* $\varepsilon_{\frac{\partial\varphi}{\partial\theta}} \equiv -\frac{\partial}{\partial s^*}\left\{\frac{\partial\varphi}{\theta}\right\}\frac{s^*}{\partial\varphi/\partial\theta}$.

Proof: See Appendix.

A Specific Query Cost Function. From proposition 3b, we can see that as quality becomes cheaper to produce (i.e. as α decreases), the price of the good that the search engine induces may increase or decrease, depending on the form of the query cost function. This example is designed to illustrate under which type of conditions the price moves in each direction. As this example shows, an important issue is whether an increase in search quality reduces the query cost equally for all users, or whether users of one type are affected more strongly than users of another type.

Let us assume that the query cost function, $\varphi(\theta_i, s)$ is given by:

$$\varphi(\theta_i, s) = \frac{\theta_i - s^{1-\gamma}}{s^\gamma}$$

where $\gamma \in (0,1]$ is an exogenous parameter. Under this specification, for low values of γ, an increase in search quality s brings about a decrease in query costs that is relatively uniform

[5] To be more precise, $\varphi^{-1}(\cdot,\cdot)$ is increasing in the first argument, $\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}$, which itself is decreasing in price.

for users of all types. Meanwhile, for high values of γ, the impact of an increase in s causes a decrease in the cost of querying that is much more sizeable for high-type users than for low-type users.

A natural interpretation is to consider high-type users to be “novice” searchers and low-type users to be “expert” searchers. Thus, if γ is high, then technology is such that increasing s reduces the difficulty of using the search engine for those who are less “internet-savvy” but has relatively little impact on more experienced internet surfers. On the other hand, if γ is low, then expert users reap the benefits of increase in s nearly as much as novice users.

Another interpretation is to consider the user’s type, θ_i, to be her degree of unfamiliarity with the good. Under this interpretation as well, a user of a higher type must perform more “work” than a user of a lower type, in order to learn her valuation. However, seen in this way, the type does not correspond to a user’s level of proficiency in using the search engine, but rather to the ease with which she can judge the good’s potential usefulness in her particular circumstances.

With this particular form of the query cost function, φ, we have the following expression:

$$\varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial\varphi}{\partial\theta}} = \frac{2\gamma - 1}{s^{2\gamma-1}\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p} + 1} \tag{21}$$

As the denominator is clearly positive, the sign of this expression depends wholly on whether γ is greater than or less than $1/2$. Thus we see that for values of $\gamma < 1/2$ – cases where an increase in quality has a relatively uniform effect on users’ query costs – a technological improvement (decrease in α) causes an increase in the price of the good. On the other hand, for $\gamma > 1/2$ – cases where an increase in quality has a more pronounced effect on high-type users – when technology improves, the price of the good drops.

3.5 Example: Price and Quality Strategic Complements.

We look at a case where an increase in quality uniformly reduces users’ query costs and thus the values of the search engine’s optimal price and quality move in the same direction as one another. Let us assume that both sets of private values are uniformly distributed over the unit interval: $\theta_i \sim \mathrm{U}[0,1]$ and $v \sim \mathrm{U}[0,1]$. Furthermore, assume that merchants produce the good at zero marginal cost, so $c = 0$. For the search engine, let the cost of producing a quality level s be given by:

$$h = s^2/2$$

(here, for simplicity, let us ignore the technology parameter, α). Finally, let the users' query cost, φ, be given by:

$$\varphi(\theta_i, s) = \theta_i - s$$

Note that this corresponds to the extreme case, where $\gamma = 0$, of the query costs function discussed in the previous subsection. Since, in this case, a change in quality affects all users' query cost in a uniform manner, we are in a setting where quality can be seen as a tool whose use encourages the search engine to induce a higher price for the good. This aspect can be further appreciated by examining the search engine's profit maximization problem, in terms of s and p, which writes as:

$$\max_{s,p} \left[\Pi^{SE} = \underbrace{\left(s + \frac{[1-p]^2}{2} \right)}_{m(s,p)} (1-p)p - \frac{s^2}{2} \right] \tag{22}$$

The solutions to this problem, p^* and s^*, are given by:

$$p^* = \frac{1}{3} \quad \text{and} \quad s^* = \frac{2}{9} \tag{23}$$

For comparison, note that, here, when $s = 0$, we are in a special case of the example given at the end of the previous section, where query costs were exogenous and maximization was done only with respect to price. Recall that in that section, the search engine's optimal price was given by $\frac{v+3c}{4}$ which equals 1/4 under our current assumptions. The price in the exogenous case is less than the search engine's optimal price, here, when it can also affect things via changes in quality.

Turning our attention to the number of advertisers, here we have:

$$n^* = 2 \tag{24}$$

It is thus clear that, in this case, the search engine relies less on competing merchants to attract users than the case of exogenous query costs. Instead, of needing to lure in all searches by the promise of positive expected utility from the potential purchase of the good, the search engine effectively is able to "pay" users to search by offering them a more enjoyable search experience. This ability then makes it optimal for the search engine to make the market for the good less competitive.

Indeed, under this setup, there is a portion of users (those who's θ_i is on the interval $\left[0, \frac{2}{9}\right]$)

who, in equilibrium, have negative query costs. In other words, there is a set of users for who search quality is high enough they would query even if, for some reason, they were told in advance that it would be impossible for them to purchase the good.

4 Competition from the Left Side

In the previous section, I have considered "search quality" to be something on which the search engine must spend resources in order to produce but which is, in principle, independent from the number of merchants who appear on the search results page. In this section, I question this assumption. Here, I take into account the fact that since a typical search results page has both "sponsored" links and "organic" links, there is clearly the potential for the latter to steal business away from the former.

This potential could come in both direct and indirect ways. This type of competition could come directly through the presence of links to the websites of merchants, among the organic, "left side" search results. Quite obviously, if there are such unpaid links to merchants' websites, then a user has less incentive to pay attention to or to click on the paid advertisements on the right "right side" of the page.

At first glance, this sort of direct competition may seem unlikely. This is because of the fact that it would appear to be better for the search engine to favor, among the organic results, links to informational websites, rather than links to non-advertising merchants. However, a similar form of competition can also arise indirectly, *via such informational websites.* Since the owner of a website is able to perceive the fact that he is receiving hits from people searching for a given search term, any site that receives a significant number of hits from internet users searching for a commercial good has incentive to take advantage of this.

Moreover, taking advantage of this commercial opportunity is not likely to be very difficult. On the one hand, the owner of a website with significant informational value regarding a particular good may well have the ability to produce and sell the good as well. Even if this is not the case, the website owner can place an ad link to another merchant who does sell this good. As a result, even if the links on the left side of a search results appear to be informational, they are still in a position that gives them terrific potential to take traffic – and demand – away from the paid advertisements on the right side.

4.1 The Model

To capture this issue, in this section I modify the model presented in the previous section. Here, instead of supposing that the search engine must bear explicit costs in order to reach a certain level of search quality, I suppose that adding quality has an implicit cost. Specifically, I suppose that a higher level of search quality implies that the search results page includes a higher number of relevant unpaid links; futhermore, I suppose that the owners of the websites towards which these links are directed take advantage of the situation, either by selling the good themselves or by advertising directly from their sites to another merchant.

Technology and Preferences. In this version of the model, the users' preferences remain unchanged from those described in section 3.1. The explicit cost to the search engine of producing quality is zero: $h = 0$. The key assumption of this section is that quality, s, and the number of "left-side" merchants, who do not pay to advertise vary in the same direction. More precisely, I suppose that there are l non-advertising merchants who compete in the final stage of the game, and that this number is a function of the search engine's quality choice:

$$l = l(s)$$

where $l'(s) > 0$.

I also assume that the set of n merchants who compete in the final stage of the game is composed of these l merchants as well as r merchants who pay fee A to the search engine in order to advertise on the right side. Thus we have:

$$n = l(s) + r$$

Timing.

1. Nature assigns a type, θ_i, to each user, from distribution $F(\cdot)$.
2. The search engine sets the advertising fee, A and the quality, s.

   Links to $l(s)$ merchants are placed on the left side.
3. Merchants choose whether to advertise, and users choose whether to search.
4. Each user who searches learns her valuation for the good, v_i, from distribution $D(\cdot)$.
5. The $n = l + r$ merchants who appear compete *à la* Cournot, to sell to the users who searched.
6. Users who searched either buy or don't buy the good from an advertising merchant.

4.2 Equilibrium Results

As in the previous sections, I focus directly on the search engine's profit maximization problem. One way to see this problem is with respect to the variables that the search engine chooses directly, the search quality, s, and the advertising fee, A. As such it can be written:

$$\max_{A,s} \left[\Pi^{SE} = r(A,s)m(s,p)D(p)(p-c)\right] \tag{25}$$

where p is a function of the number of merchants on both sides:

$$p = p(l(s) + r(A,s))$$

From this expression, we see that under the assumptions of this section, the search engine's profits are equal to the number of advertising merchants, r, multiplied by the rent each individual merchant (of either side) extracts from users in the final stage.

As before, the problem can be made more palatable by looking at things as though the search engine were directly maximizing over price. In order to do this, observe that the proportion of total "industry" profits that the search engine gets are equal to the proportion of total merchants who pay to advertise. Thus define the variable ν such that:

$$\nu \equiv \frac{l}{l+r}$$

Since we have assumed that the function $l(s)$ is a one-to-one correspondence, we can define a new function for the mass of users who search, $\tilde{m}$, where:

$$\tilde{m} = \tilde{m}(l,p) = \tilde{m}(\nu n(p), p)$$

We are now in a position to rewrite the search engine's profit maximization program as the simpler problem:

$$\max_{\nu,p} \left[\Pi^{SE} = (1-\nu)\tilde{m}(\nu n(p), p)D(p)(p-c)\right] \tag{26}$$

Here the search engine can be seen to jointly select its desired final-stage price, p^*, and its desired proportion of the rent from the final stage, given by $(1-\nu^*)$.

Taking the first-order condition of this profit function with respect to price gives the familiar equation:

$$\frac{p^* - c}{p^*} = \frac{1}{\varepsilon_{\tilde{m}}^{tot} + \varepsilon_D} \tag{27}$$

where, now, the first term in the denominator, $\varepsilon_{\tilde{m}}^{tot}$, is the "total" price-elasticity of the mass of users, defined as:

$$\varepsilon_{\tilde{m}}^{tot} \equiv -\frac{d\left[\tilde{m}(\nu n(p^*), p^*)\right]}{dp^*} \frac{p^*}{\tilde{m}(\nu n(p^*), p^*)}$$

Meanwhile, the first-order condition with respect to ν gives:

$$\frac{\nu^*}{1 - \nu^*} = \varepsilon_{\tilde{m}}^l \tag{28}$$

where

$$\varepsilon_{\tilde{m}}^l \equiv \frac{\partial m}{\partial\{\nu^* n(p)\}} \frac{\nu^* n(p)}{\tilde{m}(\nu^* n(p), p)}$$

Combing the equations given by the two first-order condition allows us to derive proposition 4, which draws attention to the crucial new ingredient added in this section.

Proposition 4. *When increasing search quality implies displaying more unpaid, competing links on the results page, the optimal price-quality combination for the search engine to induce must satisfy:*

$$\frac{p^* - c}{p^*} = \frac{1}{\frac{\nu^*}{1-\nu^*}\varepsilon_n + \varepsilon_{\tilde{m}} + \varepsilon_D} \tag{29}$$

where

$$\varepsilon_n \equiv -\frac{dn}{dp^*}\frac{p^*}{n(p^*)}\ , \quad \varepsilon_{\tilde{m}} \equiv -\frac{\partial \tilde{m}}{\partial p^*}\frac{p^*}{m(\nu^* n(p^*), p^*)} \quad \text{and} \quad \varepsilon_D \equiv -\frac{dD}{dp^*}\frac{p^*}{D(p^*)}$$

Proof: See Appendix.

From proposition 4, we see that there is an important new component facing the search engine when choosing the price level to induce. Here, the profit margin induced by the search engine is limited *not only* by the direct influence of price on users' expected utility from potentially buying the good – via $\varepsilon_{\tilde{m}}$ – *but also* by the indirect influence of price on users' query costs. This is because of the fact that, for a given ν, a higher price corresponds to a lower number of left side merchants. Thus, as price increases, search quality decreases and the mass of users who search decreases. As the coefficient weighting ε_n indicates, the importance of this indirect effect is proportional to the ratio non-advertising of merchants to advertising merchants, since only the former contribute to lowering users' query costs. The following corollary provides another way of seeing the result of proposition 4.

Corollary to Proposition 4. *When increasing search quality implies displaying more unpaid, competing links on the results page, the optimal price-quality combination for the search engine to induce must also satisfy:*

$$\frac{p^* - c}{p^*} = \frac{1 - 2\nu^*}{(1 - \nu^*)\varepsilon_{\tilde{m}} + \nu^* \varepsilon_{D'} + \varepsilon_D} \tag{30}$$

where $\varepsilon_{D'} \equiv \frac{D''(p^*)p^*}{D'(p^*)}$.

Proof: See Appendix.

One immediate insight that this expression gives is that whenever the denominator is positive (which would be implied, for instance, by the assumption that $D'' \leq 0$), then the number of merchants appearing on the left side can be no greater than the number appearing on the right side. This is because such an event would imply $\nu^* > 1/2$, which would, in turn imply $p^* < c$, which clearly cannot hold in equilibrium.

4.3 Movement of Price and Quality

In section 3, I show that, in that section's model, price and quality vary in the same direction, in response to technological change, if and only if the direct influence of a price change on the number of users who search is greater than its indirect influence through the marginal effect of quality. Regarding this section's model, I have not been able to derive such a clean result; however, it appears that, in most situations, price and quality move in opposite directions.

From the first-order condition, with respect to ν of the search engine's profit maximization problem, we have:

$$\nu^* n(p^*) = n(p^*) - \underbrace{\frac{m(\nu^* n(p^*), p^*)}{\frac{\partial m}{\partial \nu^* n(p^*)}}}_{\Psi} \tag{31}$$

Differentiating both sides with respect to p^* and noting that the left-hand side of the above equation varies in the same direction as quality, we have that price and quality move in opposite directions whenever:

$$n'(p^*) < \frac{d\Psi}{dp^*} \tag{32}$$

A sufficient but not necessary condition for this to hold is for the following inequality to be satisfied:

$$\frac{\partial^2 m}{\partial l^2} \nu^* n'(p^*) + \frac{\partial^2 m}{\partial l \partial p} \geq 0 \tag{33}$$

The intuition for why price and quality should vary in opposite directions is straightforward and is as follows. An increase in price implies that the total number of merchants decreases. So, for quality to increase at the same time as price, the *proportion* of merchants placed on the left side would have to increase by enough to offset the loss in the overall number of merchants. Thus, for "well-behaved" query cost functions and smooth distributions of types, we should expect price and quality to vary in opposite directions. The next subsection gives an example that further illustrates this feature of the model.

4.4 Example: Left and Right Side Merchants

Here we look a simple example that can be easily compared with the one found in section 3.3. The setup is in this example is precisely the same except for the fact that, here, the provision of quality, s, implies the presence of left-side merchants, *l(s)*. For simplicity, let us assume:

$$l(s) = s$$

As before, private values θ_i and v_i are independently, uniformly distributed over the unit interval and merchants' unit cost is assumed to be zero.

Using the fact that in section 3.3:

$$m(s,p) = \frac{s^{\gamma}(1-p)^2}{2}$$

In addition, using the fact that final-stage Cournot competition gives:

$$\frac{p-c}{p} = \frac{1}{n} \cdot \frac{1}{\varepsilon_D}$$

We can thus write the mass of users directly in terms of ν and p:

$$m(s,p) = \tilde{m}(\nu n, p) = \left(\nu \cdot \frac{1-p}{p}\right)^{\gamma} \frac{(1-p)^2}{2} \tag{34}$$

Hence, the search engine's maximization problem can be written as:

$$\max_{\nu,p} \left[\Pi^{SE} = (1-\nu)\nu^{\gamma}(1-p)^{3+\gamma}p^{1-\gamma}\right] \tag{35}$$

Solving this problem, the optimal price to induce, p^*, and the optimal total number of merchants to allow, n^* are given, as functions of γ by:

$$p^*(\gamma) = \frac{1-\gamma}{4} \quad \text{and} \quad n^*(\gamma) = \frac{3+\gamma}{1-\gamma} \tag{36}$$

First, note that, here, the search engine induces a more competitive overall environment than in the section 3.3 example, as the price is strictly smaller than 1/4 and the number of merchants strictly greater than 3. Moreover, note that the intensity of competition is increasing in γ. This can be interpreted in two ways. One way involves thinking of γ as a measure of users' sensitivity to quality. The other way entails thinking of γ as an inverse measure of the rate at which merchants appear on the left side as quality increases.[6] Viewing things in either way, when γ is high, the search engine places relatively more weight on drawing in users by providing high quality compared to when γ is low.

Turning our attention to the choice of search quality, the optimal proportion of merchants placed on the left side, ν^*, is:

$$\nu^*(\gamma) = \frac{\gamma}{1+\gamma} \tag{37}$$

Thus, the optimal search quality, s^*, (which, here, is the same quantity as the optimal number of left-side merchants) is given by:

$$s^*(\gamma) = \frac{\gamma(3+\gamma)}{1-\gamma^2} \tag{38}$$

Correspondingly, the optimal number of right-side merchants, r^*, is:

$$r^*(\gamma) = \frac{3+\gamma}{1-\gamma^2} \tag{39}$$

A striking feature of this example compared to the one in section 3.3 is the dramatic increase in sensitivity to the parameter γ. In that example, the optimal number of merchants from the search engine's perspective was always 3. Here, on the other hand, by introducing a positive correlation between search quality and the number of left-side merchants, the search engine's optimal number of final-stage participants can potentially be any positive number at least this great. The key point is that provided users are sufficiently sensitive to search quality *and* provided that increasing this opens the door to non-paying merchants, the incentive for the search engine to induce a high markup can disappear completely.

[6] This latter interpretation is supported by the fact that one finds the same results when one lets $l(s) = s^{1/\gamma}$ and $\varphi = \theta_i/s$

5 Conclusion

This paper begins by hypothesizing that a key to understanding search engines lies in understanding the ways in which they differ from, more generally, other platforms, and more specifically, traditional "directory services". In the body of the paper I develop three versions of an original model. The first is designed to represent a traditional directory service, while the second and third are designed to show how a search engine is different. The central message is that search engines' enormous commercial success owes to the fact that they are designed to balance "local" commercial considerations with "global" ones. In doing so, they have an incentive, on the one hand, to offer results which are not as relevant as they possibly could be, even if improving them would come at no technical cost. On the other hand, they have incentive to expose merchants who pay for advertisement to competition from outside competitors who offer the search engine no direct compensation.

The stylized features of this model leave more than ample room for a variety of interesting extensions. In particular, this paper is populated by a homogenous set of merchants whose main strategic decision is whether or not to enter according to a zero-profit condition. A natural step forward would be include heterogeneous merchants, thereby connecting this framework to the literature on search engine auctions. In the same spirit, it would also be interesting to consider situations in which individual potential advertisers could themselves potentially appear among the organic links. Such scenarios create, *a priori*, a complex set of imaginable outcomes: on the one hand, merchants would seem to have less incentive to advertise were possible to appear on the search results page without paying. On the other hand, the search engine may have a conflicting incentive to combat this by "blacklisting" merchants who refuse to buy sponsored links from showing up in the unpaid results.

Finally, the search engine could serve as a very nice example for theorists of organizations. From this standpoint, the model in this paper can be seen as akin to "team theory", since the search engine's unique goal, when performing its ad pricing and setting its left side quality is to maximize overall profits. It would be very interesting to study in detail the divisional structures within the search engine operators and to examine the task-overlap and the alignment of incentives between those in charge of algorithm-related issues and those in charge of the ad auctions.

6 Appendix

Proposition 1. *The problem facing the search engine when choosing its advertising fee, A, can be rewritten as a problem of choosing the equilibrium price, p, of the good. This problem can be written as:*

$$\max_p \left[\Pi^{SE} = m(p)D(p)[p-c]\right] \tag{40}$$

where $m(\cdot)$ is the mass of users who search as a function of the price. Moreover, the expression for the mass of users, $m(p)$, is given by:

$$m(p) = F\left(\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right) \tag{41}$$

Proof: By choosing an appropriate fee, A, the search engine can induce any positive number of merchants to advertise. Since, as the number of merchants grows large, industry profits approach zero, the search engine can induce any price in the interval $(c, p^m]$, where p^m is the price that maximizes $D(p)[p-c]$. Since the price that maximizes search engine profits is always in this interval, we can write the problem unconstrained as above.

To show the second part, we note that user i chooses to search if and only if:

$$[E[v|\text{buy}] - p] \times \Pr[\text{buy}] \geq \theta_i \tag{42}$$

The user's valuation is drawn from distribution $G(\cdot)$, and we define the demand function as $D(\cdot) \equiv 1 - G(\cdot)$. The user's expected valuation for the good, given that he buys – i.e. given that his valuation is higher than the price is given by:

$$\frac{\int_p^{\bar{v}} G'(\tilde{p})\tilde{p}\, d\tilde{p}}{1 - G(p)} = \frac{-\int_p^{\bar{v}} D'(\tilde{p})\tilde{p}\, d\tilde{p}}{D(p)} \tag{43}$$

The numerator can be simplified using integration by parts. Differentiating $\tilde{p}$ and integrating $D'(\tilde{p})d\tilde{p}$ gives:

$$-\int_p^{\bar{v}} D'(\tilde{p})\tilde{p}\, d\tilde{p} = -\left[\bar{v}\underbrace{D(\bar{v})}_{0} - pD(p) - \int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right] = pD(p) + \int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p} \tag{44}$$

So, applying the result from the integration by parts to the expression in (43) gives:

$$\frac{-\int_p^{\bar{v}} D'(\tilde{p})\tilde{p}\, d\tilde{p}}{D(p)} = p + \frac{\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}}{D(p)} \tag{45}$$

Using (44) and the definition of $D(\cdot)$, we can rewrite the condition in (42) as:

$$\left(p + \frac{\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}}{D(p)} - p\right) D(p) = \int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p} \geq \theta_i \tag{46}$$

Therefore we have shown that $m(p) = F\left(\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right)$. *Q.E.D.*

Proposition 2. *The equilibrium price of the good, p^*, induced search engine, must satisfy the following dual inverse-elasticy rule:*

$$\frac{p^* - c}{p^*} = \frac{1}{\varepsilon_m + \varepsilon_D} \tag{47}$$

where ε_m is the elasticity of the mass of users who search with respect to price, $\varepsilon_m = \frac{dm(p^*)}{dp^*}\frac{p^*}{m(p^*)}$, and where ε_D is the final-stage elasticity of demand with respect to price, $\varepsilon_D = \frac{dD(p^*)}{dp^*}\frac{p^*}{D(p^*)}$.

The expression for ε_m is given by:

$$\varepsilon_m = \frac{f\left(\int_{p^*}^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right)}{F\left(\int_{p^*}^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right)} D(p^*)p^* \tag{48}$$

Proof: Log-differentiation of the search engine's maximization problem stated in proposition 1 yields the first-order condition:

$$\frac{m'(p^*)}{m(p^*)} + \frac{D'(p^*)}{D(p^*)} + \frac{1}{p^* - c} = 0 \tag{49}$$

Multiplying this equation by p^* and rearranging gives the expression in (47).

Since $m(p) = F\left(\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right)$, we have, using Leibnitz's Rule, that:

$$m'(p) = f\left(\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}\right) \cdot [-D(p^*)] \tag{50}$$

Therefore have the expression given in (48). *Q.E.D.*

Corollary to Proposition 2. *The search engine chooses the number of advertisers, n^*, to satisfy:*

$$n^* = \frac{\varepsilon_m}{\varepsilon_D} + 1 \tag{51}$$

Proof: We combine the result of proposition 2 with the standard result in n-player Cournot that:

$$\frac{p^* - c}{p^*} = \frac{1}{n} \cdot \frac{1}{\varepsilon_D} \tag{52}$$

Simplifying, we easily obtain the expression for n^* given by (51). *Q.E.D.*

Proposition 3a. *As technology improves and the search engine's optimal choice of quality, s^*, increases, then the optimal choice of price, p^*, varies according to the following condition:*

$$p^* \text{ increases if } \varepsilon_m > \varepsilon_{\frac{\partial m}{\partial s^*}} \text{ and decreases if } \varepsilon_{\frac{\partial m}{\partial s^*}} > \varepsilon_m$$

where $\varepsilon_{\frac{\partial m}{\partial s^*}} \equiv -\frac{\partial}{\partial p^*}\left\{\frac{\partial m(s^*,p^*)}{\partial s^*}\right\}\frac{p^*}{\partial m(s^*,p^*)/\partial s^*}$.

Proof: Let us define Ψ as:

$$\Psi(p,s) \equiv \frac{\partial \Pi^{SE}(p,s)}{\partial p} \tag{53}$$

The first-order condition with respect to price of the search engine's profit maximization problem is:

$$\Psi(p^*, s) = 0 \tag{54}$$

Writing the maximizing price, p^* as a function of s^* and totally differentiating, profit maximization requires that:

$$\frac{\partial \Psi(p^*(s^*), s^*)}{\partial p^*} \cdot \frac{dp^*(s^*)}{ds^*} + \frac{\partial \Psi(p^*(s^*), s^*)}{\partial s^*} = 0 \tag{55}$$

Since the second-order condition of the profit maximization problem requires that $\partial\Psi/\partial p^* < 0$, we have that:

$$sign\left\{\frac{dp^*(s^*)}{ds^*}\right\} = sign\left\{\frac{\partial \Psi(p^*, s^*)}{\partial s^*}\right\} \tag{56}$$

And:

$$\frac{\partial \Psi(p^*, s^*)}{\partial s^*} = \left(\frac{\partial^2 m}{\partial p^* \partial s^*} D(p^*) + D'(p^*) m(p^*, s^*)\right)(p-c) + \frac{\partial m}{\partial s^*} D(p^*) \tag{57}$$

Straightforward algebra yields that:

$$\frac{\partial \Psi(p^*, s^*)}{\partial s^*} > 0 \Leftrightarrow \frac{p^* - c}{p^*} > \frac{1}{\varepsilon_{\frac{\partial m}{\partial s^*}} + \varepsilon_D} \tag{58}$$

But, since:

$$\frac{p^* - c}{p^*} = \frac{1}{\varepsilon_m + \varepsilon_D} \tag{59}$$

We have that:

$$\frac{dp^*(s^*)}{ds^*} > 0 \text{ if and only if } \varepsilon_m > \varepsilon_{\frac{\partial m}{\partial s^*}} \tag{60}$$

Therefore, we have the result of proposition 3a. *Q.E.D.*

Proposition 3b. *When user types, θ_i, are uniformly distributed, as technology improves and the search engine's optimal choice of quality, s^*, increases, then the optimal choice of price, p^*, varies according to the following condition:*

$$p^* \text{ increases if } \varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial\varphi}{\partial\theta}} < 0 \text{ and decreases if } \varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial\varphi}{\partial\theta}} > 0$$

where $\varepsilon_{\varphi^{-1}} \equiv -\frac{\partial\varphi^{-1}}{\partial s^*}\frac{s^*}{\varphi^{-1}}$ *and where* $\varepsilon_{\frac{\partial\varphi}{\partial\theta}} \equiv -\frac{\partial}{\partial s^*}\left\{\frac{\partial\varphi}{\theta}\right\}\frac{s^*}{\partial\varphi/\partial\theta}$.

Proof: In view of the result shown in proposition 3a, it suffices to show, here, that when $\theta_i \sim \mathrm{U}\left[\underline{\theta}, \overline{\theta}\right]$, we have:

$$\varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial\varphi}{\partial\theta}} < 0 \Leftrightarrow \varepsilon_m > \varepsilon_{\frac{\partial m}{\partial s^*}} \quad \text{and} \quad \varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial\varphi}{\partial\theta}} > 0 \Leftrightarrow \varepsilon_{\frac{\partial m}{\partial s^*}} > \varepsilon_m \tag{61}$$

Generally, we have:

$$m(s,p) = F\left(\varphi^{-1}\left(\int_p^{\bar{v}} D(\tilde{p})\,d\tilde{p}, s\right)\right) \tag{62}$$

Hence, when $\theta_i \sim \mathrm{U}\left[\underline{\theta}, \overline{\theta}\right]$, we have:

$$m(s,p) = \frac{1}{\Delta\theta}\cdot\varphi^{-1}\left(\int_p^{\bar{v}} D(\tilde{p})\,d\tilde{p}, s\right) \tag{63}$$

$$\frac{\partial m}{\partial p} = \frac{1}{\Delta\theta}\cdot\frac{-D(p)}{\partial\varphi/\partial\theta} \tag{64}$$

$$\frac{\partial m}{\partial s} = \frac{1}{\Delta\theta}\cdot\frac{\partial\varphi^{-1}}{\partial s} \tag{65}$$

And:

$$\frac{\partial^2 m}{\partial p\partial s} = \frac{1}{\Delta\theta}\cdot D(p)\cdot\frac{\frac{\partial^2\varphi}{\partial\theta\partial s}}{\left(\frac{\partial\varphi}{\partial\theta}\right)^2} \tag{66}$$

Plugging these values into $\varepsilon_m >$ and $\varepsilon_{\frac{\partial m}{\partial s^*}}$ and rearranging, one easily obtains the desired result. *Q.E.D.*

Proposition 4. *When increasing search quality implies displaying more unpaid, competing links on the results page, the optimal price-quality combination for the search engine to induce must satisfy:*

$$\frac{p^* - c}{p^*} = \frac{1}{\frac{\nu^*}{1-\nu^*}\varepsilon_n + \varepsilon_{\tilde{m}} + \varepsilon_D} \tag{67}$$

where

$$\varepsilon_n \equiv -\frac{dn}{dp^*}\frac{p^*}{n(p^*)}\ , \quad \varepsilon_{\tilde{m}} \equiv -\frac{\partial \tilde{m}}{\partial p^*}\frac{p^*}{m(\nu^* n(p^*), p^*)} \quad \text{and} \quad \varepsilon_D \equiv -\frac{dD}{dp^*}\frac{p^*}{D(p^*)}$$

Proof: The first-order conditions of the search engine's profit maximization problem are, first, with respect to p:

$$\frac{p^* - c}{p^*} = \frac{1}{\varepsilon_{\tilde{m}}^{tot} + \varepsilon_D} \tag{68}$$

And, second, with respect to ν:

$$\frac{\nu^*}{1-\nu^*} = \varepsilon_{\tilde{m}}^{l} \tag{69}$$

To obtain result of this proposition, we must thus show that $\varepsilon_{\tilde{m}}^{tot} = \frac{\nu^*}{1-\nu^*}\varepsilon_n + \varepsilon_{\tilde{m}}$. Using the fact that:

$$\varepsilon_{\tilde{m}}^{tot} = -\left[\frac{\partial \tilde{m}}{\partial\{\nu^* n(p^*)\}}\nu^* n'(p^*) + \frac{\partial \tilde{m}}{\partial p^*}\right] \cdot \frac{p^*}{\tilde{m}} \tag{70}$$

An the fact that rearranging the first-order condition with respect to ν, (69), gives:

$$\frac{\partial \tilde{m}}{\partial\{\nu^* n(p^*)\}} = \frac{\tilde{m}(\nu^* n(p^*), p^*)}{(1-\nu)n(p^*)} \tag{71}$$

Inserting the right-hand side of (71) into (70) and simplifying gives the equation we are trying to show. *Q.E.D.*

Corollary to Proposition 4. *When increasing search quality implies displaying more unpaid, competing links on the results page, the optimal price-quality combination for the search engine to induce must also satisfy:*

$$\frac{p^* - c}{p^*} = \frac{1-2\nu^*}{(1-\nu^*)\varepsilon_{\tilde{m}} + \nu^*\varepsilon_{D'} + \varepsilon_D} \tag{72}$$

where $\varepsilon_{D'} \equiv \frac{D''(p^*)p^*}{D'(p^*)}$.

Proof: From proposition 4, we have that:

$$p^* - c = \frac{-1}{\frac{\nu^*}{1-\nu^*}\frac{n'(p^*)}{n(p^*)} + \frac{\partial \tilde{m}/\partial p^*}{\tilde{m}(\nu^* n(p^*), p^*)} + \frac{D'(p^*)}{D(p^*)}} \tag{73}$$

In addition, the equilibrium of the final stage n-merchant Cournot game entails:

$$\frac{p^* - c}{p^*} = \frac{1}{n}\cdot\frac{1}{\varepsilon_D} \Leftrightarrow n(p^*) = \frac{D(p^*)}{D'(p^*)(p^* - c)} \tag{74}$$

Log-differentiating the function $n(p)$ then yields:

$$\frac{n'(p^*)}{n(p^*)} = \frac{D'(p^*)}{D(p^*)} - \frac{D''(p^*)}{D'(p^*)} - \frac{1}{p^* - c} \tag{75}$$

Inserting the right-hand side of (75) into (73) and simplifying gives the desired result. *Q.E.D.*

7 References

Athey, Susan & Glenn Ellison (2007) “Position Auctions with Consumer Search”. Unpublished Working Paper.

Baye, Michael & John Morgan (2001) “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets”. *American Economic Review*, Vol. 91, No. 3, pp. 454-474.

Borgers, Tilman, Ingemar Cox, Martin Pesendorfer & Vaclav Petricek (2007) “Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence”. Unpublished Working Paper.

Edelman, Benjamin, Michael Ostrovsky & Michael Scharz (2007) “Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords”. *Amercian Economic Review*, Vol. 97, No. 1, pp. 242-259.

Hagiu, Andrei (2007) “Merchant or Two-Sided Platform”. *Review of Network Economics*, Vol. 6, No. 2.

Hagiu, Andrei & Bruno Jullien (2007) “Designing a Two-Sided Platform: When To Increase Search Costs?”. Unpublished Working Paper.

Rochet, Jean-Charles & Jean Tirole (2007) “Two-Sided Markets: A Progress Report”. *RAND Journal of Economics*.

Varian, Hal R. (2007) “Position Auctions”. *International Journal of Industrial Organization*.