

Search Engines: Left Side Quality versus Right Side Profits*

Alexander White[†]

Draft: December 2009
(First Version: May 2008)

Abstract

High quality search results have the potential to eat into a search engine's profits. However few people would use a lousy search engine. While there is clearly a tradeoff governing a search engine's quality choice, it is far from straightforward. A search engine must attract users, direct them to relevant websites, some of which are paid ads and some of which are not, while creating conditions that are profitable for the former.

This paper examines this tradeoff. It shows that the quality choice is closely tied with users' *post-search* interaction with advertisers. Quality serves as a demand booster for users, causing them to tolerate a search engine on which advertisers do not offer the lowest possible prices. It argues, however, that websites that do not *pay* to appear still have an incentive to compete in the same market as advertisers, and shows that if this is the case, then adding quality reduces the prices advertisers will in fact charge. Thus, even if raising quality were costless, in order to allow advertisers to charge a positive markup, the search engine must set quality below the highest possible level. Finally, the paper considers the merger of Google and DoubleClick in light of the model.

*Many thanks to Eduardo Azevedo, Klênio Barbosa, Stefan Behringer, Guido Friebel, Dominik Grafenhofer, Andrei Hagiu, Marti Hearst, Kamal Jain, Bruno Jullien, Dmytro Kylymnyuk, German Lambardi, Jérôme Mathis, Volker Nocke, Martin Peitz, Paul Seabright, Carl Shapiro, Joacim Tåg, Glen Weyl, and especially to Jacques Crémer for very helpful discussions and/or comments on previous drafts. I also thank participants at the ZEW ICT Conference, in July 2008, where this work was first presented, as well as seminar audiences at Toulouse School of Economics, UCL and participants in Meetings of the Econometric Society in Rio de Janeiro and Barcelona. I bear full responsibility for all errors.

[†]Toulouse School of Economics & Télécom ParisTech; alexander.white@univ-tlse1.fr

The goals of the advertising business model do not always correspond to providing quality search to users. [...] we expect that advertising funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers.

– Sergey Brin and Larry Page, founders of Google, before Google was advertising funded¹

1 Introduction

Google’s founders seem to believe that they face mixed motives. As Page and Brin suggest, there are costs and benefits to a search engine of showing users the best possible search results. However, precisely what these are is by no means obvious. This paper explores the way a search engine’s choice of quality fits in with its incentives to maximize profits.

Search engines allow web users to enter a “search term” or “query” and, in response, display a “results page”. A potential conflict of interest facing a search engine arises from the combination of two types of results that are displayed. One type are those typically found on the left side of the results page. These “organic” or “unpaid” results are links to other websites that the search engine selects using a complicated set of criteria, which I refer to as its “algorithm”. The other type of result, typically found on the right side, are paid advertisements, linking to websites that have offered to pay the search engine in order to appear when a user enters a particular term. Each time a user clicks on a right side link, the search engine receives a payment; when a user clicks on a left side link, the search engine receives nothing.

In light of this arrangement and for a broad set of other reasons, it is important to understand the forces determining whether a search engine chooses to provide users with high quality results. The vast majority of Internet users rely on search engines to get information, which can be on any topic imaginable. A dominant search engine is thus in a position to wield a great deal of influence. This position is made particularly strong by at least two factors. First, a search engine projects a rather false sense of objectivity. By primarily pointing to other websites, rather than providing its own content, a search engine may not seem explicitly to express opinion or bias, but the selection of results can implicitly incorporate such leanings.

Second, as search engines get more sophisticated, variation in their results gets harder to track and to interpret in a systematic way. Since many factors, such as a user’s location and recent search history affect which results are displayed, there is often some variation in the results that appear for searches for the same term, performed under nearly identical conditions. Thus, were a search engine to vary results in some way that was somehow

¹Brin and Page (1998).

“manipulative”, there would be no clear way to distinguish this from ordinary, *prima facie* helpful variation.

Given a search engine’s potential to be both influential and opaque,² the theoretical framework of this paper helps to clarify the forces driving the choice of search quality it provides to users. In particular, the paper illustrates the way that these forces relate to search advertisement. It does so with the help of a theoretical model developed over three sections. To begin, in section 2, I take the search engine’s quality to be exogenously given and focus specifically on the search engine’s advertising. Since the search engine connects users with advertisers selling a clearly defined type of good that is *pre-specified* via users’ queries, the search engine seems to wield a great deal of control over who gets to sell to the search engine’s users and at what price. I derive formulae for the prices the search engine would optimally *want* advertisers to charge for their goods and for the number of advertising positions it would like to have filled, and compare this market structure with other sorts of multi-sided platforms.

In the subsequent two sections, I take search quality to be set endogenously by the search engine. The tradeoff determining this quality choice takes two different forms in the two sections. In section 3, the search engine’s quality decision is very straightforward. In order to provide users with high search quality, the search engine must make a costly investment. I derive the optimal quality and optimal price for the advertisers’ good, from the standpoint of the search engine, and I pin down the condition determining when it considers these two are complements. Interestingly, this hinges on the form of user heterogeneity. In particular, price and quality tend to be complements except when a change in quality level has a disproportionate effect on the willingness to search of “novice” searchers.

Section 4 then focuses on a form of interaction that I argue is likely to occur between high quality organic results and advertisement. This version of the model considers the fact that relevant websites appearing among the organic links have incentive to participate in the same market as paying advertisers, and it examines the search engine’s optimal price and quality in this setting. Crucially, I find that in the presence of this interaction between organic results and advertisement, price and quality are much more likely to be substitutes for the search engine. Since, high quality of organic results mean advertisers face more competition, increasing quality puts downward pressure on prices.

Together, sections 3 and 4 show that the search engine’s incentives depend greatly on whether this interaction between organic links and advertisement in fact occurs. In the absence of this phenomenon, providing users with high quality allows the search engine

²For a wide-ranging, fascinating discussion of these issues, see the book, *The Search* by John Battelle (2005) as well his “searchblog”, at www.battellemedia.com.

to charge advertisers more by attracting more users. In this case, when users benefit from high quality search results they typically also face advertisers who charge high prices for their goods. By contrast, when this interaction between the left side and the right side is significant, high quality search results typically go hand in hand with low prices for advertised goods.

This result appears to be of significance to the intensifying regulatory concern over search engines, particularly regarding Google's increasing dominance. Section 5 considers this result in the context of the Google/DoubleClick merger. Significantly, it shows that when one takes into account the incentives for a search engine to adjust quality endogenously, an important aspect of the debate may be reversed. Finally, section 6 reviews related literature and section 7 concludes.

2 Baseline Case

I begin by considering the simplest case in order to illustrate the overarching problem the search engine faces. Search engines are platforms that make money by charging websites in exchange for connecting them with web surfers. At the same time, its service to users is two-fold: by showing to them a set of websites, as a function of their query,³ a search engine not only points to a set of potential trading partners, it also allows users to learn a great deal about the subject of their query terms. The model in this section captures these key characteristics.

Two features of the model are of particular note. First, the search engine has very precise control over how many merchants users can interact with. This represents the idea that, in choosing the links to be shown as well style and format of the results page, the search engine has significant influence over the degree of market power held by individual websites.⁴

Second, users learn their valuation for the good only *after* they have searched and incurred a "query cost". This reflects the fact that, by searching for an item and looking at the results, a user can "get up to speed" about the item and figure out whether or not it is of interest. In the model, query costs vary from one user to another. This can be interpreted in two ways: it can be thought to reflect the fact that different people are more or less skilled at using a search engine. Alternatively, it can be seen to capture the fact that people search

³In addition to using the terms in a user's query, search engines increasingly use additional information, such as the user's location and browsing history as the basis for the group of websites that are displayed on the results page. For more on the search engine's tradeoff between privacy versus personalization, see Krause and Horvitz (2008).

⁴This is in contrast with the Baye and Morgan (2001) model in which Bertrand competition among sellers limits the precision of the "gatekeeper's" control.

for goods in varying degrees of ignorance as to what the goods are, and that when one is more ignorant about a particular good, searching for it (and deciding how much its worth) requires more effort.

In this section, I take users' query costs to be exogenously given. In sections 3 and 4, which analyze the search engine's incentive to provide left side, or organic search results, query costs are endogenous.

2.1 A Simple Model

The model has three kinds of players: a "search engine", an arbitrarily large set of "merchants", and a continuum of "users" of mass one. In the game, the search engine sets an advertising fee, merchants decide whether to advertise and how much of a good to produce, and users decide whether to search and whether or not to buy the good.

Preferences. Each user i 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, user i 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 doesn't and gets a payoff of zero. I assume that merchants produce the good at a common marginal cost, c . So, merchant j who advertises receives profits

$$(p - c)q_j - A.$$

Here, q_j denotes the quantity j sells, and p denotes the price of the good. Advertising merchants engage in Cournot competition and thus the price of the good is a function of the total quantity merchants sell.

Search engine profits come exclusively from advertising and are given by

$$n \times A,$$

where n , the number of merchants who advertise, is determined by a “zero profits” condition.⁵ A crucial feature of the model is that, in setting A , the search engine affects not only the number of advertisers but also the number of users who search. In effect, the search engine is a monopolist in a particular kind of two-sided market, an issue that I discuss in the next subsection.

Timing. The stages of the game are as follows:

1. The search engine sets the advertising fee, A .
2. Merchants choose whether to advertise, and users choose whether to search.
3. Each user who searches learns her valuation for the good, v_i .
4. The n merchants who advertise compete *à la* Cournot, to sell to the users who searched.
5. Users who searched either buy or don’t buy the good from an advertising merchant.

Note that the assumption that merchants compete *à la* Cournot is not essential to the model. It is, rather, a convenient way to represent that idea that price decreases as the number of competing advertisers increases, but, that it does not drop to marginal cost as soon as there are two merchants as in Bertrand competition.⁶ The model can be readily written in a more general form in which one assumes only that price decreases with the number of advertisers. All of the following results carry through in such a model except for the analytic solution for the equilibrium number of firms. Also, throughout the rest of the paper, I ignore the constraint the number of firms should be an integer and instead treat this as a continuous quantity.

2.2 Equilibrium Analysis

I look for pure-strategy equilibria in which merchants act symmetrically and assume throughout that the search engine’s problem has a unique, interior solution.⁷ Define the function $D(\cdot) \equiv 1 - G(\cdot)$. In the fourth-stage Cournot game, n advertisers face demand given

⁵In this regard, the present model builds upon those of “free entry” with fixed costs. See, for instance, Mankiw and Whinston (1986).

⁶On this issue, see, in particular, Ellison and Ellison (2009), including the supplemental materials associated with that article, which discuss this issue and give a different, somewhat richer form of competition that could be used here to the same effect.

⁷To avoid distraction, I assume that all players have beliefs that are consistent with maximal participation. This eliminates, for instance, the uninteresting equilibrium in which $A = 0$, no merchants advertise and no users search. Note that the equilibrium I select leads to the outcome that would prevail if the search engine could set its merchant fee as a function of the mass of users that search. For a general treatment of such “insulating tariffs” in the context of a monopoly platform, see Weyl (forthcoming).

by $mD(\cdot)$, where m denotes the mass of users who searched. Since m is set in a prior stage, the equilibrium price of this Cournot game depends only on $D(\cdot)$. Thus, in any equilibrium, the price, p^* , must satisfy

$$\frac{p^* - c}{p^*} = \frac{1}{n \varepsilon_D},$$

where ε_D is the price-elasticity of demand. Let $p^*(n)$ denote the equilibrium price as a function of the number of advertisers. Then, when n merchants advertise and m users search, each merchant's profits are equal to

$$\frac{m}{n} D(p^*(n)) [p^*(n) - c] - A.$$

A key feature of the model is the search engine's control over the equilibrium price that emerges from the merchant competition. In setting the advertising fee, A , the search engine determines both the number of advertisers, n , and the mass of users, m . Consequently, we can view the search engine as if were solving a monopoly pricing problem. The following lemma establishes this.

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

$$\max_p \{m(p)D(p)[p - c]\}, \quad (1)$$

where $m(p)$, the mass of users who search is

$$m(p) = F\left(\int_p^{\bar{v}} D(\tilde{p}) d\tilde{p}\right). \quad (2)$$

Proof. See Appendix A.1. □

Setting the problem up in this way allows us to isolate the key issues facing search engine, illustrated by Proposition 1.

Proposition 1. *The equilibrium price of the good, p^* , induced by the search engine, must satisfy the dual inverse-elasticity rule,*

$$\frac{p^* - c}{p^*} = \frac{1}{\varepsilon_m + \varepsilon_D}, \quad (3)$$

where ε_m is the elasticity of the mass of users who search with respect to price.

Proof. See Appendix A.2. □

This proposition provides the basis for an interesting comparison between a search engine and what one might call a “conventional” two-sided market. In contrast to the latter, here, just one kind of agent – advertisers – benefits directly from the presence of more of the other kind of agent. Users benefit only *indirectly* from the presence of more merchants, through their effect on the price of the good. Meanwhile, rather than levying separate charges on both sides of the market, the search engine’s sole instrument is the advertising fee for merchants.

In setting this fee, however, the issues the search engine faces include those faced by other multi-sided platforms. The optimal fee, A^* , reflects merchants’ willingness of merchants to pay, satisfying

$$A^* = -\frac{n}{dn/dA}.$$

At the same time, by inducing a price for the good that satisfies equation (3), the search engine takes on the role of attracting users. Since it is costly for users to search, they do so only if the market price for the good is low enough to compensate for their query 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 1. *As search demand becomes more elastic with respect to purchasing demand, the search engine’s optimal number of advertisers increases. Specifically, this number, n^* , is given by*

$$n^* = 1 + \frac{\varepsilon_m}{\varepsilon_D}. \quad (4)$$

Proof. See Appendix A.3. □

This corollary reflects the search engine’s use of competition as a way to commit to offering users a sufficiently low price to justify searching. Note that if the number of users were set first, and the search engine could choose its advertising fee, *afterwards*, 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 expression (2) to calculate ε_m as well as the definition of ε_D gives the following expression for the ratio of these elasticities. Omitting the arguments of the functions, we have,

$$\frac{\varepsilon_m}{\varepsilon_D} = -\frac{f D^2}{F D'}. \quad (5)$$

Equation (5) 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, confirms 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,

$$(E[v_i|v_i \geq p] - p) \Pr[v_i \geq p] = \int_p^{\bar{v}} D(\tilde{p}) d\tilde{p},$$

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 distributed over some non-negative intervals, $[\underline{\theta}, \bar{\theta}]$ and $[\underline{v}, \bar{v}]$, consistent with an interior solution. With these assumptions, the game can readily be solved.

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

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

where $\Delta v \equiv \bar{v} - \underline{v}$. This implies the inverse-demand function, $p(\cdot)$, is 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

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

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

$$p^*(n) = \frac{\bar{v} + nc}{n + 1}.$$

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

$$\int_p^{\bar{v}} D(\tilde{p}) d\tilde{p} = \frac{[\bar{v} - p]^2}{2\Delta v} \geq \theta_i.$$

The mass of users who search is then given by

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

where $\Delta \theta \equiv \bar{\theta} - \underline{\theta}$.

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 $m(p)D(p)(p - c)$, with respect to p . Plugging in the specific functions $m(p)$ and $D(p)$, the search engine’s problem writes as

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

The solution to this is, $p^* = (\bar{v} + 3c)/4$, and search engine profits are $\frac{27/512}{(\Delta v)^2 \Delta \theta} [\bar{v} - c]^4$. Since $p^* = p^*(n)$, we can back out the optimal number of advertisers for the search engine, n^* , which is

$$n^* = 3. \tag{6}$$

Thus we see that with uniformly distributed valuations and query costs, the optimal number of advertisers does not vary with changes in the intervals from which they are drawn. This is of course consistent with the result given in equation (4), since, with uniform distributions,

it is always the case that $\varepsilon_m/\varepsilon_D = 2$.

3 Endogenous Search Quality

Thus far, users' query costs have been exogenous; I now drop this assumption. Clearly, the various characteristics over which the search engine has control, such as the scope and relevance of its results, its speed and the organization and features of its results page, affect the amount of effort one must exert and the level of (in)convenience one experiences when searching. In the framework of this model, it is appropriate, as an approximation, to consider a one-dimensional variable called "search quality", which, as it increases, lowers the query cost for each user.

I now add, to the baseline model, the additional decision for the search engine of setting this search quality variable. Here, I assume the trade-off to be 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 search quality, the search engine must undertake a costly investment. This costly investment could be thought to take the form, for instance, of the building of additional data centers or the employment of additional algorithm designers.

The key lessons of this section are the following. If increasing search quality reduces users' query costs in a relatively uniform fashion, then, when a technological improvement causes the search engine to make a quality improvement, it will, at the same time, have an incentive to increase in the price of the good. If, on the other hand, an increase in search quality causes a greater improvement for users with high query costs, then the opposite may hold – when the search engine increases quality, it will also have an incentive to make the market for the good more competitive. In other words, in the most straightforward case, where search quality affects users proportionally, the two are "Edgeworth complements",⁸ whereas in the subtler case of disproportionate query cost reduction, the two can be substitutes. Proposition 2a gives a general result regarding this issue, and Proposition 2b makes one simplifying assumption in order to give a richer intuition.

3.1 The Model: Investment in Quality

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,

⁸For a detailed discussion of this concept, see Milgrom and Roberts (1995).

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

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. Thus, 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 its advertising fee, A .

3.2 Equilibrium Analysis

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. We can thus write its profit maximization problem as,

$$\max_{s,p} \{m(s,p)D(p)[p - c] - h(s, \alpha)\}. \quad (7)$$

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.⁹ Second, the search engine incurs higher costs, given by h , the higher the search quality it sets.

Looking first at the choice of price, it is apparent that the first-order condition of this problem with respect to p yields a dual inverse-elasticity pricing rule analogous to the one given by Proposition 1,

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

⁹Note 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. Meanwhile, the optimal search quality, s^* , must satisfy the first-order condition,

$$\frac{\partial m(s^*, p^*)}{\partial s} D(p^*) [p^* - c] = \frac{\partial h(s^*, \alpha)}{\partial s}. \quad (8)$$

Expression (8) states that quality’s marginal contribution of to the mass of users, times the profits per unit of mass must be equal to the marginal cost of producing quality.

A particularly interesting issue arises when the search engine must set both quality and price, regarding the co-movement of the two. Are price and quality (Edgeworth) complements, so that, when the search engine increases quality it will also want to increase the price at which the good is sold? Proposition 2a characterizes their relationship.

Proposition 2a. *Price and quality are complements as long as a change in quality does not have too great an effect on the price-elasticity of the mass of users. Precisely, when an exogenous shock directly causes a change in only the search engine’s optimal quality, s^* , or only it’s optimal price, p^* , the search engine responds by adjusting the unaffected variable in the same direction if and only if*

$$\varepsilon_m > \varepsilon \frac{\partial m}{\partial s}. \quad (9)$$

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 B.1 □

To understand this result, consider first, the case in which price and quality are complements. Suppose that α decreases, representing a technological improvement in the quality technology. This induces the search engine to increase its quality, boosting users’ willingness to query. Due to this increased query demand, the search engine will increase the fee it charges to merchants, causing fewer to join. Since fewer merchants join, the price of the good increases.

When price and quality are substitutes, the story is only slightly different. Once again, imagine that α decreases. Then, as before, the search engine increases quality and users’ query demand increases. Due to the increased query demand, the search engine increases

its advertising fee. However, in spite of the increase fee, because of the increase in users, more merchants pay to advertise, causing the price to fall.

As it turns out, either of these two scenarios can potentially occur. Which one *does* depends on whether condition (9) holds. This above discussion describes the complementary issue in terms of the function $m(\cdot, \cdot)$. The next subsection expands on this issue by relating it to the primitives of the model.

3.3 Interpretation in Terms of Query Costs

In this subsection, I present Proposition 2b. This result enhances our intuition for the co-movement of price and quality by stating the issue 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,

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

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

$$\varphi^{-1}\left(\int_p^{\bar{v}} D(\tilde{p}) d\tilde{p}, s\right) = \theta^*. \quad (10)$$

where θ^* is the type of the "threshold user" who is indifferent between searching and not searching, for a given 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 price¹⁰ 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: $\theta_i \sim U[\underline{\theta}, \bar{\theta}]$. We can then derive the result given by Proposition 2b.

Proposition 2b. *When user types, θ , are uniformly distributed, price and quality are complements as long as the effect of a change in quality is not too concentrated on users with high query costs.*

¹⁰To 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.

Specifically, s^* and p^* vary in the same direction if and only if

$$\varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial \varphi}{\partial \theta^*}} < 0, \quad (11)$$

where $\varepsilon_{\varphi^{-1}} \equiv -\frac{\partial \varphi^{-1}}{\partial s} \frac{s^*}{\varphi^{-1}(\int_{p^*}^{\bar{p}} D(\tilde{p}) d\tilde{p}, s^*)}$ and where $\varepsilon_{\frac{\partial \varphi}{\partial \theta}} \equiv -\frac{\partial}{\partial s} \left\{ \frac{\partial \varphi}{\theta} \right\} \frac{s^*}{\partial \varphi(\theta^*, s^*) / \partial \theta}$.

Proof. See Appendix B.2 □

To interpret (11), note that the first term, $\varepsilon_{\varphi^{-1}}$ is negative, and that the key issue is whether the second term, $\varepsilon_{\frac{\partial \varphi}{\partial \theta}}$, is sufficiently positive. This depends crucially on the cross derivative of the query cost function,

$$\frac{\partial}{\partial s} \left\{ \frac{\partial \varphi(\theta^*, s^*)}{\theta} \right\} = \frac{\partial}{\partial \theta} \left\{ \frac{\partial \varphi(\theta^*, s^*)}{s} \right\}. \quad (12)$$

When (12) is negative, this means that as the type of the “threshold searcher” increases, the marginal effect on her query cost of a quality increase becomes more important. For price and quality *not* to be complements, this term must be sufficiently negative.

In section 2, I suggest that a natural interpretation of users’ types is as a parameter of their level of expertise at using a search engine. Following this interpretation, we can think of high types as “novice” searchers. Thus, Proposition 2b shows that unless quality changes disproportionately affect novice users, quality and price are complements.

A Specific Query Cost Function. Here I give an example to elucidate this interpretation. 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 for users of all types. Meanwhile, for high values of γ , an increase in s causes a decrease in the cost of querying that is more sizeable for high-type users than for low-type users.

This particular form of the query cost function, φ , implies,

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

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, 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.

3.4 Example: Price and Quality Complements

This subsection examines 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. Suppose that both sets of private values are uniformly distributed over the unit interval: $\theta_i \sim U[0,1]$ and $v \sim U[0,1]$. 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, I ignore the technology parameter, α .) Finally, let the users' query cost be $\varphi(\theta_i, s) = \theta_i - s$.

Note that this corresponds to the extreme case, where $\gamma = 0$, of the query costs function discussed in subsection 3.3. 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,

$$\max_{s,p} \left\{ \underbrace{\left(s + \frac{[1-p]^2}{2} \right)}_{m(s,p)} (1-p)p - \frac{s^2}{2} \right\}.$$

The solutions to this problem are $p^* = 1/3$ and $s^* = 2/9$. For comparison, note that, here, when $s = 0$, we are in a special case of the example given at the end of the subsection 2.3, where query costs were exogenous and maximization was done only with respect to price. Recall that in that example, the search engine's optimal price was given by $(v + 3c)/4$ 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 the quality instrument.

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

$$n^* = 2.$$

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

searchers 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 is on the interval $[0, \frac{2}{9}]$) who, in equilibrium, have negative query costs. In other words, there is a set of users for whom search quality is high enough they would query even if 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 assume that the search engine must spend *its own* resources in order to produce “search quality”. In this section, I modify this assumption. Here, I take into account the fact that a search engine’s “content” largely consists of “organic results” – links to *other*, freely available websites. While high quality organic results draw users to the search engine, they also compete with sponsored links

This type of competition could come directly, from the presence of non-advertising merchants’ links among the organic, “left side” search results. If the unpaid results do link to merchants’ websites, then a user has less incentive to pay attention to or to click on the paid advertisements on the “right side” of the page. This sort of competition may seem unlikely, on the basis that the search engine has little incentive to display links to non-paying merchants. Instead of displaying links to additional merchants, one might expect the search engine to display links to other types of sites that are more “informational”.

Doing so, however, does not eliminate the issue of competition from the left side, which can also arise indirectly, *via* informational websites. In particular, informational sites that are of high quality, and thus of value to users, need to make money in order to support themselves.¹¹ Thus they have a strong incentive to market the good themselves, either by selling it directly, or by placing advertisement links on their own sites to merchants who sell the good.¹²

¹¹Indeed, sometimes seemingly independent informational websites are maintained *by merchants*, in such a way so as to appear prominently in organic search results and to direct traffic to the merchant’s site. One such example is the Internet Movie Database (imdb.com). This site, which maintains painstakingly detailed pages of data and reviews for nearly every movie ever released, often appears near the top of the organic part of search engine results pages when one enters a given movie title as a query. It is owned by Amazon, and, in addition to providing information about each movie, it also features a link to buy the movie on DVD (at amazon.com). Since it is not my goal in this paper to explain organic sites’ existence, in this section, I do not model this issue of investment in informational sites but rather take it as given that they exist.

¹²Websites of all descriptions can be found to have some space devoted to advertising, of which two forms

To capture this issue, I extend the model to include left side merchants. Instead of supposing that the search engine must bear explicit costs in order to provide a certain level of search quality, I suppose that adding quality imposes an implicit cost – the presence of non-paying merchants that compete in the Cournot game with the advertisers. Under this assumption, the results change in two ways. First, as shown in Proposition 3, left side merchants serve as an additional force driving down the price of the good that the search engine induces. Second, unlike the case without left side merchants, quality and price appear typically to be *substitutes* for the search engine.

4.1 The Model: Left Side Merchants

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 $l(s) \geq 0$ non-paying "left side" merchants who compete with r "right side" merchants that pay, as in previous sections, to advertise. The total number of firms, n , that compete to sell the good to users is now given by $n = r + l$. The timing of the game is unchanged from section 3.1. Note that when the search engine sets a given level of quality, \hat{s} , in the first stage, it is common knowledge that $l(\hat{s})$ will be present in the competition stage, in addition to those who pay to appear.

As in the previous sections, I focus 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, search quality, s , and the advertising fee, A . As such it can be written

$$\max_{A,s} \left\{ r(A,s) \frac{m(s,p)D(p)(p-c)}{l(s) + r(A,s)} \right\}, \quad (13)$$

where $p = p(l(s) + r(A,s))$ is a function of the number of merchants on both sides. From (13), we see 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, I simplify the problem by writing the search engine's maximization problem

are most prominent: "display" or "banner" and "contextual". The former typically uses catchy imagery to draw user attention, while the latter takes a more subtle, text-based form that is similar in appearance to search advertisement. An interesting exception is Wikipedia, the online, user-generated encyclopedia, which does not show ads. This section's model suggests a reason why Wikipedia entries frequently appear at the top of search engines' organic results.

directly with respect price. In order to do this, let $\lambda \equiv l/(l+r)$ denote the proportion of total “industry” profits that the search engine receives. Since the function $l(s)$ is a one-to-one correspondence, we can define a new function for the mass of users who search, $\tilde{m} = \tilde{m}(l, p) = \tilde{m}(\lambda n(p), p)$. An equivalent way to write the search engine’s profit maximization problem is then

$$\max_{\lambda, p} \{(1 - \lambda)\tilde{m}(\lambda n(p), p)D(p)(p - c)\}, \quad (14)$$

where the search engine jointly selects its desired final-stage price, p^* , and its desired proportion of the rent from the final stage, given by $(1 - \lambda^*)$.

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

$$\frac{p^* - c}{p^*} = \frac{1}{\varepsilon_{\tilde{m}}^{tot} + \varepsilon_D}. \quad (15)$$

Note that here, the first term in the denominator, $\varepsilon_{\tilde{m}}^{tot}$, is the total price-elasticity of the mass of users, defined as

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

Meanwhile, the first-order condition with respect to λ is

$$\tilde{m}(\lambda^* n(p), p) = (1 - \lambda^*) \frac{\partial \tilde{m}}{\partial \{\lambda n(p)\}} n(p). \quad (16)$$

Combining the two first-order conditions in (15) and (16) allows us to derive Proposition 3, which illustrates the effect of left side merchants on the price the search engines induces.

Proposition 3. *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*

$$\frac{p^* - c}{p^*} = \frac{1}{\frac{\lambda^*}{1-\lambda^*} \varepsilon_n + \varepsilon_{\tilde{m}} + \varepsilon_D}, \quad (17)$$

where

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

Proof. See Appendix C.1. □

This pricing formula highlights the “three-sidedness” of the search engine’s problem. Even with the Web freely available to display, the search engine will optimally exercise restraint in showing it. In order to maximize profits, the search engine *sometimes* connects users with a third type of non-paying agent in order to have a better opportunity to *other*

times connect them with paying advertisers. In deciding the relative frequency with which it connects users to each type of merchant, the search engine must take into account the fact that the presence of merchants on the left cuts into the price that paying advertisers charge users.

To see this formally, note that in (17), in addition to $\varepsilon_{\bar{m}}$ and ε_D , whose roles here are unchanged from earlier discussion (see Proposition 1), when left side merchants are present, the term ε_n – the elasticity of the total number of merchants with respect to price – also affects the search engine’s optimal price for the good. This is due to the relationship between search quality and, on the one hand, users’ query costs, and on the other hand, the number of firms on the left side. This third elasticity is weighted by $\frac{\lambda^*}{1-\lambda^*}$, reflecting the fact that the solution to the search engine’s tradeoff between inducing a high markup and reducing user query costs depends on the proportion of merchants that are on the left side. Subsection 4.2 looks in more detail at the relationship between price and the number of left side merchants.

4.2 Movement of Price and Quality

In the model of section 3, price and quality are complements, so long as changes in quality do not have too disproportionate an effect on users with high query costs. As Proposition 4 shows, the addition of left side merchants to the model pushes price and quality in the direction of being substitutes.

Proposition 4. *When there are non-paying left side merchants, price and quality are complements only when the increase in the number of non-paying merchants caused by an increase in quality is optimally met by a compensating decrease of greater size in the number of advertising merchants. This is the case if and only if*

$$\varepsilon_{\bar{m}} > \varepsilon_{\frac{\partial \bar{m}}{\partial l}} + \psi, \quad (18)$$

where $\varepsilon_{\frac{\partial \bar{m}}{\partial l}} \equiv \frac{\partial}{\partial p} \left\{ \frac{\partial \bar{m}(l^*, p^*)}{\partial l} \right\} \frac{p^*}{\partial \bar{m}(l^*, p^*) / \partial l}$ and $\psi \equiv -n'(p^*) \frac{\frac{\partial \bar{m}}{\partial l} p^*}{\bar{m}(l^*, p^*)} > 0$.

Proof. See Appendix C.2. □

The condition in (18) contains analogues of the two elements that play a role in Proposition 2a, as well as an additional term, ψ . This term, in effect, sets the bar higher for price and quality to be complements. To see why, imagine that there is an exogenous change causing an increase in the optimal price of the good but not directly affecting the optimal division of merchants between the left and right side. Due to the increase in price, the total number of merchants decreases. If, with this decreased number of total merchants, the search engine were to keep the same proportion on the left, quality would suffer.

Thus, in order for a quality increase to accompany such a price increase, the search engine must meet the decrease in total number of merchants with a corresponding increase in the proportion that are on the left side. Moreover, this increase in the left side proportion must be large enough to increase the total number non-paying merchants. Expression (18) shows that for this to be optimal for the search engine, the direct effect of a change in price on the mass of users must be enough greater than its indirect effect (via changes in the marginal contribution of quality) *to offset* the change in total number of merchants brought on by a shift in price.

4.3 Example: Left and Right Side Merchants

Here I give a simple example that isolates the new force present under the assumptions of section 4 that pushes price and quality to be substitutes. Assume that user i 's valuations for the good, v_i , is uniformly distributed over the interval $[0, 1]$, and her type, θ_i , is uniformly distributed over the interval $[0, \bar{\theta}]$, where $\bar{\theta} > 0$ is sufficiently large to give rise to an interior solution.¹³ Also, let $c = 0$, and let the users' query cost function, φ , be given by $\varphi(\theta_i, s) = \theta_i/s^\gamma$, 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,

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

This implies that

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

Thus, if the cost of producing quality were solely explicit – as in section 3 – and there were no left side merchants, then, this set of assumptions would imply that the optimal price for the search engine would not be affected by an independent change in the optimal quality. Here, however, I assume that the cost of providing quality is implicit. Specifically, let $l(s) = s$.

To write the search engine's maximization problem directly in terms of λ and p , I use the fact that the outcome of the n -merchant Cournot game entails

$$p = \frac{1}{n+1} \Leftrightarrow n = \frac{1-p}{p}.$$

¹³The condition for this is $\bar{\theta} > \gamma^\gamma(3+\gamma)^{2+\gamma}/32(1-\gamma^2)^\gamma$. If this condition does not hold, the search engine will induce all users to search.

This gives that the mass of users who search is

$$m(s, p) = \tilde{m}(\lambda n, p) = \left(\lambda \cdot \frac{1-p}{p} \right)^\gamma \frac{(1-p)^2}{2\bar{\theta}}. \quad (19)$$

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

$$\max_{\lambda, p} \left\{ \frac{1}{2\bar{\theta}} (1-\lambda) \lambda^\gamma (1-p)^{3+\gamma} p^{1-\gamma} \right\}.^{14} \quad (20)$$

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

$$p^*(\gamma) = \frac{1-\gamma}{4} \quad \text{and} \quad n^*(\gamma) = \frac{3+\gamma}{1-\gamma}. \quad (21)$$

Note that the intensity of competition among merchants that the search engine induces is increasing in γ . Turning our attention to the choice of search quality, the optimal proportion of merchants placed on the left side, λ^* , is

$$\lambda^*(\gamma) = \frac{\gamma}{1+\gamma}. \quad (22)$$

Thus, the optimal search quality, s^* , is

$$s^*(\gamma) = \frac{\gamma(3+\gamma)}{1-\gamma^2}. \quad (23)$$

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

$$r^*(\gamma) = \frac{3+\gamma}{1-\gamma^2}. \quad (24)$$

The expression for the search engine's equilibrium profits, as a function of γ is

$$\frac{\gamma^\gamma (3+\gamma)^{3+\gamma} (1-\gamma)^{1-\gamma}}{2^9 \bar{\theta} (1+\gamma)^{1+\gamma}}. \quad (25)$$

¹⁴Note that when this problem is written with p and s as the search engine's decision variables, it becomes

$$\max_{p, s} \left\{ \frac{1}{2\bar{\theta}} s^\gamma \left(\frac{1-p}{p} - s \right) (1-p)^2 p^2 \right\}.$$

Thus, an exogenous change in γ affects the first-order condition with respect to p only indirectly through its effect on the optimal value of s .

Figure 1 plots several of these results.

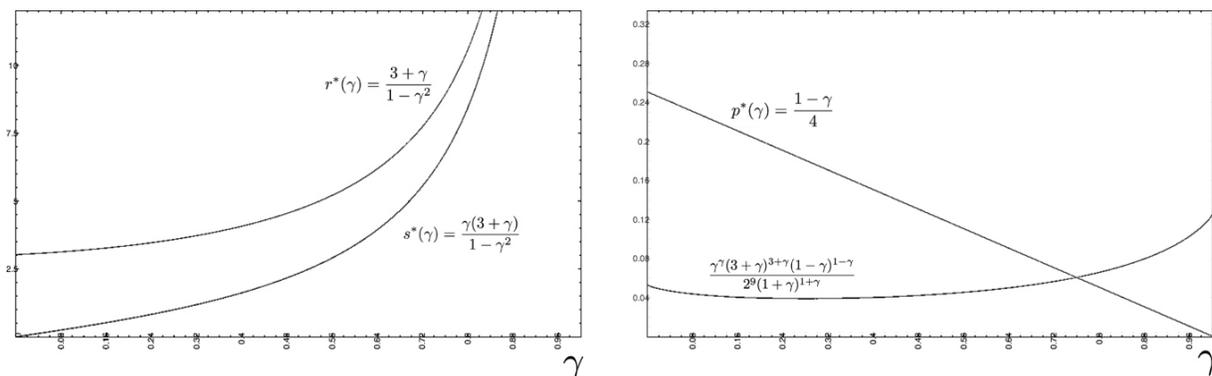


Figure 1: Left: number of advertisers, r^* , and quality, s^* ; Right: price, p^* , and search engine profits (scaled by $\bar{\theta}$).

A striking feature of this example is the responsiveness, in opposite directions, of price and quality to the parameter γ . For very small values of γ , quality is not very important for users,¹⁵ causing the search engine to set it at a low level. As γ approaches 0 from above, so does the number of left side merchants, and the outcome of the game resembles that of the example of subsection 2.3, in which the search engine's optimal number of merchants is 3. As γ increases, the search engine places relatively more weight on drawing in users by providing high quality. In order to implement this quality increase, the search engine increases both the total number of merchants, n^* , and the proportion of merchants on the left side, λ^* . Thus, we see that in this example, as left side quality becomes sufficiently important for users, the incentive for the search engine to induce a high markup can disappear completely.

5 Google, DoubleClick and Conventional Wisdom

Using the framework developed above, in this section I draw a thumbnail sketch of Google's recent acquisition of the ad-serving firm DoubleClick. An issue in the regulatory debate surrounding that acquisition was whether advertisers consider as substitute goods search ads and non-search ads on other websites. Parties against the merger argued that these two forms of advertisement *are* substitutes, while parties in favor of the merger argued

¹⁵An alternative interpretation of γ involves thinking of it as an inverse measure of the rate at which merchants appear on the left side as quality increases and holding fixed users' sensitivity to quality. This view is supported by the fact that when one lets $l(s) = s^{1/\gamma}$ and $\varphi = \theta_i/s$, one obtains the same results.

that they *are not*. For instance, the “Statement of Federal Trade Commission Concerning Google/DoubleClick”, in which the Commission explains its decision not to impose any conditions on the merger, contains the assertion that “the evidence in this case shows that the advertising space sold by search engines is not a substitute for space sold directly or indirectly by publishers or vice versa” (2007, p. 3). The belief was clearly that the danger of anti-competitive behavior would be greater if the two forms of advertising were substitutes.

The model presented in this paper indicates, however, that the ill effects on consumers of anticompetitive behavior may be smaller in the case that these two forms of advertisement are in some ways substitutes. I show this using the analysis of sections 3 and 4 to compare the likely effects of the merger when the two forms of advertising are not substitutes and when they are.

First, if search ads and non-search web ads (hereafter just “web” ads) are never substitutes in the eyes of advertisers, then users will enjoy higher search quality but will be harmed by decreased competition among search advertisers. In this case, advertisers on the search engine cannot potentially attract consumers for the same good using web advertisement. As a result, there’s no direct competition between organic results and sponsored links, and in this case, the assumptions of section 3 provide a reasonable approximation the search engine’s incentives.

By merging with a web ad-serving firm, the search engine would benefit from various technological improvements that reduce the cost of providing users with search quality. In particular, it would gain access to more data on users’ behavior when interacting with non-search webpages, which is useful for determining pages’ relevance. Thus, let us assume that the merger would cause a decrease in the search engine’s technology parameter, α . As a result of this change, the search engine would increase the level of search quality. In the most probable case – i.e. where price and quality are complements for the search (see Proposition 2a) – the search engine would also charge advertisers higher fees, by sufficiently much to induce an increase in the price consumers pay for the advertisers’ good.

By contrast, in the case where search ads and web ads *are* substitutes from the perspective of advertisers, as a result of the merger, users would likely see higher search quality *and* lower prices for advertisers’ goods. In this case, there is direct competition in the market for a given good among search and web advertisers. Under this scenario, the assumptions of section 4 give a more accurate depiction of the search engine’s incentives.

To examine the effects of the merger, in this case, let us write the search engine’s profits as a generalization of expression (14),

$$(1 - \rho\lambda)\tilde{m}(\lambda n(p), p)D(p)(p - c).$$

In the case prior to the merger, $\rho = 1$, since the search engine makes profits only from the merchants paying for search advertisements. Post-merger, however, the combined entity also engages in web ad-serving, and thus recoups some profits from transactions between users and left side merchants.

Accordingly, after the merger, $0 < \rho < 1$. Given this lowered value of ρ , the search engine would have incentive to provide improved search quality. Meanwhile, in the more likely case – i.e. where price and quality are substitutes for the search engine (see Proposition 4) – the search engine would also have incentive to induce a reduction in the price that advertisers charge consumers.

6 Related Literature

This work adds to a quite new body of research on Internet search engines and web advertisement. Much of this literature examines the auction mechanisms that search engines use to sell their paid links. Several papers that focus on this issue are Börgers et al. (2008), Edelman and Ostrovsky (2007), Edelman et al. (2007), Katona and Sarvary (forthcoming), Milgrom (2008) and Varian (2007). Another strand of the literature focuses on consumer search problems that arise when Internet users are uninformed of what lies behind individual links. Armstrong et al. (2009), Ellison and Ellison (2009) and Taylor (2009) concentrate on this issue. A set of papers that incorporate both of these issues into a single theoretical framework are Athey and Ellison (2009), Chen and He (2006) and Rayo and Segal (2009), while Ghose and Yang (2009) study empirically the way such interaction depends on the search term entered by the user. Pollock (2009) and Telang et al. (2004) study the organization of the search engine industry.

This paper also contributes to the literature on two-sided markets.¹⁶ In particular, it considers a platform with specific control over the interaction between the two sides of the market (Baye and Morgan (2001), Hagiu (2007), Hagiu and Jullien (2009) and Nocke et al. (2007)) and investigates the platform's incentives to provide quality to consumers. To do so, it proposes two variants of the model, one of which (section 3) is an extension of a classical model of monopoly quality provision (Sheshinski (1976) and Spence (1975)) to a two-sided environment. The latter of these two variants (section 4) uncovers an additional force governing the relationship between a search engine's quality and the profitability of its advertisement.

More generally, search engines are a form of ad-supported media, studies of which

¹⁶For more general coverage of this topic, see Armstrong (2006), Caillaud and Jullien (2003), Evans (2003), Rochet and Tirole (2003), Rochet and Tirole (2006), Weyl (forthcoming).

include Anderson and Coate (2005), Ellman and Germano (2009) and Reuter and Zitzewitz (2006). This additional force, whereby “content” competes with advertisement not only for user attention, but also directly in the advertisers’ market, has, to the best of my knowledge, not been discussed in this literature. I believe this to be indicative of the fact that this phenomenon appears much more likely to affect a search engine than it does traditional forms of media such as a newspaper or a television channel. An interesting issue is whether it will also be present in the context of other, constantly developing forms of ad-supported web (and mobile web) media.

7 Conclusion

Increasingly, a single search engine determines what information society sees, and thus it seems important that the incentives the search engine faces be clearly understood. They are, however, extremely complex: the web and the way people use it are constantly changing, and, even at any given moment, the world of online advertising is exceedingly complicated. This paper represents a small step towards a sharper understanding of a search engine’s motives and the likely effects of its behaving in accordance with them.

In particular, this paper explores the costs and benefits to a search engine of providing Internet users with high quality unpaid (“organic” or “left side”) search results in addition to showing them paid advertisements. Simply put, a search engine receives queries from users and, in response, shows lists of results or “links”. Some of the links on these lists are selected to appear by the search engine’s algorithm, whereas others appear thanks to monetary agreements between the search engine and the respective linked websites.

A search engine’s profitability thus stems from its ability to sell market power to advertisers. The reason that a search engine finds it profitable to show better rather than worse results on the left side is that doing so makes the market power it sells worth more. Since the search engine is effectively a three-sided platform, the relationship between its choice of quality and its sale of market power turns out to be particularly interesting, and the paper identifies two key factors that govern it.

The first such factor is that advertisers who pay the search engine *to obtain* market power also *compete it away* from one another. As a result, while increasing quality allows the search engine to charge higher advertising fees, it is ambiguous whether this causes advertisers to charge higher or lower prices in their subsequent interactions with users. The paper shows that when improvements in search quality benefit all users equally, advertisers will charge a higher price. However, when improvements in search quality provide a greater benefit to novice searchers, advertisers will charge a lower price.

The second factor is that advertisers who pay for market power also compete it away with sites appearing *among the organic results*. Moreover, in many situations, when the search engine provides users with higher quality left side results, more of the links on this list point to websites that sell the same product, or fulfill the same need as the advertisers. Thus, when this is the case, an increase in quality puts downward pressure on the price that advertisers can charge consumers.

While further research is required to determine which of these effects are most dominant, the crucial message is that the a search engine's profits depend highly on the type of organic results it shows to users. This issue should be studied in more detail and, in particular, richer models should be developed to understand the relationship between the effects discussed here and the auctions that search engines use in practice to allocate their ad slots. Currently, despite the regulatory eyebrows Google has been raising, there is remarkable silence over the incentives to "manipulate" the left side. When not a single mention of this issue appears in the "Statement of Federal Trade Commission Concerning Google/DoubleClick" (F.T.C. (2007)), nor in the accompanying "Dissenting Statement" (Jones Harbour (2007)), nor in any of the articles published in *Global Competition Policy* in response to the FTC's decision¹⁷ – all documents that mull over lists of potentially bad things Google might want to do – one wonders whether there's an elephant in the room.

¹⁷See *Global Competition Policy Magazine*, March, 2008, Release One.

Appendices

A Proofs for Section 2: Optimal Pricing

A.1 Lemma 1

Proof. The profits of an individual advertiser, π_j , can be written as a function of the advertising fee, A , and the number of merchants, n , as

$$\pi_j(A, n) = \frac{m(p^*(n))D(p^*(n))(p^*(n) - c)}{n} - A. \quad (26)$$

Suppose the search engine sets a fee $\tilde{A} > 0$ in the first stage. In the equilibrium of the continuation game, the number of advertiser, \tilde{n} , must satisfy

$$\begin{aligned} \frac{m(p^*(\tilde{n}))D(p^*(\tilde{n}))(p^*(\tilde{n}) - c)}{\tilde{n}} &= \tilde{A} \\ \Leftrightarrow m(p^*(\tilde{n}))D(p^*(\tilde{n}))(p^*(\tilde{n}) - c) &= \tilde{n}\tilde{A} = \Pi_{SE}, \end{aligned} \quad (27)$$

where Π_{SE} denotes the search engines profits.

Since we have assumed that the search engine's problem has a unique, interior solution, it must be the case that \tilde{n} is a decreasing function of \tilde{A} in the neighborhood of the optimal fee, A^* . Therefore, by setting the fee equal to A^* , the search engine determines, uniquely, the equilibrium number of advertisers and the equilibrium price, p^* , so as to maximize the left-hand side of expression (27).

For the second part, we have that user i chooses to search if and only if

$$[E[v_i|v_i > p] - p] \times \Pr[v_i > p] = \int_p^{\bar{v}} D(\tilde{p}) d\tilde{p} \geq \theta_i.$$

□

A.2 Proposition 1

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

$$\frac{m'(p^*)}{m(p^*)} + \frac{D'(p^*)}{D(p^*)} + \frac{1}{p^* - c} = 0. \quad (28)$$

Multiplying this equation by p^* and rearranging gives the expression in (3). □

A.3 Corollary to Proposition 1

Proof. We combine the result of Proposition 1 with the standard result in symmetric n -player Cournot that

$$\frac{p^* - c}{p^*} = \frac{1}{n} \cdot \frac{1}{\varepsilon_D}. \quad (29)$$

Simplifying, we obtain the expression for n^* given by (4). \square

B Proofs for Section 3: Endogenous Quality

B.1 Proposition 2a

Proof. Search engine profits are given by $\Pi(s, p) = m(s, p)D(p)[p - c] - h(s, \alpha)$. Price and quality are complements if

$$\frac{\partial^2 \Pi(s^*, p^*)}{\partial s, \partial p} = \left(\frac{\partial^2 m}{\partial p \partial s} D + \frac{\partial m}{\partial s} D' \right) (p^* - c) + \frac{\partial m}{\partial s} D > 0 \quad (30)$$

and substitutes otherwise, where s^* and p^* denote the optimal values of these variables for the search engine. Algebra yields that the condition in (30) holds, when $\varepsilon_{\frac{\partial m}{\partial s}} + \varepsilon_D > (<)0$, if and only if

$$\frac{p^* - c}{p^*} < (>) \frac{1}{\varepsilon_{\frac{\partial m}{\partial s}} + \varepsilon_D}.$$

Since

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

we have the desired result. \square

B.2 Proposition 2b

Proof. In view of the result shown in proposition 2a, it suffices to show, that when $\theta_i \sim \text{U}[\underline{\theta}, \bar{\theta}]$, we have $\varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial \varphi}{\partial \theta}} < (>)0$ if and only if $\varepsilon_m > (<)\varepsilon_{\frac{\partial m}{\partial s}}$. Generally, we have

$$m(s, p) = F \left(\varphi^{-1} \left(\int_p^{\bar{p}} D(\tilde{p}) d\tilde{p}, s \right) \right). \quad (31)$$

Hence, when θ is uniformly distributed, we have

$m(s, p) = \frac{1}{\Delta\theta} \cdot \varphi^{-1} \left(\int_p^{\bar{p}} D(\tilde{p}) d\tilde{p}, s \right)$	$\frac{\partial m}{\partial p} = \frac{1}{\Delta\theta} \cdot \frac{-D(p)}{\partial \varphi / \partial \theta}$
$\frac{\partial m}{\partial s} = \frac{1}{\Delta\theta} \cdot \frac{\partial \varphi^{-1}}{\partial s}$	$\frac{\partial^2 m}{\partial p \partial s} = \frac{1}{\Delta\theta} \cdot D(p) \cdot \frac{\partial^2 \varphi}{\partial \theta \partial s} / \left(\frac{\partial \varphi}{\partial \theta} \right)^2$

Plugging these values into ε_m and $\varepsilon_{\frac{\partial m}{\partial s}}$ and rearranging, one obtains the stated result. \square

C Proofs for Section 4: Left Side Merchants

C.1 Proposition 3

Proof. To obtain result of this proposition, we must thus show that $\varepsilon_{\tilde{m}}^{tot} = \frac{\lambda^*}{1-\lambda^*}\varepsilon_n + \varepsilon_{\tilde{m}}$. Note that

$$\varepsilon_{\tilde{m}}^{tot} = - \left[\frac{\partial \tilde{m}(\lambda^* n(p^*), s^*)}{\partial \{\lambda n(p)\}} \lambda^* n'(p^*) + \frac{\partial \tilde{m}(\lambda^* n(p^*), s^*)}{\partial p} \right] \cdot \frac{p^*}{\tilde{m}(\lambda^* n(p^*), s^*)}, \quad (32)$$

and that, from the first-order condition with respect to λ , (16), we have

$$\frac{\partial \tilde{m}(\lambda^* n(p^*), s^*)}{\partial \{\lambda n(p)\}} = \frac{\tilde{m}(\lambda^* n(p^*), p^*)}{(1-\lambda^*)n(p^*)}. \quad (33)$$

Inserting the right-hand side of (33) into (32) and simplifying gives the result. \square

C.2 Proposition 4

Proof. The search engine's first-order condition, with respect to λ , (16), implies that

$$\Phi(l, p) \equiv (n(p) - l) \frac{\partial \tilde{m}(l, p)}{\partial l} - \tilde{m}(l, p) = 0 \quad (34)$$

Let $l^*(p)$ denote the search engine's optimal number of left side merchants as a function of p . By the implicit function theorem and the second-order condition when taken with respect to l , price and quality are complements if

$$\frac{\partial \Phi(l^*(p), p)}{\partial p} = (n(p) - l^*(p)) \frac{\partial^2 \tilde{m}}{\partial l \partial p} + n'(p) \frac{\partial \tilde{m}}{\partial l} - \frac{\partial \tilde{m}}{\partial p} > 0 \quad (35)$$

and substitutes otherwise. By (34) we have

$$n(p) - l^*(p) = \frac{\tilde{m}(l^*(p), p)}{\partial \tilde{m} / \partial l}, \quad (36)$$

Using (36), the condition in (35) can be rewritten

$$\frac{\partial \Phi(l^*(p), p)}{\partial p} = \frac{\tilde{m}(l^*(p), p)}{\partial \tilde{m} / \partial l} \frac{\partial^2 \tilde{m}}{\partial l \partial p} + n'(p) \frac{\partial \tilde{m}}{\partial l} - \frac{\partial \tilde{m}}{\partial p} > 0,$$

which holds if and only if $\varepsilon_{\tilde{m}} > \varepsilon_{\frac{\partial \tilde{m}}{\partial l}} + \psi$. \square

References

- Anderson, Simon P. and Stephen Coate (2005), "Market Provision of Broadcasting: A Welfare Analysis." *Review of Economic Studies*, 72, 947–972.
- Armstrong, Mark (2006), "Competition in Two-Sided Markets." *The RAND Journal of Economics*, 37, 668–691.
- Armstrong, Mark, John Vickers, and Jidong Zhou (2009), "Prominence and Consumer Search." *The RAND Journal of Economics*, 40, 209–233.
- Athey, Susan and Glenn Ellison (2009), "Position Auctions with Consumer Search." *Mimeo.*, URL <http://www.nber.org/papers/w15253.pdf>.
- Battelle, John (2005), *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Portfolio New York.
- Baye, Michael R. and John Morgan (2001), "Information gatekeepers on the internet and the competitiveness of homogeneous product markets." *The American Economic Review*, 91, 454–474.
- Börger, Tilman, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek (2008), "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence." *Mimeo.*, URL <http://econ.lse.ac.uk/staff/mpesend/papers/EquiBids.pdf>.
- Brin, Sergey and Lawrence Page (1998), "The Anatomy of a Large-Scale Hypertextual Web Search Engine." In *Proceedings of the Seventh International World-Wide Web Conference*.
- Caillaud, Bernard and Bruno Jullien (2003), "Chicken & Egg: Competition Among Intermediation Service Providers." *The RAND Journal of Economics*, 34, 309–328.
- Chen, Yongmin and Chuan He (2006), "Paid Placement: Advertising and Search on the Internet." *Mimeo.*, URL <http://ssrn.com/abstract=936472>.
- Edelman, B., M. Ostrovsky, and M. Schwarz (2007), "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords." *American Economic Review*, 97, 242–259.
- Edelman, Benjamin and Michael Ostrovsky (2007), "Strategic Bidder Behavior in Sponsored Search Auctions." *Decision Support Systems*, 43, 192–198.
- Ellison, Glen and Sara Fisher Ellison (2009), "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, 77, 427–452.
- Ellman, Matthew and Fabrizio Germano (2009), "What do the papers sell? a model of advertising and media bias." *Economic Journal*, 119, 680–704, URL <http://dx.doi.org/10.1111/j.1468-0297.2009.02218.x>.
- Evans, David S. (2003), "The Antitrust Economics of Multi-Sided Platform Markets." *Yale Journal of Regulation*, 20, 325–431.

- F.T.C. (2007) (2007), "Google/DoubleClick, F.T.C. File No. 071-0170: Statement of Federal Trade Commission Concerning Google/DoubleClick." URL <http://www.ftc.gov/os/caselist/0710170/071220statement.pdf>.
- Ghose, Anindya and Sha Yang (2009), "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets." *Management Science*, 55, 1605–1622.
- Hagiu, Andrei (2007), "Merchant or Two-Sided Platform?" *Review of Network Economics*, 37, 115–133.
- Hagiu, Andrei and Bruno Jullien (2009), "Why Do Intermediaries Divert Search?" *Mimeo.*, URL <http://www.people.hbs.edu/ahagiu/Hagiu%20Jullien%20revision%2002122009%20HBS%20WP.pdf>.
- Jones Harbour, P. (2007), "Google/DoubleClick, F.T.C. File No. 071-0170: Dissenting Statement of Commissioner Pamela Jones Harbour." URL <http://www.ftc.gov/os/caselist/0710170/071220harbour.pdf>.
- Katona, Zsolt and Miklos Sarvary (forthcoming), "The Race for Sponsored Links: A Model of Competition for Paid Placement on a Search Engine." *Marketing Science*.
- Krause, Andreas and Eric Horvitz (2008), "A Utility-Theoretic Approach to Privacy and Personalization." In *Proceedings of the Twenty-Third Conference on Artificial Intelligence (AAAI-08)*.
- Mankiw, N. Gregory and Michael D. Whinston (1986), "Free Entry and Social Inefficiency." *The RAND Journal of Economics*, 17, 48–58.
- Milgrom, Paul (2008), "Simplified mechanisms with an application to sponsored-search auctions." *Games and Economic Behavior*, forthcoming.
- Milgrom, Paul and John Roberts (1995), "Complementarities and Fit Strategy, Structure, and Organizational Change in Manufacturing." *Journal of Accounting and Economics*, 19, 179–208.
- Nocke, Volker, Martin Peitz, and Konrad Stahl (2007), "Platform Ownership." *Journal of the European Economic Association*, 5, 1130–1160.
- Pollock, Rufus (2009), "Is Google the Next Microsoft? Competition, Welfare and Regulation in Internet Search." *Mimeo.*, URL <http://ssrn.com/abstract=1265521>.
- Rayo, Luis and Ilya Segal (2009), "Optimal Information Disclosure." *Mimeo.*, URL <http://www.stanford.edu/~isegal/disclosure.pdf>.
- Reuter, Jonathan and Eric Zitzewitz (2006), "Do Ads Influence Editors? Advertising and Bias in the Financial Media." *Quarterly Journal of Economics*, 121, 197–227.
- Rochet, Jean-Charles and Jean Tirole (2003), "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association*, 1, 990–1029.

- Rochet, Jean-Charles and Jean Tirole (2006), "Two-Sided Markets: A Progress Report." *RAND Journal of Economics*, 37, 645–667.
- Sheshinski, Eytan (1976), "Price, Quality and Quantity Regulation in Monopoly Situations." *Economica*, 43, 127–137.
- Spence, A. Michael (1975), "Monopoly, Quality, and Regulation." *The Bell Journal of Economics*, 6, 417–429.
- Taylor, Greg (2009), "Competing Search Engines With Utility Maximising Consumers and Ordinary Search Results." *Mimeo.*, URL <http://ssrn.com/abstract=1233282>.
- Telang, Rahul, Uday Rajan, and Tridas Mukhopadhyay (2004), "The Market Structure for Internet Search Engines." *Journal of Management Information Systems*, 21, 137–160.
- Varian, Hal .R. (2007), "Position Auctions." *International Journal of Industrial Organization*, 25, 1163–1178.
- Weyl, E. Glen (forthcoming), "A Price Theory of Multi-Sided Platforms." *American Economic Review*, URL http://www.people.fas.harvard.edu/%7Eweyl/ptNsms_FINAL.pdf.