Search Engines: Left Side Quality versus Right Side Profits*

Alexander White†

April 2013

Abstract

Search engines face an interesting tradeoff in choosing the way to display their results. While providing high quality unpaid, or "left side" results attracts users, doing so can also cannibalize the revenue that comes from paid ads on the "right side". This paper examines this tradeoff, focusing, in particular, on the role of users' *post-search* interaction with the websites whose links are displayed. In the model, high quality left side results boost demand from users, causing them to tolerate a search engine on which advertisers do not offer the lowest possible prices for the goods that they sell. However, because websites appearing on the left side still have an incentive to compete in the same market as advertisers, an increase in quality on the left side may *reduce* advertisers' equilibrium prices. I analyze the circumstances under which this will occur and discuss the model's potential implications for antitrust policy.

Keywords: Search Engines, Economics of the Internet, Two-sided Markets, Monopolist Quality Choice, Media Bias

JEL Codes: D21, D42, D43, D83, L12, L13, L40

*Thanks to Tobias Kretschmer, Pai-Ling Yin and two anonymous referrees, as well as Eduardo Azevedo, Klênio Barbosa, Stefan Behringer, Guido Friebel, Avi Goldfarb, 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, Jean Tirole, and especially to Jacques Crémer and Glen Weyl. I gratefully acknowledge the financial assistance of Project 20121087916 supported by the Tsinghua University Initiative Scientific Research Program and Project 71203113 supported by the National Natural Science Foundation of China. This paper is based on a chapter of my Ph.D. dissertation at the Toulouse School of Economics and was first presented at the ZEW ICT Conference in 2008. I am fully responsible for any errors.

†Tsinghua University School of Economics and Management; awhite@sem.tsinghua.edu.cn

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 funded[1]

1 Introduction

In 2012, advertisers paid Google $44 billion, a large but undisclosed portion of which was for their links to appear among its search results. However, users are roughly five times more likely to click on links that the search engine selects by itself than on those for which advertisers pay.[2] Thus, a full economic view of search engines must incorporate understanding of unpaid as well as of paid links.

Systematic observation and analysis of unpaid search results poses significant challenges, however. In particular, as search engines get more sophisticated, variation in their results gets harder links attract more people to use the search engine. On the other hand, high quality organic links compete with sponsored links and can thus undermine the search engine's profitability. While this tradeoff may seem similar to one facing any provider of ad-supported media, it is, in fact, more complex. This is because all of the links on a results page focus on roughly the same subject, and many of them are essentially competing with one another to sell a particular product or service to users. Consequently, understanding the tradeoff facing a search engine requires taking into account this market in which users and websites interact.

This paper develops a model to clarify this tradeoff facing search engines and to explore its economic implications. In the model, users visit a search engine both to learn about a particular good and to locate merchants that sell it. Merchants, meanwhile, can pay to have a link to their website appear on the right side of the results page. The equilibrium price of the good depends on the number of merchants. This, in turn, depends on two choices the search engine makes. One of these choices is how much to charge merchants to appear on the right side,[4] and the other is the quality level of the organic results. A novel feature of the model is that additional non-paying merchants may appear among the organic results, thereby influencing the market for the good.

Specifically, the paper addresses and provides findings on two main questions. First, it asks what distortions arise when there is a profit-maximizing monopoly search engine that chooses the level of search quality to offer its users. There are two. One is a market power distortion. For a search engine to earn profits, its results page must allow advertisers to earn rents by charging high prices for the goods they sell – rents which the search engine can then extract. The other is a distortion in the quality of results that is not mentioned by Brin and Page but is analogous to the one identified in Spence's (1975) analysis of a monopolist in a conventional market. When setting search quality, the search engine has incentive to cater to the tastes of only its marginal searchers, rather than to those of all of its users, as would be socially optimal.

The second question regards the comparative statics between the level of quality the search engine offers users and the price of goods located via the search engine. Here I find that the results hinge crucially on the details of the model for search quality. In the "simplest possible" model, the search engine improves quality by investing more money. In this case, price levels for goods

[4] For simplicity, I assume merchants are homogenous and thus do not model an auction mechanism of the type that search engines use in reality. Below, I mention several papers that focus on such mechanisms.

and quality levels tend to be complements for the search engine. That is, when the search engine increases quality, it will typically find it optimal to increase the level of market power its advertisers exert. A subtler version of the model allows non-paying "left side" merchants to appear among the organic results. In this case, price levels and quality levels can easily become substitutes. This issue of whether the search engine perceives price and quality levels as complements or substitutes is of particular significance in analyzing the impact on consumer surplus resulting from changes in the underlying environment. If the two are complements, a shock that causes an increase in search quality can end up hurting some or all users. If they are substitutes, such a shock leaves all users weakly better off.

Following the above analysis, I use the model to shed light on the debate surrounding Google's 2007 acquisition of the "display advertising" firm, DoubleClick. In particular, I illustrate the potential implications of the aforementioned comparative statics results regarding price and quality. In conventional markets, a standard argument implies that, the more substitutable two firms' goods are, the less likely it is that they should be allowed to merge. This argument is shakier, however, when the two goods in question are two different forms of advertisement. I make the case that, with such firms, the opposite relationship may in fact be true: because of the impact it will have on both the quality of organic results and the prices for goods that users face, a merger between a search engine and another seller of ad space is *more* desirable to policymakers if the ad space the two firms sell is substitutable in the eyes of merchants.

The paper is organized as follows. The remainder of this section describes the way that the current work fits into the literature. Section 2 sets up the basic model and illustrates the first, *market power* distortion. Section 3 endogenizes search quality using a simple investment model. It illustrates the second, *quality* distortion. It then analyzes the equilibrium comparative statics under this version of the model. Section 4 then introduces non-paying "left side" merchants into the model and shows their tendency to reverse the comparative statics found in section 3. Section 5 discusses the Google/DoubleClick merger. Finally, section 6 concludes.

Related Literature

While this work adds to a growing body of research on Internet search engines and web advertisement, it is the first to consider the relationship between organic results and the market power of advertisers. 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 (2010), Milgrom (2010) 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), de Cornière (2011), Eliaz and Spiegler (2011), Ellison and Ellison (2009) and Taylor (2012) concentrate on this issue, with the latter most closely related to this paper in that it considers both paid and unpaid links. A set of papers that incorporate both sponsored link auctions and consumer search into unified theoretical frameworks are Athey and Ellison (2011), Chen and He (2011) and Rayo and Segal (2010), while Ghose and Yang (2009) study empirically the way such interaction depends on the search term entered by the user. Pollock (2010) and Telang et al. (2004) study the organization of the search engine industry.

This paper also contributes to the literature on two-sided markets.[5] 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 (2011) and Nocke et al. (2007)) and investigates the platform's incentives to provide quality to consumers. More generally, search engines are a form of ad-supported media, studies of which 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.

[5] For more general coverage of this topic, see Caillaud and Jullien (2003), Evans (2003), Rochet and Tirole (2003), Rochet and Tirole (2006), Armstrong (2006), Weyl (2010). Also see Gomes (2012), which provides a connection between the two-sided markets literature and the aforementioned auctions.

2 A Simple Model

This section develops a maximally straightforward model of a monopolistic search engine. In doing so, it illustrates the relationship between the distortion caused by a search engine with market power and that caused by a conventional monopoly. It also establishes the basic framework that will be used throughout the paper.

A search engine is a platform that earns 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,[6] 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 the number of merchants that 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.[7]

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 learn about the item and figure out whether or not it is of interest. In the model, query costs are heterogenous across users. This can be interpreted in two ways: it can reflect the idea 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 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 it's worth) requires more effort. This section takes query costs to be exogenously given, and subsequent sections relax this assumption.

[6]In addition to using the terms in a user's query, search engines 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, e.g., Krause and Horvitz (2010).

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

2.1 The Setting

The model has three kinds of players: a "search engine", an arbitrarily large set of homogenous "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 $[\underline{\theta}, \overline{\theta}]$, where $0 \leq \underline{\theta} < \overline{\theta}$. 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 $[\underline{v}, \overline{v}]$, where $0 \leq \underline{v} < \overline{v}$. I assume that the two distributions are independent of one another. User i's (risk-neutral) utility function is given by

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

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. 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.

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.[8]
5. Users who searched either buy or don't buy the good from an advertising merchant.

Note that I ignore the constraint that the number of firms should be an integer and, instead, treat this as a continuous quantity throughout the rest of the paper.[9]

2.2 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.[10] Define the function $D(\cdot) \equiv 1 - G(\cdot)$. In the fourth-stage Cournot game, n advertisers face demand given 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 the well known formula,

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

where ε_D is the price-elasticity of demand. Let $p^*(n)$ denote the equilibrium price of the good 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-quantity pair that emerges from the merchants' oligopoly 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

[8] The assumption of Cournot competition 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. All of the following results carry through under a more general specification, except for the explicit solution for the equilibrium number of firms. On this issue, see Ellison and Ellison (2009), including the supplemental material associated with that article, which discusses this issue and gives a different, somewhat richer form of competition that could be used here to the same effect. Also see de Cornière (2011) and Eliaz and Spiegler (2011), both of which develop models in which sellers retain market power under competition due to an underlying consumer search process. Note, however, that the current model cannot be seen as a "reduced form" representation of these models, as they involve heterogeneous sellers and thus are concerned with quality of the matches between particular sellers with particular buyers.

[9] In assuming the equilibrium price in an oligopoly game to be a smooth, decreasing function of the number of firms, the current analysis follows in the spirit of Mankiw and Whinston's (1986) study of free entry.

[10] I also ignore any possible issues of multiple equilibria arising in the "coordination game" between users and merchants in stage 2. In reality, search engines typically use "pay-per-click" pricing, which helps to overcome this problem. A theoretical analogue to such pricing mechanisms, which can justify the approach taken here, are "insulating tariffs" (Weyl, 2010; White and Weyl, 2012).

can view the search engine as if it were solving a monopoly pricing problem. The following lemma establishes this.

Lemma 1. *The problem facing the monopoly 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*

$$\max_p m(p)D(p)\left(p - c\right), \tag{1}$$

where the size of the mass of users who search is $m(p) = F\left(\int_p^{\overline{v}} D(x)\,dx\right)$.

Proof. See Appendix A.1.

Setting the problem up in this way gives rise to a clean illustration of the distortion stemming from the search engine's market power, given by Proposition 1.

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

$$\frac{p^* - c}{p^*} = \frac{1}{\varepsilon_m + \varepsilon_D}, \tag{2}$$

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

Proof. See Appendix A.2.

One way to view this result is as a simple application of Lerner (1934)'s classic monopoly pricing rule, where the firm's demand function is given by $m(p)D(p)$. However, given the timing of the game and the presence, among users, of both a searching margin and a buying margin, a more detailed interpretation can be of interest. To this end, consider, as a benchmark, the problem facing a total surplus-maximizing social planner. His objective function can be written

$$\max_m m\left(\int_0^{Q(m)} P(x)dx - cQ(m)\right) - \int_{\underline{\theta}}^{\theta(m)} xf(x)dx, \tag{3}$$

where $P(\cdot)$ is the inverse function of $D(\cdot)$, and $Q(m)$ denotes the same quantity as $D(\cdot)$, expressed as a function of m, the size of the mass of users who find it optimal to search. The derivative of (3) with respect to m is given by $(Q + mQ')(P - c)$, where all terms in the left hand brackets are strictly positive. This implies that total welfare is maximized when $p = c$.

Therefore, the monopolist search engine resembles a conventional monopolist in that it sets a price that restricts efficient trade. Nevertheless, compared to a conventional monopolist selling in a market whose demand is given by $D(\cdot)$, the search engine induces a lower price, because it takes into account the additional margin of users who decide whether or not to query. Moreover, in order to implement its desired "monopoly" price as prescribed by (2), it finds it optimal to foster *competition* among advertisers. The following corollary translates this intuition into the search engine's choice for the number of merchants.

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*

$$n^* = 1 + \frac{\varepsilon_m}{\varepsilon_D}. \tag{4}$$

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.[11] 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 standard Lerner formula, $\frac{p-c}{p} = \frac{1}{\varepsilon_D}$. Here, however, the mass of demand, $m \;\; p$, that is present in the market depends on the anticipated equilibrium price. Consequently, the intensity of competition the search engine 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.*

3 Endogenous Search Quality

Building on the model described above, I now examine the search engine's choice of quality. In this section, I represent this decision in the simplest way possible. Specifically, I assume that the search engine must choose an investment level and that this choice will determine the distribution of users' query costs.

[11] For instance, when both θ and v are uniformly distributed, $\varepsilon_m/\varepsilon_D = 2$, and thus the search engine induces three merchants to enter.

This approach provides a reasonable description of various relevant issues facing the search engine. One important characteristic of Google's search product is that it typically displays results very quickly, and this speed is determined by factors such as the number and location of data centers it operates and the efficiency of its software. For example, in 2010, Google introduced a feature called "Google Instant", which predicts the query that users are in the midst of typing and displays results before they have finished. Along similar lines, search engines often display information that is relevant to a particular query at the top of a results page (in addition to providing links to other websites). The range of queries for which a search engine does this and the accuracy of the information it shows naturally affect users' query costs, but developing such technology is costly.

A key issue that I ignore in this section is the potential for direct competition between organic links and paid links. While such competition

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

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

Under this specification, the individual characteristic of user i that influences her query cost, θ_i, (hereafter, her "type") is still drawn from an interval $[\underline{\theta},\overline{\theta}]$ with cdf $F(\cdot)$. Now, however, her query cost also depends negatively on search quality. Formally, I assume, $\frac{\partial\varphi}{\partial\theta_i}>0>\frac{\partial\varphi}{\partial s}$.

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 The Quality Distortion

I now compare the search engine's incentives to set quality with those of a total surplus-maximizing planner, holding fixed m, and thus the set of users who search. This exercise, which is a natural extension of Spence's (1975) analysis of a quality-choosing, "traditional" monopolist,[13] makes sense to consider in the context of Brin and Page (1998)'s claim, quoted at the beginning of the paper, regarding the inherent bias of an ad-funded search engine. We can write the search engine's profits as

$$mQ\left(m,s\right)\left(P\left(Q\left(m,s\right)\right)-c\right)-h\left(s,\alpha\right), \tag{5}$$

where $Q\left(\cdot,\cdot\right)$ is defined analogously as before but now depends on s as well as m. Differentiating (5) with respect to s yields first-order condition

$$m\frac{\partial Q}{\partial s}\left(P-c+QP'\right)=\frac{\partial h}{\partial s}. \tag{6}$$

[13] Also see Sheshinski (1976).

Since m is held fixed, the type, θ^*, of the marginal searcher is as well. Hence, $\partial Q/\partial s$ can be found by differentiating both sides of the equation $\varphi(\theta^*, s) = \int_0^{Q(m,s)} P(x)\,dx - P(Q(m,s))\,Q(m,s)$, which gives $\partial\varphi/\partial s = -P'Q \cdot \partial Q/\partial s$. Moreover, since $-P'Q = D/(-D')$, equation (6) can be written

$$-\frac{\partial\varphi(\theta^*, s)}{\partial s} - \left(-\frac{\partial\varphi(\theta^*, s)}{\partial s}\frac{-D'}{D}(P-c)\right) = \frac{\frac{\partial h}{\partial s}}{m}. \tag{7}$$

The socially optimal level of quality is found by differentiating total surplus,

$$m\left(\int_0^{Q(m,s)} P(x)\,dx - cQ(m,s) - \int_{\underline{\theta}}^{\theta^*(m)} \varphi(x,s)\,f(x)\,dx\right) - h(s,\alpha), \tag{8}$$

with respect to s. Plugging in the same quantities as above gives first-order condition

$$-\frac{\int_{\underline{\theta}}^{\theta^*(m)} \frac{\partial\varphi}{\partial s} f(x)\,dx}{m} - \left(-\frac{\partial\varphi(\theta^*, s)}{\partial s}\frac{-D'}{D}(P-c)\right) = \frac{\frac{\partial h}{\partial s}}{m}. \tag{9}$$

Equations (7) and (9) differ only in their first terms. In (7), this term measures the benefit to the marginal searcher of a small, "one-unit" increase in search quality. In (9), this term measures the average benefit of such an increase across all searchers. The second term in both expressions measures the reduction in gains from trade from marginal *buyers* that arises when search quality increases and the mass of users is held fixed. Holding fixed m, an increase in s necessarily implies an increase in the price of the good and thus a decrease in the probability that any user who searches indeed buys the good. This loss, however, is internalized by the search engine, since it is lost profit driven by fewer sales to users with valuations for the good, v, that are very close to the good's price, p.

The Spencian logic thus extends to this context in the following way. When the marginal *searcher* places higher-than-average importance on improvements in the quality of the search engine, the latter will have an incentive to over-invest, whereas when the marginal searcher is relatively indifferent, it will have an incentive to under-invest. These results suggest that, independently of Brin and Page's claim, which applies to the point about market power of advertisers, analyzed in section 2, a search engine's bias in its choice of quality level will be in favor of marginal searchers (e.g. novice Internet users; see section 3.4) over inframarginal ones (e.g. expert users).[14]

3.3 Comparative Static Analysis

I now address the question of whether search quality and the price of the good are complements or substitutes for the search engine. To address this question, it is convenient to write the search engine's profit maximization problem as

$$\max_{s,p}\; m(s,p)D(p)\left(p-c\right) - h(s,\alpha). \tag{10}$$

Here, 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. Proposition 2a characterizes the comparative statics between p and s.

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 its optimal price, p^*, the search engine responds by adjusting the unaffected variable in the same direction if and only if*

$$\varepsilon_m > \varepsilon_{\frac{\partial m}{\partial s}}\,. \tag{11}$$

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.

[14] This basic point seems quite robust. However, it also seems likely that, in a model with richer preferences for both users and advertisers, the search engine will have an incentive to bias its design towards the *particular* marginal users who are most valued by *marginal advertisers*. See Veiga and Weyl (2012) for a framework that could be useful for extending this model to study such effects.

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 increased 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 (11) 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.

3.4 Interpretation in Terms of Query Costs

In this subsection, I state Proposition 2b, which enhances the 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,

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

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

$$\varphi^{-1}\left(\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}, s\right) = \theta^*. \tag{12}$$

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 price[15] and increasing in quality.

Using this notation, the mass of users is given by

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

For the rest of this subsection, assume that the θ's are uniformly distributed over some positive interval: $\theta_i \sim U\left[\underline{\theta}, \overline{\theta}\right]$. I can then derive the result given by Proposition 2b.

[15]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.

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

$$\varepsilon_{\varphi^{-1}} + \varepsilon_{\frac{\partial\varphi}{\partial\theta}} < 0, \tag{13}$$

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

Proof. See Appendix B.2

To interpret (13), 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,

$$\frac{\partial}{\partial s}\left(\frac{\partial\varphi(\theta^*,s^*)}{\partial\theta}\right) = \frac{\partial}{\partial\theta}\left(\frac{\partial\varphi(\theta^*,s^*)}{\partial s}\right). \tag{14}$$

When (14) 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 a natural interpretation of users' type parameter, θ, to be their level of expertise at using a search engine. Following this interpretation, we can think of high types as "novice" searchers. Thus, as Figure 1 illustrates, Proposition 2b implies that unless quality changes disproportionately affect novice users, quality and price are complements.

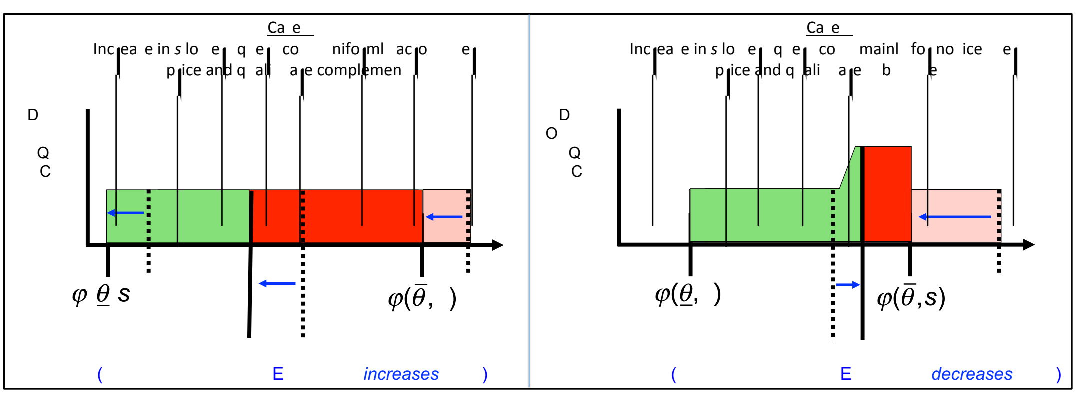

Figure 1: Graphical intuition underlying Proposition 2b

A Specific Query Cost Function. To supplement the intuition given by Figure 1, I give an example to illustrate, in a formal yet concrete setting, how the *relative* impact on different users of an increase in search engine investment affects the complementarity of price and quality. For this purpose, a useful exercise is to compare a query cost function under which an increase in search quality lowers all users' query costs by the same amount with one under which such an increase lowers them differentially. A simple specification that allows for such a comparison in a continuous parameter space is given by the query cost function $\varphi(\theta_i, s) = \frac{\theta_i - s^{1-\gamma}}{s^\gamma}$, 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 novice (large θ) users than for expert (small θ) users.

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

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

As the denominator is unambiguously 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 (i.e., a 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.

4 Competition from the Left Side

In this section, I focus on the fact that a search engine's content consists largely 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.[16] 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.[17]

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 assume 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 Proposition 3 shows, left side merchants serve as an additional force driving down the price of the good that the search engine induces. Second, as Proposition 4 illustrates, unlike in the case without left side merchants, here 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, 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.

[16] Indeed, 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.

[17] Websites of all descriptions can be found to have some space devoted to advertising, of which two forms 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.

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

$$\max_{A,s} r(A,s)\frac{m(s,p)D(p)(p-c)}{l(s)+r(A,s)}, \tag{15}$$

where $p = p(l(s) + r(A,s))$ is a function of the number of merchants on both sides. From (15), 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 directly with respect to the induced price of the good. In order to do this, let $\lambda \equiv l/(l+r)$ denote the proportion of total merchant profits that the search engine does not extract because they are earned by merchant websites appearing among the organic results. Since the function $l(s)$ is a one-to-one correspondence, I 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

$$\max_{\lambda,p} (1-\lambda)\tilde{m}(\lambda n(p), p)D(p)(p-c), \tag{16}$$

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

$$\frac{p^*-c}{p^*} = \frac{1}{\varepsilon_{\tilde{m}}^{tot} + \varepsilon_D}, \tag{17}$$

which is similar to the formula of Proposition 1. Note that here, however, the first term in the denominator, $\varepsilon^{tot}_{\tilde{m}}$, is the total price-elasticity of the mass of users, defined as

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

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

$$\tilde{m}(\lambda^* n(p), p) = (1 - \lambda^*) \frac{\partial \tilde{m}}{\partial \{\lambda n(p)\}} n(p). \tag{18}$$

Combining the two first-order conditions in (17) and (18) leads to Proposition 3, which illustrates the effect of left side merchants on the price for the good that 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*

$$\frac{p^* - c}{p^*} = \frac{1}{\frac{\lambda^*}{1-\lambda^*} \varepsilon_n + \varepsilon_{\tilde{m}} + \varepsilon_D}, \tag{19}$$

where

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

Proof. See Appendix C.1.

To interpret (19), note that, in addition to $\varepsilon_{\tilde{m}}$ 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, weighted by λ^*

4.2 Comparative Statics with Left Side Competition

In the model of section 3, price and quality are complements, so long as changes in quality do not have too disproportionate an eff

4.3 Example with 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.[18] 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 $p = \frac{1}{n+1} \Leftrightarrow n = \frac{1-p}{p}$. This implies that the mass of users who search is

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

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

$$\max_{\lambda,p} \frac{1}{2\bar{\theta}}(1-\lambda)\lambda^{\gamma}(1-p)^{3+\gamma}p^{1-\gamma}.$$ [19]

Solving this problem, the optimal price, p^*, and the optimal total number of merchants to allow, n^*, are given, as functions of γ, by $p^*(\gamma) = \frac{1-\gamma}{4}$ and $n^*(\gamma) = \frac{3+\gamma}{1-\gamma}$. Note that the intensity of competition among merchants that the search engine induces is increasing in γ. Turning to the

[18] The 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.

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

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

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.

choice of search quality, the optimal proportion of merchants placed on the left side is $\lambda^*(\gamma) = \frac{\gamma}{1+\gamma}$. Thus, the optimal search quality is $s^*(\gamma) = \frac{\gamma(3+\gamma)}{1-\gamma^2}$. Meanwhile, the optimal number of right-side merchants is $r^*(\gamma) = \frac{3+\gamma}{1-\gamma^2}$. The expression for the search engine's equilibrium profits, as a function of γ, is $\frac{\gamma^\gamma(3+\gamma)^{3+\gamma}(1-\gamma)^{1-\gamma}}{2^9\bar{\theta}(1+\gamma)^{1+\gamma}}$. Figure 2 plots several of these results.

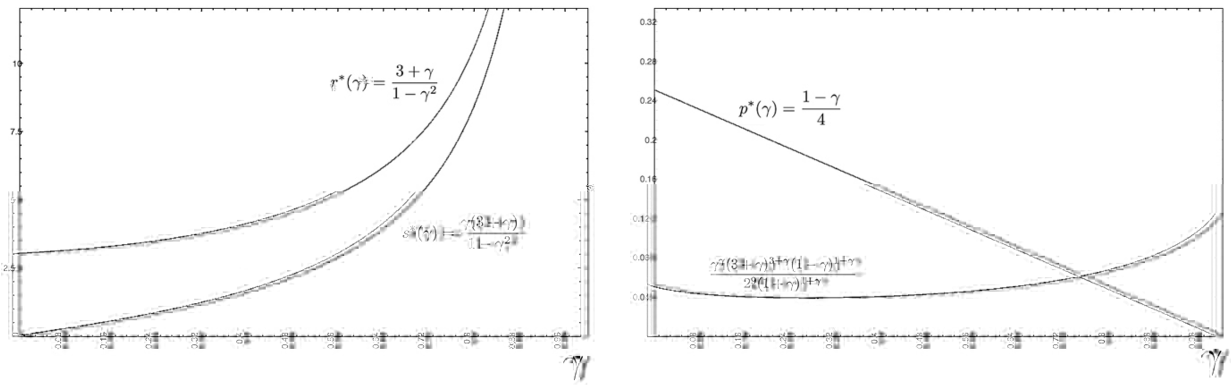

Figure 2: Left: number of advertisers, r^*, and quality, s^*; Right: price, p^*, and search engine profits (scaled by $\bar{\theta}$).

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,[20] 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. 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.

4.4 Discussion

In order to interpret the above more broadly, it is important to consider the role of the model's stylized assumptions. In particular, while the homogenous goods framework I use provides significant analytical convenience, how might it drive the results? Relatedly, what is the impact of compressing all of the "learning" in the model into the update of a single valuation that each consumer experiences

[20] An 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.

upon searching? I believe that the main message, that the more important competition is from organic results, the more search quality and advertisers' market power will tend to be substitutes, is quite robust.

One way to think about this is to imagine the following. As users go about clicking through the results of a search page, on each page they view, with some probability, they find the information that they need to know in order to be ready to make a purchase. (For example, "Is this the correct replacement cartridge for my sink's water filter?") If the search engine can display sites that are totally noncommercial that allow the user to accomplish this, then, as the search engine improves its quality, the market power of the advertisers is preserved. If, on the other hand, the web is such that most sites that contain such information (e.g., blogs or forums about home hardware) also contain their own links to sites that sell the relevant goods, then, as the search engine improves its quality, the market power of the search engine's advertisers is eroded. This, of course, raises a point about a search engine's incentives also to sell ads on non-search page, which I discuss in the next section.

5 Application: Google's Acquisition of DoubleClick

This section briefly considers the possible implications of the above comparative static analysis on the economics of acquisitions by search engines of other firms operating in the ecosystem of online advertising. For concreteness, I focus on the prominent, 2007 case of Google's purchase of DoubleClick.[21] An issue in the antitrust debate surrounding that acquisition was whether advertisers consider as substitute goods search ads and non-search ads on other websites. Consistently with standard debating procedure for mergers in conventional industries, parties against the merger argued that these two forms of advertisement *are* substitutes, while parties in favor of the merger argued 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

[21] According to Wikipedia, between 2001 and June, 2012, Google acquired 114 firms. In the area of online advertising, more recently, in 2011, it paid a reported $400 million to acquire Admeld, a firm specializing in display advertising. See `http://en.wikipedia.org/wiki/List_of_acquisitions_by_Google`.

greater if the two forms of advertising were substitutes.

The model presented in this paper indicates, however, that possible harms to consumers resulting from such a merger 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, suppose that search ads and non-search web ads (hereafter just "web" ads) are never substitutes in the eyes of 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 (16),

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

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.

Clearly this discussion is no more than a thumbnail sketch of the relevant issues for thinking about such a merger. In particular, the above discussion focuses only on the direction of the model's comparative statics and not on how to quantitatively measure the effects of such a merger on total/consumer surplus.[22] However, at the very least, it illustrates the difficulties and potential perils of trying to make antitrust arguments related to this industry without seriously taking into account the structure of the market. Of particular importance is the basic point that, while a scarcity of *sponsored* links available for purchase will drive ad prices up, leading advertisers to complain, consumers likely care more about the contents of organic results.

6 Conclusion

Increasingly, a single search engine determines what information society sees, and, thus, it seems important that the incentives it 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 are selected by the search engine's algorithm, whereas others appear because they pay to be there. A search engine's profitability thus stems from

[22]Quantitatively predicting such changes in response to mergers, even in more "conventional" industries, is still a relatively immature topic. See Jaffe and Weyl (2013) for a recent contribution in this area, and see section 6.2 of White and Weyl (2012) for an extension of the former article's approach to analyzing mergers in platform industries.

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.

The relationship between its choice of quality and its sale of market power turns out to be

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

$$\pi_j(A,n) = \frac{m(p^*(n))D(p^*(n))(p^*(n)-c)}{n} - A. \tag{21}$$

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

$$\frac{m(p^*(\tilde{n}))D(p^*(\tilde{n}))(p^*(\tilde{n})-c)}{\tilde{n}} = \tilde{A}$$

$$\Leftrightarrow m(p^*(\tilde{n}))D(p^*(\tilde{n}))(p^*(\tilde{n})-c) = \tilde{n}\tilde{A} = \Pi_{SE}, \tag{22}$$

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 (22).

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

$$\underbrace{\left(E[v_i|v_i > p] - p\right) \times \Pr[v_i > p]}_{\text{Expected Consumer Surplus}} = \int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p} \geq \theta_i.$$

A.2 Proposition 1

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

$$\frac{m'(p^*)}{m(p^*)} + \frac{D'(p^*)}{D(p^*)} + \frac{1}{p^*-c} = 0. \tag{23}$$

Multiplying this equation by p^* and rearranging gives the expression in (2).

A.3 Corollary to Proposition 1

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

$$\frac{p^*-c}{p^*} = \frac{1}{n} \cdot \frac{1}{\varepsilon_D}. \tag{24}$$

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

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

$$\frac{\partial^2\Pi(s^*,p^*)}{\partial s,\partial p} = \left(\frac{\partial^2 m}{\partial p \partial s}D + \frac{\partial m}{\partial s}D'\right)\left(p^* - c\right) + \frac{\partial m}{\partial s}D > 0 \tag{25}$$

and substitutes otherwise, where s^* and p^* denote the optimal values of these variables for the search engine. Algebra yields that the condition in (25) 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.

B.2 Proposition 2b

Proof. In view of the result shown in proposition 2a, it suffices to show, that when $\theta_i \sim \mathrm{U}\left[\underline{\theta}, \bar{\theta}\right]$, 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

$$m(s,p) = F\left(\varphi^{-1}\left(\int_p^{\bar{v}} D(\tilde{p})\, d\tilde{p}, s\right)\right). \tag{26}$$

Hence, when θ is uniformly distributed, we have

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

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

C Proofs for Section 4: Left Side Merchants

C.1 Proposition 3

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

$$\varepsilon^{tot}_{\tilde{m}} = -\left[\frac{\partial\tilde{m}(\lambda^* n(p^*),p^*)}{\partial\{\lambda n(p)\}}\lambda^* n'(p^*) + \frac{\partial\tilde{m}(\lambda^* n(p^*),p^*)}{\partial p}\right]\frac{p^*}{\tilde{m}(\lambda^* n(p^*),p^*)}, \tag{27}$$

and that, from the first-order condition with respect to λ, (18), we have

$$\frac{\partial\tilde{m}(\lambda^* n(p^*),p^*)}{\partial\{\lambda n(p)\}} = \frac{\tilde{m}(\lambda^* n(p^*),p^*)}{(1-\lambda^*)n(p^*)} \tag{28}$$

Inserting the right-hand side of (28) into (27) and simplifying gives the result.

C.2 Proposition 4

Proof. The search engine's first-order condition, with respect to λ, (18), implies that

$$\Phi(l,p) \equiv (n(p)-l)\frac{\partial\tilde{m}(l,p)}{\partial l} - \tilde{m}(l,p) = 0 \tag{29}$$

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

$$\frac{\partial\Phi(l^*(p),p)}{\partial p} = (n(p)-l^*(p))\frac{\partial^2\tilde{m}}{\partial l\partial p} + n'(p)\frac{\partial\tilde{m}}{\partial l} - \frac{\partial\tilde{m}}{\partial p} > 0 \tag{30}$$

and substitutes otherwise. By (29) we have

$$n(p) - l^*(p) = \frac{\tilde{m}(l^*(p),p)}{\partial\tilde{m}/\partial l}, \tag{31}$$

Using (31), the condition in (30) can be rewritten

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

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

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 (2011), "Position Auctions with Consumer Search." *Quarterly Journal of Economics*, 126, 1213–1270.

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.

Börgers, 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 (2011), "Paid Placement: Advertising and Search on the Internet." *Economic Journal*, 121, 309–328.

de Cornière, Alexandre (2011), "Search Advertising." *Mimeo.*, URL `https://sites.google.com/site/adecorniere/Home/research`.

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.

Eliaz, Kfir and Ran Spiegler (2011), "A Simple Model of Search Engine Pricing." *Economic Journal*, 121, 329–339.

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.

Evans, David S. (2003), "The Antitrust Economics of Multi-Sided Platform Markets." *Yale Journal of Regulation*, 20, 325–431.

Ghose, Anindya and Sha Yang (2009), “An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets.”

Spence, A. Michael (1975), "Monopoly, Quality, and Regulation." *The Bell Journal of Economics*, 6, 417–429.

Taylor, Greg (2012), "Search Quality and Revenue Cannibalisation by Competing Search Engines." *Journal of Economics & Management Strategy*, Forthcoming, 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.

Veiga, André and E. Glen Weyl (2012), "Multidimensional Product Design." *Mimeo.*, URL `http://ssrn.com/abstract=1935912`.

Weyl, E. Glen (2010), "A Price Theory of Multi-Sided Platforms." *American Economic Review*, 100, 1642–1672.

White, Alexander and E. Glen Weyl (2012), "Insulated Platform Competition." *Mimeo.*, URL `http://ssrn.com/abstract=1694317`.