

Review of Network Economics

Volume 9, Issue 4

2010

Article 4

Is Google the Next Microsoft: Competition, Welfare and Regulation in Online Search

Rufus Pollock, *University of Cambridge*

Recommended Citation:

Pollock, Rufus (2010) "Is Google the Next Microsoft: Competition, Welfare and Regulation in Online Search," *Review of Network Economics*: Vol. 9: Iss. 4, Article 4.

DOI: 10.2202/1446-9022.1240

Is Google the Next Microsoft: Competition, Welfare and Regulation in Online Search

Rufus Pollock

Abstract

The rapid growth of online search and its centrality to the ecology of the Internet pose important questions: why is the search engine market so concentrated, and will it evolve towards monopoly? What implications does this concentration have for consumers, search engines, and advertisers? Does search require regulation, and, if so, in what form? This paper supplies empirical and theoretical material with which to examine these questions. In particular, we (a) show that the already large levels of concentration are likely to continue; (b) identify the consequences, negative and positive, of this outcome; and (c) discuss the regulatory interventions that policy-makers could use to address these.

KEYWORDS: search engine, regulation, competition, antitrust, platform markets

Author Notes: Rufus Pollock, Emmanuel College and Faculty of Economics, University of Cambridge; email: rp240@cam.ac.uk or comments@rufuspollock.org. I thank the organizers of 2007 IDEI Toulouse Conference on the *Economics of the Software and Internet Industries*, whose search round-table greatly stimulated my thoughts on this topic as well as participants at seminars in Cambridge and at the 2009 RES conference. Particular thanks are also due to the editor and an anonymous referee who provided an excellent set of comments that served to greatly improve the paper. Finally, I would like to thank www.searchenginewatch.com and Hitwise who both generously provided me with access to some of their data. This work is licensed under a Creative Commons Attribution License (v3.0, all jurisdictions).

1 Introduction

Internet search (or perhaps more accurately ‘web-search’) has grown enormously in recent years, rising in line, or even faster, than the general development of the Internet and World-Wide-Web.¹ Beginning from practically nothing 12 years ago, today search is a multi-billion dollar business. Search engine providers such as Google and Yahoo! have become household names² and use of a search engine, like use of the web, is now a part of everyday life.

As the amount of information pouring onto the web has grown, the utility, importance, and power of search engines has grown concomitantly: with ever more information available, a user is faced with finding a ‘needle’ in an ever larger ‘haystack’ – and has therefore become ever more dependent on the filtering facilities provided by search engines. With this process of information accumulation showing little sign of slowing, let alone stopping, the continued growth of search engines, and their importance, seems assured.

Apart from its wider societal importance there are several noteworthy features of the search engine *business*. Most importantly, the fact that users (almost always) do *not* pay – that is to say, the service provided by web search engines are free (to use).³ Where then do web search engines find their revenue? In one word: advertising. When search engines provide ordinary users with a ‘free’ service they gain something very valuable in exchange: attention. Attention is a rival good, and one in strictly limited supply – after all, each of us have a maximum of 24 hours of attention available in any one day (and usually much, much less). Access to that attention is correspondingly valuable – and is likely to become ever more so – especially for those who have products or services to advertise. Thus, while web search engines do not charge users, they can retail the attention generated by their service to those are willing to pay for access to it. In so doing such companies have built multi-billion dollar businesses.

It is also noteworthy, that the skills and resources acquired in developing the basic search engine, particularly the skills in optimizing the selection of advertising to show, are now proving valuable outside of their original context. For example, by the second quarter of 2007, 35% of Google’s total revenue (\$1.35 billion)

¹It is important to remember that while the World-Wide-Web traffic now represents one of the largest sources of Internet traffic it is by no means the only one.

²While Google has been almost entirely search-focused throughout its existence, the same is not true of Yahoo! which has long positioned itself as a web ‘portal’, devoting substantially less attention to its search business.

³We make this qualification because the term ‘free’ particularly in the context of ‘free software’ or even, increasingly, ‘free’ services denotes something which is both ‘free’ to use but also which one is ‘free’ to copy and modify. Here, for clarity, where such a distinction needs to be drawn we will usually talk of an ‘open service’ or an ‘open system’.

came from provision of advertising on 3rd party sites via its AdSense programme, while 64% (\$2.49 billion) of its revenue came from sites it owned and operated.⁴ Similarly, in the same time period, 35% of Yahoo!'s total revenue (\$599 million of \$1,698 million) came from affiliates while just over 52% of Yahoo!'s revenue (\$887 million) came from sites it owns and operates.⁵

Another major feature of the search engine market is its high levels of concentration. As of August 2007 the top four search engines had a combined market share 97% in the US with the top firm (Google) having 65%.⁶

The rapid growth of online search, its concentration and its growing centrality to our societies raise a variety of important questions for economists to answer. Why is the search engine market so concentrated? Will concentration increase or decrease over time, and will a single firm come to dominate the market? What are the implications for different 'players' (consumers, search engines, advertisers) both under the current market structure and under its likely future evolution? Does the fact that search engines act as 'information gatekeepers', determining, in effect, what can be found on the web, mean that there may be need for regulation quite apart from standard commercial and welfare considerations? Finally, what issues does the search market raise for antitrust/competition policy? Specifically does the search market require regulation, and, if so, in what form?⁷

This article addresses several of these questions. In section 2 we provide empirical evidence on the levels of concentration in the search engine market, both over time and across jurisdictions. This data clearly shows that the search engine market is indeed highly concentrated and has grown more so over time. Sections 3, 4 and 5 form the core of the paper. In section 3 we introduce a basic model of the search engine market and use it in section 4 to explain why the search engine market is so concentrated – and likely to grow even more so. In addition, we discuss in detail the question of contestability – that is whether the market might remain competitive (contestable) even if one firm were very dominant. We suggest that there are a variety of reasons why, even if one thinks the market is contestable now, it is likely to grow less so over time.

⁴See http://www.google.com/intl/en/press/pressrel/revenues_q207.html, visited 2007-09-24.

⁵Yahoo! Q2 2007 Earnings release available online at: <http://yhoo.client.shareholder.com/results.cfm>

⁶Concentration in other markets was if anything even higher. For example in the UK Google held over 80% market share as of August 2007. More details on market shares and their changes over time are in Section 2 available below.

⁷Additionally web search provides a fascinating case study for a student of technology and innovation. After all web search is clearly a new product, and one which is developing and evolving rapidly, with very large R&D spends by the major players.

This motivates the second piece of theoretical analysis in section 5. Building on the framework of the previous sections, we introduce social welfare and use it to analyze the performance of a monopolist. We show that monopoly can result in either over-provision or under-provision of quality relative to the social optimum. However, as we discuss, there are various reasons why it is more likely that under-provision occurs. In particular, we identify two particular effects, ‘substitution’ (organic results substitute for paid ones) and ‘antagonism’ (organic results may provide information that deter people from using paid ones), which both unambiguously operate to reduce the monopoly-provided level of quality compared to the socially optimal one.

This conclusion that a monopolist is likely to under-provide quality – whether relative to the social optimum or a more competitive environment – leads naturally into the last section of the paper which discusses possible actions to address this deficiency. We argue that the evidence on increasing concentration and the theoretical results earlier in the paper suggest that some form of intervention is needed. However, the informational and legal difficulties of direct regulation are substantial. We therefore focus on the indirect approaches a policy-maker could take. In particular, we point out that search engines have a natural division into ‘service’ and ‘software’ sections, with large competitive and technological differences between the two (in particular, the former has much greater resemblance to a natural monopoly than the latter). This suggests analogies with experience in other utility markets such as telecoms and electricity where a similar upstream/downstream division have proved useful in the design of regulatory intervention.

1.1 Related Literature

Much related work, particularly in theoretical areas, is discussed later in the paper in the modelling sections. Nevertheless, we briefly discuss here some of the wider context in which this work is situated.

The majority of the existing literature focuses on the advertising side of search engines. For example, there is significant work on ad-auctions, e.g. Edelman, Ostrovsky, and Schwarz (2007), Varian (2007), work on seller’s strategies in an online, search-based environment, see e.g. Ellison and Ellison (2004), work on the impact of advertising (‘paid-placement’, ‘sponsored-results’ etc) on facilitating (or hindering) consumer search, see e.g. Chen and He (2006), Athey and Ellison (2007), White (2008).⁸

⁸Most of these papers are theoretical but there is also growing amount of empirical work, see e.g. Ghose and Yang (2007), Goldfarb and Tucker (2007).

With their focus on advertising many of these papers see Internet search as some form of improved ‘yellow-pages’.⁹ In particular, search engines are seen primarily as a way for consumers to find *commercial* services or products they want. This contrasts with the approach taken here where ‘organic’ results are primary with ‘paid’ or ‘sponsored’ links secondary – at least for users.¹⁰

Of course, search engines pay for providing the quality of their ‘organic’ results using money gained from ‘sponsored ones’ and hence the two parts are, in many ways, symbiotic. Nevertheless, it is important to keep in mind that the major benefits generated by search engines are in connecting people with information from which no commercial transaction is likely to result – at least in the near-term.¹¹ This point will be central to our analysis and it is this focus, together with the explicit attention we give to questions on market structure and welfare, which differentiate our analysis from much of this existing literature.

There has been some limited work more directly related to what is presented in this paper, particularly on issues of market share. Gandal (2001), did (very) early empirical work which examined changes in market share in the late 1990s. Telang, Mukhopadhyay, and Rajan (2004), probably the closest paper to ours in its theoretical approach, also looked at market share and sought to explain the persistence of low-quality firms in a market where prices are zero.¹²

2 Concentration in the Search Engine Market: The Data

As already mentioned, one of the most noteworthy aspects of the search market is the very high levels of concentration already evident. Table 1 gives data from Autumn 2007 on the share of major search engines in several different countries. As can be seen, the C_4 values (the combined market share of the top 4 firms) are

⁹We should mention here Baye and Morgan (2001), one of the first papers to formally analyze a ‘yellow-pages’ (information-gatekeeper) model in an online environment and which also connects this area directly into the older and larger literature on general consumer search.

¹⁰This ordering also reflects the initial development of search engines themselves in which ‘pure’ search came first.

¹¹One could argue that *all* search has *some* impact on commercial activities over the long-term – and clearly not all advertising is directed at stimulating purchases right now. However, in most cases, this connection is so tenuous that we feel it can be ignored.

¹²Evans (2008) is an interesting related work that looks at search engines in the context of a broader discussion of online advertising and examines, as we do below, the extent to which the search market exhibits a winner-takes-all dynamic.

Company	United Kingdom	United States	Australia	Hong Kong
Google	81.1	59.1	84.0	36.2
Yahoo!	3.9	19.3	3.2	33.1
Microsoft	4.1	7.7	5.8	3.2
Ask.com	2.7	2.8	0.0	0.0
Sogou	0.0	0.0	0.0	5.6
Baidu	0.0	0.0	0.0	2.7

Table 1: Percentage Market Shares of the Largest Search Engine Operators (in week ending the 29th of September 2007). Note that these figures amalgamate for a given operator both traffic to their local site and their generic one (e.g. both google.co.uk and google.com) and traffic across different site types (e.g. images and video as well as normal text search). Source: Hitwise.

over 90% in all jurisdictions except Hong Kong.¹³ Even more significantly, in all cases except Hong Kong, the market share of the largest operator is substantially larger than its nearest competitor, and in the UK and Australia this dominance has reached the point where the largest operator (Google) has over 80% of the market – a level an order of magnitude higher than its nearer competitor.¹⁴

Also interesting is the question of how market shares have evolved over time. Obtaining good (comparable) market share data over a reasonable period is difficult. In particular, in the late 90s and early 2000s the only information recorded was the number of visits to a particular website. Since many providers of search also ran ‘portals’ it can be difficult to distinguish pure search from simple visits. In addition, early data frequently only records the number of unique visitors a month rather than giving a breakdown of the number of hits and this can severely distort results since pure-search providers (such as Google) are much more likely to have multiple visits from the same user than more portal-like sites. Matters are further complicated by the fact that in the late 1990s and early 2000s many search *sites* had their search powered by a third-party provider. For example, up until 2004, Yahoo! did not have their own search engine but ‘bought-in’ results, first from Inktomi (up

¹³It may be useful here to compare recent data from China which put Baidu at over 60%, with Google in second place at around 26% and Yahoo! third at around 10% implying a $C_4 \geq C_3 = 96\%$ (see <http://blog.searchenginewatch.com/blog/080229-230636>).

¹⁴Perhaps even more significantly, Google’s market share among younger users (University and High School) is even greater: over 90% according to Hitslink (<http://marketshare.hitslink.com/articles.aspx>, retrieved 2008-03-10). Compared to the 60% figure estimated for the overall US market this indicates a much, much higher level of concentration among the future user population than among the present one.

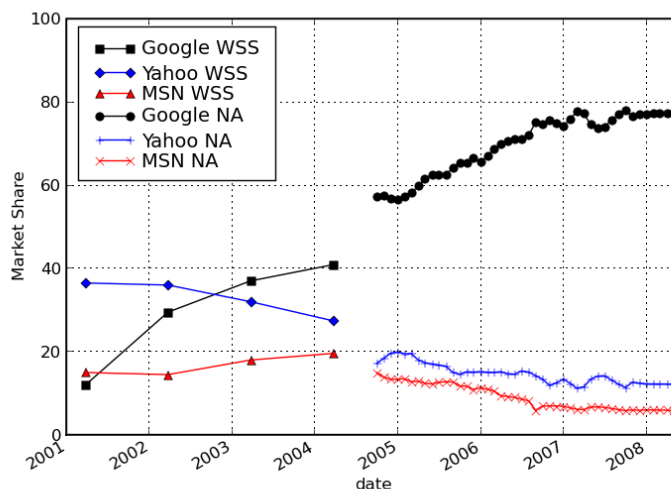


Figure 1: Search Engine Market Shares from 2001-2008 for the top 3 providers (as of 2008). NA indicates NetApplications data, WSS WebSideStory data. MSN figures incorporate both MSN site and Live.

until 2000) and then Google. Figure 1 is our effort to address this by combining data from NetApplications and WebSideStory (now part of Omniture). Both firms source their data from web analytic applications installed on customers' sites and NetApplications appears to be more global in its customer-base than WebSideStory (which may partially explain the non-exact match between the two datasets apparent in the 2004 values).¹⁵ The graph shows a simple story: a single firm (Google) emerges to dominate the market. In terms of general concentration, it is noteworthy that even in 2002, when Google was not yet as dominant as it is today, the top two firms (Google and Yahoo!) accounted for over 70% of the market while adding in Microsoft pushes this up to close to 90% (and of course at that point Yahoo!'s search was being powered by Google and MSN's by LookSmart and Inktomi).

Should these high market shares be cause for concern? After all, most competition/antitrust authorities, including for example the EU's, normally take a mar-

¹⁵This source of data differs from that found in the likes of Nielsen's NetRatings, comScore's MediaMetrix. Those products get their data from the users themselves (directly or indirectly via ISPs) rather than from websites they visit. In this sense they may be more reliable sources of data. However, it has proved difficult to obtain continuous time-series data for these providers for more than a couple of years – and for that period the trend they show is very similar to that found in the data shown.

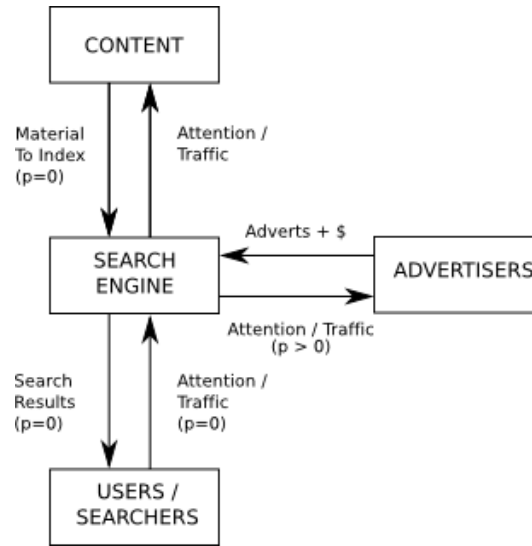


Figure 2: The structure of the search engine business.

ket share over 50% to be indicative of a dominant position. There are two distinct issues to consider in assessing whether high concentration is problematic: a) is the search market competitive even though one company (or a few companies together) has/have a very large market share b) even if the market is not competitive (in the extreme case a monopoly), given the structure of the search market and, in particular, the zero charges to search users, does a lack of competition imply harm to social welfare.¹⁶ To address these questions properly we need to develop a more detailed model of the search engine market and so it is to that task that we now turn.

3 Modelling the Search Engine Market

The core structure of the search engine market is displayed schematically in Figure 2. As can be seen it has a basic ‘three-sided’ aspect in which the search engine acts as a ‘platform’ intermediating between ‘content providers’ (who want ‘users’), ‘users/searchers’ (who want ‘content’), and ‘advertisers’ (who want access to ‘users’). Closely related to this structure of connections between agents is the associated pricing (and supply) structure – also illustrated in the Figure and which we will discuss further as we develop our model below.

¹⁶The search market is R&D intensive and so classic Schumpeterian arguments could be made that increased concentration will have a positive effect on R&D and hence on overall social welfare.

Given the three-sided nature of search, the literature on two-sided, platform, markets (see e.g. Rochet and Tirole (2003), Armstrong (2006), Nocke, Peitz, and Stahl (2007)) provide one obvious analytical toolset. However, the search engine business does not fit as comfortably as it might within this paradigm: the two primary groups a search engine sits between are ‘users’ and ‘content providers’ neither of whom pay to participate, while it is a third group ‘advertisers’ who are charged.¹⁷ This means that the central concern of a two-sided model, namely the pricing structure, is rather secondary since a price is only set for one of the three groups, and that, furthermore, with least relevance to the two-sided framework.¹⁸

The other strand of general literature is that related to oligopolistic competition, particularly models of Bertrand competition and vertical product differentiation Shaked and Sutton (1983), Sutton (1991). As we shall see this immediately provides some simple predictions (convergence to monopoly) which seem borne out by current data – though we will also discuss why the model is unlikely to fit exactly.

3.1 A Formal Model

There are four types of agents in our model: ‘users’ (U), ‘advertisers’ (A), ‘content providers’ (C), and search engines (S). We start with some terminology and basic assumptions:

1. The pool of material made available by ‘content providers’ is available to all search engines and is available for free. As such, ‘content providers’ can be ignored as (strategic) agents in this model leaving us to focus solely on the other three types.¹⁹

¹⁷Here, we are focused on ‘advertisers’ who advertise on the search engine only – not on any content provider. Several search engine companies also provide ‘ad-brokerage’ – selling advertising ‘space’ on search results also provides the tools (and customer base) to sell advertising ‘space’ on general sites. ‘Ad-brokerage’ fits very well in the two-sided framework as there the two sides (‘content providers’ and ‘advertisers’) care directly about the size of the other group and the ‘ad-broker’ naturally takes a platform role.

¹⁸There is some work on newspapers within the two-sided framework and newspapers do display a similar structure to search engines – users care about content and newspapers use advertising revenue (advertisers care about users) to create and improve content.

¹⁹This largely reflects the world as it is – search engines are able to freely index the vast majority of the web. Of course, there are some exceptions: some websites have restricted access to search engines, either because of concerns about caching and reuse or out of a desire to be remunerated for the content they provide (see e.g. Copiepress v. Google <http://www.groklaw.net/articlebasic.php?story=20070726152837334>).

2. Search engine quality is reflected in a single variable v , which all users value positively. That is, all attributes of a search engine such as the amount of material it indexes (positive), the up-to-dateness of the index (positive), the relevance of search results (positive), the number of advertisements (zero or negative), can be incorporated into a single overarching variable named ‘quality’.
3. Each ‘user’ uses a single search engine and it is the one that offers the highest utility. Note that it is straightforward, and perhaps even more logical, to interpret ‘users’ in this context as ‘usage’, that is as denoting individual queries not individuals themselves. Not only does this obviate debate about the correctness of assuming that individuals use a single search engine,²⁰ but it also fits better with the data – most market share information is measured in terms of usage (‘hits’/queries on the website), not as the share of individual users. Thus in what follows whenever one reads ‘user’ one can, if one wishes, substitute, ‘usage’ or ‘query’.

There are N search engines: S^1, \dots, S^N . Search engine i has quality v^i and charges price p_u^i to users. There are a continuum of potential users represented by the interval $[a, b]$ and indexed by t (without loss of generality we may take $a = 0, b = \infty$ and thereby map potential users one to one the positive real line). A user’s utility from using search engine i is given by:²¹

$$U_t^i = U_t(v^i, p_u^i) = u(t, v^i, p_u^i)$$

It is assumed that utility is increasing in quality for all users – $u(t, v^i, p_u^i)$ is increasing in v for all t .²² The user’s outside option will be normalized to 0 and users use the search engine which delivers the highest utility.

Total user demand for search engine i is q_u^i and corresponds to the set of users whose utility from search engine i is greater than their utility from any other search engine.

²⁰There is some degree of evidence that users do use multiple search engines. For example, in 2006 Search Engine Watch <http://searchenginewatch.com/showPage.html?page=3598011> report figures of Harvest Digital which showed that, of ‘experienced’ internet users, fully 20% regularly use four or more search engines. However it appears that most users use only one search engine most of the time.

²¹A specific form that is similar to that used in the vertical differentiation literature would be $U_t^i = \theta_t v^i - k_t - p_u^i$ where k_t is a user specific cost of using the engine, p_u^i is the price charged by search engine i to users and $\theta_t = \theta(t)$ is user-specific value for quality (assumed, wlog, to have $\theta' > 0$).

²²The form chosen implicitly assumes that there is no variation in the valuation of quality across search engines – that is users just care about the level of quality not which search engine it is associated with (note, however, that quality may of course be valued differently by different users).

$$q_u^i = \{t | U_t^i \geq 0, U_t^i > U_t^j \forall j\}$$

Formally, q_u^i is a set, however when no ambiguity arises, we may equate it with the measure of this set, i.e. the total number users using search engine i .²³

Finally, note that search engine user demand, q_u^i , will be a function of own quality, v^i and of price, p_u^i as well as all the qualities and prices of other search engines: $q_u^i = q_u^i(v^i, v^{-i}, p_u^i, p_u^{-i})$. For notational convenience, when no confusion will arise, we will frequently drop the search engine index i and write $U_t = u(t, v), q_u$ etc.

3.1.1 Advertising

Advertising will be modelled using a reduced form approach as follows. First let the advertising revenue generated by user t at search engine i be denoted by $a(t, v^i, q_u^i)$ – which becomes $a(t, v, q_u)$ without the i index. Total advertising revenue at search engine i is then given by the sum of this revenue across all users of that search engine:

$$R_A = \int_{q_u} a(t, q_u, v) dt = R_A(q_u, v)$$

3.1.2 Search technology

The total costs of a search engine are a function of quality, the number of users and the amount of advertising:

$$C = C(v, q_u, R_A) = C(v, q_u, R_A(v, q_u)) = \bar{C}(v, q_u)$$

It will be useful to divide C up into two parts as $C = c + c_A$ where $c = c(v, q_u) = C(v, q_u, 0)$ are ‘core’ or ‘user’ costs and $c_A(v, q_u) = \bar{C} - c$ are ‘advertising’ costs (i.e. those arising from managing ‘advertisers’).

Assumption. ‘Core/user’ costs are primarily fixed. In particular, the marginal cost of an additional user is assumed to be zero. Furthermore, the cost of supplying a given quality is (up to a point) independent of the number of users.²⁴

²³Being precise, one requires that an increase in the ‘size’ of demand involves a change from a set to a superset of that set (otherwise one may have the case of two sets of ‘users’ A, B with $|A| > |B|$ but, because of its composition, B is more valuable).

²⁴Search quality has several components. ‘Pure’ quality in the sense of the algorithm is nonrival across users and therefore has zero marginal cost for additional users (the costs of producing algorithm to make the index and rank results are one-off). The costs of maintaining the search service,

Justification: search engines are R&D intensive and the market generally displays high levels of innovation and obsolescence.²⁵ In addition, running a search engine service, quite apart from any R&D, is highly capital intensive. That is providing the hardware, support, monitoring etc to keep a search engine running, responsive and up-to-date requires a very significant investment, quite apart from any spending on R&D in order to improve the service. Both of these types of cost, whether related to R&D or the development and maintenance of service infrastructure, are largely fixed. At the same time the marginal cost of serving one additional user (or advertiser) is very low (almost zero in fact), especially when compared to these fixed costs. Taken together, this means that search engine cost structures display many of the characteristics of traditional (natural monopoly) utilities (on both the user and advertiser side of the market): very high fixed costs (both in terms of investment and direct supply) combined with very low (approximately zero) marginal costs.²⁶

Putting together the cost function and the revenue function we have that profits are given by:

$$\Pi = R_A - (c + c_A)$$

Finally, we make a major assumption about pricing which reflects the current reality of the search engine market:²⁷

Assumption. Search engines do not charge users: $p_u^i = 0$.

on the other hand, may marginal component – though costs of IT equipment and maintenance still have significant fixed costs there is a point at which increasing demand necessitates installing new servers, buying more bandwidth etc.

²⁵For example, Microsoft claimed to be spending over \$1bn a year on its online services (including its search engine) in 2006 (http://www.cbronline.com/article_news.asp?guid=3D810B1B-BBE0-482D-A81C-DBE60BAB97C4).

²⁶One factor not mentioned is the significant ‘learning-by-use’ component in search. Search engines learn heavily from the information provided by usage of their engine. For example from the click-through rates on both organic and paid (advertising) results and variation in this rates when different algorithms are employed.

²⁷While we do not seek to justify this outcome – it is taken simply as a reflection of reality – we observe that a zero-price to users is a common feature in several multi-sided platforms (e.g. shopping malls don’t charge consumers to enter). In this case, the choice will be driven by a combination of factors, most prominently: the highly effective method of charging advertisers for access to users’ ‘attention’ (and intentions), the (relatively) high sensitivity of advertising demand and revenue to the level of users and the large heterogeneity and uncertainty of the value of ‘queries’. We should also note that even if search engines do not charge users they could, in fact, *pay* them to use their search engine either directly or indirectly – e.g. a search engine could pay to ensure they were the default search option in a web browser. While direct payment is unlikely due to adverse selection indirect methods are a real possibility and Google, Yahoo!, Microsoft and Ask have all, at one time or another, paid either browser or computer manufacturers to have their search engine be made the default.

With this assumption, user utility becomes $U_t^i = u(t, v^i)$ and user demand for search engine i becomes $q_u^i(v^i, v^{-i})$.

4 Market Structure

In this section we formalize some of the intuitive arguments above regarding search market structure. Our basic result is that monopoly, or near-monopoly, is the likely outcome given cost and pricing structure of search. We supplement this formal result with an extensive discussion.

First, recall we can interpret search engine demand, q_u , as a scalar. This demand (taking other search engine qualities as constant) is a function of v . We can then invert and take v as a function of demand $v = v(q_u)$. Then defining $\bar{p}(q_u) = R_A(q_u)/q_u$ we have:

$$\Pi = \bar{p}(q_u)q_u - C(q_u)$$

This now looks like a classic vertical product differentiation problem in which \bar{p} now represents the price charged to a user (it is in fact the ‘derived’ price of a user in terms of advertising revenue). However there are some major differences, in particular $\bar{p}(q_u)q_u$ is guaranteed to be always increasing in q_u and it does not make sense to consider q_u as a function of \bar{p} .

Furthermore, users do not choose on the basis of price but on the basis of quality so there is no complementarity between quality and price (this would only occur here if one allowed the amount of advertising to negatively impinge on demand – in that case q_u would implicitly come to depend on \bar{p}). Specifically, as we assume that users are homogeneous in their taste for quality, our assumption of a fixed zero-price has converted the general vertical differentiation model into something very similar to a classic Bertrand setup with firms competing on quality instead of price (and higher quality being preferred by consumers rather than lower price).

Observation. Users will only use the search engine(s) with the maximum quality.

Proof. User t derives utility from search engine i :

$$U_t^i = u(t, v^i)$$

Thus, their utility from search engine i is greater than from j if, and only if, search engine i has higher quality and this holds independent of t : $U_t^i > U_t^j \Leftrightarrow v^i > v^j, \forall t$. Hence, any users (who is maximizing utility) will use only search engines with maximum quality, i.e. whose v satisfies $v \geq v^j, \forall j$. \square

In the case where several search engines offer this maximum quality we need to specify how market demand is divided. The simplest approach is to assume that all demand configurations are equally likely which implies that each of these search engines has equal (expected) revenue. To avoid trivial cases, we shall also make the following assumption:

Assumption (Basic profitability conditions). (a) firms with zero quality are inactive and earn zero profits (b) if there is only firm active, at least for one quality level $v > 0$ that firm can earn non-zero profits (i.e. it is profitable to supply search in the absence of competition from other firms).

Proposition 1. *Assuming continuity of costs in quality there is no (Nash) equilibrium in pure strategies of this simultaneous quality choice game.*

Proof. Let v be the maximum quality offered by a search engine. We must have $v > 0$ (if not, some firm can profitably deviate). Since provision of quality is costly for a search engine no search engine will offer quality in $(0, v)$ since they could either deviate to 0 or v and be strictly better off.

Assume that more than one search engine offers this top quality $v > 0$. Both must have non-zero market shares (if not then the one with zero market share must be making a loss since quality incurs a non-zero cost). Assume first that quality can be varied continuously in costs, i.e. for any $\delta > 0$ there exists an ε such that a firm can spend less than δ but increase its quality by ε . By ‘deviating’ in this way, one of the firms can offer quality $v + \varepsilon$ and thereby obtain complete market share with cost less than δ . Since for any quality increase above zero (and hence for ε) the gain in market share is equal to the combined market share of all other firms it is bounded below (at a level above zero). As (advertising) revenue is increasing in market share, then the gain in income from advertising is bounded below by some amount $A > 0$. Choosing an ε and δ such that $\delta < A$ we have that such a deviation is profitable and hence no equilibrium can exist in which more than one firm offers a non-zero quality.

Thus one firm offers non-zero quality v and garners all of the market. Let v^0 be the maximum quality such that the firm makes zero profits. Suppose this firm chooses $v' < v^0$ then another firm could enter with $v \in (v', v^0)$ and obtain positive profits (so not a NE). The firm must therefore offer v^0 . But, given that other firms are offering $v = 0$ this firm could deviate to another v and obtain positive profits and so this cannot be a NE either. \square

Remark. This problem is very similar to the problem of a R&D race with deterministic discovery functions (see Dasgupta and Stiglitz (1980)).

This non-existence result is largely the artefact of the strict simultaneity of moves and the discontinuity of payoffs it creates. It therefore makes sense to vary the setup by allowing one firm to ‘move first’ (a Stackelberg approach). We then have:

Proposition 2. *When one firm moves first (the leader) there is a single (pure-strategy) Nash equilibrium in which the leader offers a non-zero quality v and is the only search engine in use. All other search engines offer a zero quality level and have no users. The single active firm makes zero profits.*

Proof. One proceeds exactly as in the previous proof except that if the leader offers v^0 the threat of subsequent entry means that deviation is not a best-response and hence this is a Nash equilibrium. \square

4.1 Discussion

Clearly, in reality, the situation is rarely this simple and the result is rarely this stark. On the one hand, even with a very dominant firm, there are likely to be *some* other firms active in the market – i.e. a pure monopoly outcome is unlikely, and it would therefore be better to interpret this result not as predicting absolute monopoly but simply a single highly dominant firm. On the other hand, though there is monopoly, there is also ‘strong’ contestability in the sense that the active (monopoly) firm is constrained by the threat of competition to make zero profits and (associatedly) to supply the maximum feasible quality. Both predictions are central to any discussion on the competitiveness of the search market into the future. It is therefore important to consider how robust they are; in particular to evaluate whether they flow from a particular aspect of the formalism (e.g. the use of one-shot Stackelberg) or reflect deeper features of the general environment. We shall discuss each of these two items in turn.

4.1.1 Dominance

It is first worth recalling the main factors driving our formal result: (a) a cost structure which involves high fixed costs (for quality) and low marginal costs (serving additional users)²⁸ (b) pure quality competition for users (i.e. zero prices and no

²⁸Recall that this cost structure arises from two distinct aspects of the search engine model: economies of scale in the supply of the service itself, and the fixed costs of R&D. We have not distinguished these explicitly in our modelling since both contribute to the overall ‘quality’ of the experience.

user heterogeneity). In our view, any model which shares these basic features is likely to show very high levels of concentration and a single dominant firm.

In particular, high fixed cost/low marginal costs alone would imply a concentrated market. After all, as noted earlier, this cost structure is very similar to that of a classic ‘natural monopoly’ utility – a comparison that is all the more noteworthy given the basic, and crucial, infrastructural role that search engines play in the nascent ‘information society’.²⁹

This existing tendency to concentration is reinforced by the pricing structure: with a zero price, competition for users (and hence advertisers) takes the form of a winner-takes-all competition. It is this *lack* of competition on price that differentiates the current setup from the classic vertical differentiation models (see e.g. Shaked and Sutton (1983), Sutton (1991)) in which firms choose *both* quality and price. However it is noteworthy that those models, even with this price flexibility, often predict significant concentration, especially when quality (and the associated fixed costs) are ‘endogenous’ (as is the case, for example, with R&D and advertising).

Of course, the implicit assumption here is that there is a single (overall) ‘quality’ attribute which all users value positively (and that this is the only attribute differing across search engines). In reality, it is likely that there is some degree of heterogeneity across users. Brand preference is one obvious, though slightly nebulous, form of such heterogeneity. Another possibility is that search engines specialize in searching a particular kind of content.³⁰ However, any such heterogeneities are likely to be fairly limited compared to the general, homogenous, preference for ‘quality’ and, as such, unlikely to change the basic property of the existence of a single dominant firm.³¹

4.1.2 Contestability

It is not surprising then that the search engine market is already concentrated, and growing more so. However, might it still be competitive? As discussed in our model, the (credible) threat of entry means that although there is a single firm it

²⁹Just as access to, say, electricity is now considered essential, at least in most ‘developed’ countries, so we can imagine that, soon, access to the Internet and, therefore, to a search engine, will be an equally essential requirement.

³⁰For example, it is argued that part of Sogou and Baidu’s popularity come from their provision of a specific ‘MP3-search’ facility that allows users to easily search for music files on the Internet (most of which will be unauthorised copies – which perhaps explains the unwillingness of other search engines to emulate them).

³¹However, adding such ‘minor’ heterogeneities would allow the model to become more realistic by predicting the existence of several small, fringe firms.

behaves rather like it would under competition. Here, even though the fixed costs are large, because the game is static and deterministic, the threat of entry is credible. In reality, the market is dynamic with investments in quality (particularly those in R&D) being made sequentially. Thus, the question as to whether the dominant firm is insulated from the threat of competition by significant ‘barriers to entry’ is largely determined by how these dynamics interact with the large (sunk) fixed costs.³²

Generally, the question will revolve around the degree to which an incumbent can credibly ‘block’ entrants. This in turn depends on a variety of factors. Two of the most important will be (a) the size (and ‘sunkness’) of fixed costs; (b) the degree of (non-price, non-quality) ‘lock-in’ to an incumbent due, for example, to switching costs or ‘network effects’.

Let us take each of these issues in turn. First, fixed costs seem to be large and growing. Most of the major players have R&D spending in excess of \$500 million a year and the core infrastructure appears to be equally large. Furthermore, most of these incurred costs will be sunk: hardware and infrastructure have limited resale value (obsolescence is high) and the results of R&D will be highly search-specific. Hence, it would appear that, not only are the costs of entry large and growing, but that, facing the threat of entry, an incumbent can credibly commit to be ‘aggressive’ – for example via heavy R&D spend to improve quality.

On the lock-in question we focus on switching costs. If switching costs are high then, even if an actual or potential competitor offers a better quality product, they will find it hard to obtain market share (rapidly). The question of switching costs applies both to users and to advertisers as both are needed for a search engine to be successful. That said, one would expect that, if users switched, it would not be hard to persuade advertisers to switch as well, so it seems reasonable to focus on the user-side switching costs.

At first glance it would appear that switching costs are very low. After all, a search engine user can switch to an alternative by simply visiting a different website. However, switching costs may not be as low as they appear. In particular, there may be substantial brand effects as well as user adaptation to the behaviour of a particular search engine.

On the first of these points, a recent paper by Jansen, Zhang, and Ying (2007) examined the impact of brand on the evaluation of search results and found a significant impact. Specifically, they displayed an identical set of results through different ‘branded’ interfaces and elicited user evaluations of their quality (‘relevance’). Despite using these identical results they found a 25% difference in rating

³²For example, pursuing the analogy with the R&D literature, there are a variety of results (e.g. Harris and Vickers (1985)) which show that in a multi-stage race when the ‘leader’ has a large enough advantage even though ‘followers’ may exist (or could enter) the ‘leader’ can ignore this threat and behave like an (uncontested) monopolist – obtaining, for example, non-zero profits.

across engines. Along similar lines, it is interesting to note that there is significant geographical variation in search engine shares. Of course, a significant portion of this may reflect genuine heterogeneity in consumer tastes and in what search engines are offering. However, it is also likely that at least some of this reflects brand ‘stickiness’. For example Yahoo!’s core search system is likely to be the same in the UK and the US yet its market share is approximately five times larger in the US than in the UK (19.3% vs. 3.9%). Similarly, Google – the leaders in almost every other jurisdiction, trail Baidu (the first-mover) in China despite significant efforts on Google’s part.³³ While such jurisdictional heterogeneity, particularly where it relates to first-mover advantage, does not necessarily imply high switching costs,³⁴ it does, at the very least, imply that there are significant factors affecting market shares which do not arise straightforwardly from superior quality of service.

It is also important to note that an increasing number of users pursue fairly sophisticated query strategies, often refining (and refining again) their initial query if it fails to turn up what they are looking for. It seems likely (though not empirically tested, to our knowledge) that refinement strategies are search engine specific. As such, switching to a different engine is likely to involve some re-learning costs as a user adapts to the different search strategy required by the different search engine.

Moreover, an increasing number of search engines offer some form of explicit or implicit personalization. Such personalization, which could be used either to improve a user’s search experience or increase their value to advertisers, is clearly search engine specific. Naturally, this leads to increased switching-costs. Whilst these points are largely conjectural, there is some empirical evidence that users display increasing ‘loyalty’ to search engines. For example, a Jupiter Research study from 2006³⁵ looked at user behaviour when they did not find what they were looking for with their first query. They found that 41% tried again (compared to just 28% four years earlier in 2002). Of these 82% refined their query on their existing search engine and 18% switched engines, whereas four years earlier only 68% stayed with their existing engine (and 32% switched).

³³In most jurisdictions, Google should be considered the original ‘first-mover’ in which it has a lead despite not being the first to enter as all other companies to pre-date it in the search market either were not focused on search itself (for example Yahoo!) or fell out of contention before the importance of search (qua search) was recognized (e.g. Altavista).

³⁴For example it fits comfortably within the escalation models of Sutton, and in fact Sutton (1991, 1998) provides a large variety of cases where ‘random’ advantages early on in an industry have played out into permanent long-term dominance.

³⁵Reported at <http://searchenginewatch.com/showPage.html?page=3598011>.

4.1.3 Conclusion

To sum up, the monopoly (or near-monopoly) result seems reasonably robust to variations in the model structure given the underlying zero-user price/quality competition model of search. In addition, this result fits fairly well as first-order approximation as the current state of the search market in most jurisdictions (especially when dynamics are taken into account). However, the strong contestability result (and associated zero-profits outcome) is unlikely to be robust.

Thus, in examining the effect of monopoly we will focus on the case where the monopolist has some degree of flexibility in choosing variables such as the level of quality (by contrast, in the basic model above the monopolist is constrained to offer the maximum possible level of quality). Furthermore, in a dynamic model this flexibility would be likely to grow over time, concomitantly with the growth in the investment needed to rival the incumbent's quality level (it is these existing, 'sunk', costs which form the barrier to entry/competition in this market). Hence, in the next section, a fair degree of latitude will be assumed for the monopolist in regard of pricing and quality provision,³⁶ and our attention will be on how the monopolist's choice of these variables affects consumer and societal welfare, rather than on issues of market structure and market share.

5 Monopoly and Welfare

In this section we obtain formulae for social welfare and monopoly profits and, using these, values for the optimal and monopoly level of search engine quality. We then proceed to compare these outcomes in order to evaluate how well monopoly performs compared to the social optimum.

5.1 Welfare

We begin by defining a social welfare function W ³⁷:

$$\begin{aligned} W &= \text{Utility of Users} + \text{Profits of Search Engine} + \text{Profits of Advertisers} \\ &= U_U + \Pi_S + \Pi_A \end{aligned}$$

³⁶If one needed to incorporate the impact of external competition, either actual or potential, this could be imposed in the form of a minimum quality level or the like.

³⁷We accord consumer surplus and producer surplus equal weight in the social welfare function. While this is standard practice one could argue that the widespread and diverse set of users and the relatively concentrated ownership of most search engine companies might merit explicit distributional weights. We have not pursued this possibility but note that it would be relatively easy to introduce an explicit weighting into the analysis.

We observe that users' utility, search engine profits and advertisers' utility must all be inter-related: when advertisers pay money to the search engine they must expect to recoup these funds in the form of more buyers or higher prices.

Here we take a reduced form approach to connect advertising, search and users and thereby avoid a diversion into a detailed analysis of the form of the advertising market and the equilibrium conditions. First, recall that R_A is the total revenue from advertisers accruing to the search engine (which is therefore also equivalent to total payments by advertisers), and R_U the total *additional* revenue accruing to advertisers from users as a result of their advertising (that is revenue related to their advertising activities). Next let U_A be the (gain in) utility users derive as a result of advertising. Then total advertising profits (in respect of the activities under consideration here) are $\Pi_A = R_U - R_A$. Search profits are $R_A - (c + c_a)$. Meanwhile total user utility is given by the combination of the utility from search³⁸ $U_S(v, q_u) = \int_{q_u} U_t$ with the (net) utility from advertising $U_A - R_U$. With these formulations social welfare now has the form:

$$W = U_S + (U_A - R_U) + (R_A - c - c_a) + (R_U - R_A) = U_S + U_A - c - c_a$$

The final step is to specify U_A , the impact of advertising on users' utility. Here there are three options which could be put under the classic headings of advertising as:

'Good' $U_A > 0$. In this case, advertising directly improves users' welfare, perhaps by enabling better matches between consumers and producers, reducing 'search' time,³⁹ or simply directly increasing the valuation of the good advertised.

'Bad' $U_A < 0$. Advertising decreases consumer's utility, for example by reducing the quality of matches, or creating incentives for malicious behaviour.⁴⁰

³⁸As before, all superscript i indices used to index the search engine will be omitted as there exists only one search engine.

³⁹See for example, the arguments in Athey and Ellison (2007).

⁴⁰See Edelman (2006, 2007). As Edelman summarises: "Across all search terms we analyze, a Google ad is on average more than twice as likely to take a user to an unsafe site [one which installed spyware, adware and the like without fully informing the user] than is a Google organic link. At Ask, the difference is especially pronounced: Their sponsored results are almost four times as risky as their organic listings." Summed over all engines his data indicated that 'organic' results had 2.0% 'red-rated' sites and 1.1% 'yellow-rated' sites while for 'sponsored' results the rates were 6.5% and 2.0% respectively. Edelman goes on to give numerous examples of ways in which the sponsored results (advertised) on search engines may be substantially poorer than the organic results. To take one example: in May 2006 the top sponsored link for 'Skype' was download-it-free.com who, despite their name, charged \$29 to download a copy of Skype, a program that is supplied

‘Neutral’ $U_A = 0$. Advertising has a neutral effect on consumer’s utility generating neither direct benefits nor direct costs. This would correspond to the classic case of advertising as a war of attrition in which all (advertising) rents are dissipated in competition (or, in this case, payments to the search engine).

With plausible arguments on both the ‘good’ and ‘bad’ sides our approach will be to compromise and adopt the neutral perspective in which $U_A = 0$. While this is a convenient simplification we would point out that, obviously, a different assumption whether in the positive or negative direction could have a substantial impact on the overall welfare findings and this should be kept in mind by the reader. With this assumption, social welfare is (dropping ‘S’ subscript on U_S):

$$W = U(v, q_u(v)) - c - c_A = \text{User utility from search} - \text{Search engine cost}$$

Proposition 3. *The socially optimal level of advertising is zero and the socially optimal level of quality v^W solves:*

$$U'(v^W) = \underbrace{U_v}_{\text{Marg. util. on existing queries}} + \underbrace{U_q q'}_{\text{Marg. util. on new queries}} = \underbrace{c'}_{\text{Marginal cost}} \quad (1)$$

Proof. Advertising does not now enter the formulation for W except via c_A . It is therefore immediate that maximizing welfare requires $c_A = 0$ and $R_A = 0$.⁴¹ We therefore have $W = U_S(v, q(v)) - c(v)$ and the result follows. \square

5.2 Monopoly

Recall that monopoly profits are:

$$\begin{aligned} \Pi &= R_A(v, q) - c(v, q) - c_A(v, q) \\ &= R_A(v, q(v)) - c(v, q(v)) - c_A(v, q(v)) \\ &= R_A(v) - c(v) - c_A(v) \end{aligned}$$

for free by its producer (skype.com – the first ‘organic link for this search). He also discusses the possible incentives for search engines to behave in this way due to the large revenues that ‘bad’ sponsored links can generate (see e.g. <http://www.benedelman.org/news/012606-1.html>).

⁴¹This implicitly assumes the search engine could be directly funded by non-distortionary taxation. If this is not possible, advertising could be used if that were an efficient way to raise revenue. Nevertheless, the general point that ‘society’ would choose a lower level of advertising than the search engine is likely to be robust. We should also emphasize that a non-zero value of c_A has no material impact on the remainder of the welfare analysis presented below (i.e. replacing the zero value for c_A with the value for the monopolist will have no significant effects).

The monopolist's profit maximization problem is then to choose the quality level v^M that maximizes this function. We have that v^M satisfies the following first order condition:

$$R' = \underbrace{R_v}_{\text{Marg. rev. on existing queries}} + \underbrace{R_q q'}_{\text{Marg. rev. from new queries}} = \underbrace{c' + c'_A}_{\text{Marginal cost}} \quad (2)$$

Where subscripts indicate partial derivatives (the A subscript on R has been dropped), $'$ indicates a total derivative, and c'_A is shorthand for $c'_A(q(v)) = \bar{C}' - c'$ (which is necessarily positive). In text we have:

5.3 How Optimal is Monopoly?

We now compare socially optimal search quality v^W , and usage q_u^W , with that obtaining under monopoly. In particular, whether the quality level under monopoly is too high or too low compared to the socially optimal level (equivalently is search quality 'over-provided' or 'under-provided' under monopoly).⁴² The simple answer, as encapsulated in the following proposition, is that both under and over-provision are possible. However, there are several reasons, discussed in detailed below why under-provision is more likely.

Proposition 4. *A monopolist may under or over-provide quality (relative to the social optimum) depending on the form of the revenue and search utility functions. The likelihood that the monopolist under-supplies quality is greater:*

- *The smaller the advertising revenue from new users (R_q) compared to the social value of new users (U_q) (this is the classic social-private gap).*
- *The greater the positive effect of quality on the utility of existing users: U_v (this increases the socially optimal level but leaves the monopoly level unchanged).*
- *The greater the (negative) direct ('substitution' and 'antagonism') effect of quality on the monopolist's revenue R_v (this decreases the monopolist's chosen quality but leaves the socially optimal level unchanged).*
- *The higher the (marginal) cost of advertising c'_A (this decreases the monopolist's chosen quality but leaves the socially optimal level unchanged).*

⁴²The optimality of monopoly in traditional models of quality choice (and the related question of whether quality is under or over-supplied) is well-studied topic (the general result being 'it depends'). We note that our situation here is a little different in that user price is constrained to be zero and charging occurs indirectly through advertisers.

5.3.1 Under-provision

Take the approach of a ‘normal’ monopoly model: for an increase in quality a) the utility from an extra user ($U_q q'$) is larger than or equal to the revenue received by the monopolist ($R_q q'$) b) the effect on existing users would be greater for utility (U_v) than for revenue (R_v).

With an assumption of diminishing returns to quality, these functions are decreasing in v . Together with the fact that:

$$R' = R_v + R_q q' = c' + c'_A \geq c' = U_v + U_q q' \quad (3)$$

These would imply $v^W \geq v^M$, i.e. that the monopolist under-provides search quality (analogously to, but for slightly different reasons, to the way a monopolist under-supplies demand).

5.3.2 Over-provision

How can the situation differ from the ‘normal’ monopoly case? The key point is that the monopolist’s revenue function is not a ‘normal’ revenue function and the revenue from an additional user comes from the advertising revenue linked to that user not from a direct user payment. This breaks the link between utility and revenue and make it possible for marginal revenue R_q to exceed marginal utility M_q and hence (depending on the magnitudes of the direct effect of quality (R_v, U_v)) the fact that monopoly quality is higher than the socially optimal level ($v^M > v^W$).

There are two distinct mechanisms by which this can occur. First, and simplest, advertising revenue for an additional user is higher than the utility that the user gets from search (after all, a search engine chooses to charge advertisers and not users).

Second, revenue displays increasing returns in the number of users (at a greater rate than utility). Increasing returns can occur for two distinct reasons: a) economies of scale involved in advertising on a search engine, for example those that would arise from a fixed cost in generating or placing an advert b) economies of scope in advertising, arising, for example where an advertiser wishes to carry out several (related) campaigns each targeting different types of users and/or queries. In both cases, revenue would show increasing returns in the number of users (and quite independently of the utility function). This in turn means $R_q q'$ is, at least over some portion of its domain, increasing in quality rather than decreasing and hence that (total) marginal revenue R' may be larger than total marginal utility, U' , again implying that the monopolist’s quality v^M is greater than the socially optimal level v^W .

5.3.3 Likelihood of under-provision

The key comparison to make is between total marginal utility ($U' = U_v + U_q q'$), and total marginal revenue ($R' = R_v + R_q q'$) since at equilibrium we have: $U' = R' - c'_A$. With diminishing returns,⁴³ derivatives are decreasing and hence, for example, $U_v > R_v$ implies, other things being equal, that socially optimal quality is higher than monopoly quality. We now work through the likelihood claims of the proposition in turn.

It is immediate from the discussion in the previous paragraph that the smaller is marginal revenue from new users R_q compared to the utility from new users U_q (i.e. the larger ‘social-private gap’), the greater likelihood, all other things being equal, that the monopolist under-provides quality.

A similar argument applies to the direct effect of quality on utility and revenue: U_v, R_v . Lastly, since $U' = R' - c'_A$, the larger the marginal costs of advertising the lower U' relative to R' and hence, with diminishing returns, the higher the level of social quality relative to monopoly quality.

5.3.4 The direct effect of quality

It is clear that, by definition, the direct effect of search quality on utility, U_v , must be positive. However, the direct effect of quality on revenue, R_v , is less obvious and, we shall now argue, is in fact likely to be negative for two reasons: the ‘substitution’ and the ‘antagonism’ effects.

The ‘substitution’ effect arises from the fact that ‘ads’ can be seen as a method of helping consumers search. For example, if you search for ‘shoes’ or, even more explicitly, ‘buy shoes’, it may actually be useful for advertisements related to shoes and purchasing shoes, to be displayed. In this case, if a search engine is able to display advertisements (or ‘sponsored’ results) relevant to users’ search intentions, it is highly likely that the search engine is also able to display organic search results that are relevant. In this case, the advertisements and the search results are substitutes in the sense that better search means less need to click on advertisements (and vice versa). As such, improving search quality, by improving the search results the user receives for a given query, must necessarily reduce the likelihood of the user clicking on the advertisements (‘sponsored’ links) presented alongside. Conversely, worse search quality actually increases the likelihood, *for a given search*, that a user clicks on an ad rather than an ‘organic’ result.

⁴³Returns to quality will be eventually decreasing, even if they are increasing over some portion of the quality domain and second-order conditions will require the equilibrium point to be at a point of diminishing returns.

The effect also operates from the opposite, advertisers direction. If a search engine had such amazing quality that whenever one was looking to ‘buy shoes’ the ‘good’ places to buy shoes were presented as the top search results, there would be much less reason to advertise. However if the search engine does not present that information then it will be necessary for companies to advertise, and, once again, an increase in search quality reduces advertising revenue (and vice-versa).⁴⁴

The ‘antagonism’ effect, arises from the fact that, for a given query, search results may, by providing information that is ‘antagonistic’ to an advertiser, reduce the advertising revenue *for that query*. Consider the hypothetical example where a query for ‘vitamin supplements’ generates both ‘organic’ search results as well as advertisements to firms which supply such supplements and further suppose that there is new research out that demonstrates that such supplements are of no value (or even harmful). Displaying such a result high up (perhaps at the top of the search results) may increase quality for users but may well reduce the likelihood a given user clicks on an advertisement. As such by making this information prominent one reduces the amount of advertising revenue generated from that query.

Together these two effects imply that the direct impact of quality on a search engine’s revenue is negative: $R_v < 0$.⁴⁵

5.3.5 Conclusion

We established that it is possible a monopoly under or over-provides quality relative to the social optimum. However, under-provision appears much more likely for several reasons. First, the indirect effect of search engine quality on utility is likely to grow at least as fast, if not faster, than its effect on revenue (so $U_q > R_q$). Second, the direct effect of quality on utility is positive (and likely substantial) while the direct effect on revenue will be negative. Third, and least importantly, search engines have to bear advertising-related costs which increase their costs compared to the direct funding case and therefore reduces the quality provided.

The first of these effects is just the classic ‘social-private’ gap: the benefits of an extra unit of search quality to society are less than those extracted in the form of advertising revenues. The second of these effects, the ‘distortion effect’, arises

⁴⁴There are some suggestions that over time Google have downgraded search results which of are an explicitly commercial nature. Of course this could simply be to get rid of ‘spam’ or overly commercial information. However, it also forces those commercial organizations to buy advertising.

⁴⁵This does not imply increasing search quality is bad for a search engine: search quality also has an indirect impact via increasing the number of users/queries and it is likely that this effect is larger than the direct one $|R_q q'| > |R_v|$ and so the overall effect of increasing quality is positive on search engine revenue.

from the fact that quality potentially acts as substitute (or antagonist) for advertising (which is what the search engine is ultimately concerned with).⁴⁶

6 Regulation

Does Internet search require regulation – whether now or in the future? Search today is a huge business and the choices made by the primary companies involved, particularly in how to rank results and what advertises to display, affect the lives of everyone who uses the Internet. While some argue that search requires no regulation – and that any such regulation would unnecessarily impede the rapid technological progress of the industry; others have voiced concerns both about the informational integrity of search engines and the potential misuse of the vast power accumulating in commercial hands – a power to shape the information we discover and use.

We have already demonstrated both why the search engine market is so concentrated and that it is likely to become more so (converging to monopoly or near monopoly). It therefore seems unlikely that one can simply rely on ‘competition’ to avoid the need for regulatory engagement.

In considering a monopoly (or close-to-monopoly) situation the next step is to investigate whether, and how, a monopolist will behave in ways that are not socially optimal. This investigation is doubly important here. The structure of the search market, in particular the zero price faced by search engine users, often gives the misleading impression that a monopoly in the search engine market cannot result in negative consequences in same way as in other areas – areas where monopoly is explicitly associated with higher prices. This is not correct. Costs still exist here but they are indirect, operating either via the search engines charges to advertisers or via the quality of the service the search engine chooses to operate.

The model presented above permitted us to reduce welfare comparisons to a comparison of search engine quality, v and we established that it is likely a monopoly will under-provide quality with this under-provision attributable to two main factors: the ‘social-private’ gap and the ‘distortion’ effect.

What can a regulator do with regard to the first of these factors, the ‘social-private’ gap? In some ways the options are limited. After all they cannot mandate higher expenditures by private search engines and while government subsidies are a possibility they tend to bring with them a host of difficult issues: who should be awarded money; could such awards be anti-competitive if directed to a particular

⁴⁶This effect, unlike the first, is not a general one but will affect quality in specific areas where substitution and antagonism are prominent. Thus, it is likely more to ‘distort’ quality rather than unilaterally reduce it.

firm etc. If this route were to be pursued, the regulators best option would be to focus on funding basic R&D which would then be made available to all firms.

Another possibility, along similar lines, but which avoids some of the difficulties, would be the provision of a computing grid and search index upon which developers could try out different algorithms. This option points towards the fact that the provision of a search engine divides (imperfectly) into what we could term the ‘software’ provision and the ‘service’ provision. The ‘software’ includes all the main software used to run the system, including the ranking algorithm. The ‘service’ side involves all the infrastructure, data-centres, support systems etc, which run the software and actually respond to users’ queries. Obviously there is some degree of interaction between these two sides – for example developing the software requires feedback and data from actual usage, but it is also possible that the two sides could be separated to some degree.

This is important because the costs involved in algorithm development could be much smaller than the large fixed costs of infrastructure – though in the long run it may be the algorithm, extensively developed via learning-by-doing etc, that provides the real barrier to entry. Thus, decoupling the two, might allow for greater competition, innovation, and perhaps most importantly, transparency on the ‘software’ side while on the ‘service’ side there remains a monopoly or near monopoly (provided by the Government or a neutral, regulated, third-party). This would be similar to a situation in many other industries where there exists a key piece of infrastructure which, for technology and cost reasons is a natural monopoly. For example, in electricity supply, the underlying transmission network is a natural monopoly (and hence regulated) but competition is clearly possible in generation (and so less regulated). Similarly, in telecommunications it is usual for the ‘local loop’ to be a natural monopoly (and hence regulated) but for there to be competition in service provision (telephony, broadband etc) over that ‘local loop’.⁴⁷

Such an approach in which there is a division, at least from a regulatory point of view, between ‘software’ and ‘service’ would have more general benefits than allowing targeted support. First, competition in ‘software’ would increase spending and therefore quality moving society towards the socially-optimum. Second, and relatedly, it would reduce the risk of long-term lock-in to a single provider. Third, regulatory attention could be focused on the ‘service’ side which in many ways is simpler: economies of scale arise less from (field-specific) innovation and more from the sunk costs of infrastructure.

⁴⁷It is important to emphasize that these sorts of divisions are not a magic bullet. Experience in both electricity and telecommunications have shown failures as well as successes with this model of separation.

Turning now to the second factor mentioned, the ‘distortion’ effect, we observe that the ‘software/service’ division would also be beneficial by increasing transparency and competition on the ‘software’ side. However, there are other ways of dealing with this problem without taking such a major step. ‘Distortion’ could be handled, for example, by greater monitoring of search results and their relation to advertisements. Relatedly, the regulator could request confidential access to the search engine’s ranking algorithm and could also act as a review panel for those who wish to ‘appeal’ their ranking.⁴⁸ Similarly, such a regulator might also monitor the other, advertising side of search engine activities, not only in the area of advertising content but also in relation to issues such as click-fraud.

To sum up, there are both the grounds and the means for greater regulatory oversight of search engines’ activities – be such oversight formal or informal. There are a variety of ways such regulatory intervention could proceed. The most major, but also perhaps the most effective, would involve dividing search engine provision, whether conceptually or actually, into two separate ‘software’ and ‘service’ components. Less dramatically, it seems clear that, as the power of search engines grows, there will be an increasing need for independent monitoring of the quality and content of search engine results together with a body able to deal with complaints regarding search engine rankings.

7 Conclusion

This paper has provided a comprehensive introduction and analysis of the search engine market. After a basic overview of the nature of search engines, their current importance, both commercially and socially, we turned to the main empirical and theoretical questions that animate our investigation: the current and future structure of the search engine market and its implications for societal welfare.

Our empirical material demonstrated how the concentration of the search engine market has grown over time and has now reached very substantial levels though with some significant and important variation across market segments. This also formed the background for the theoretical investigations that followed and which form the core of this paper.

This theoretical work provides what is, to our knowledge, the first formal analysis of the wider search engine market and its welfare implications. The first step involved developing a basic model which captured the main features of the

⁴⁸At present all major search engines, while providing facilities with which to raise complaints, claim complete discretion in resolving any disputes over ranking. This is unlikely to prove sustainable into a future where search is increasingly important, powerful, and concentrated.

search market, in particular the ‘implied revenue’ function which gives search engine revenue as a function of users. The value of a user here is not, as in a normal case, the revenue from a direct charge to that user but is the implied value arising from the advertising revenue that user generates.

We demonstrated how the structure of the search engine market – that users care about search quality but *are not* charged, while advertisers care about users and *are* charged – explains the highly concentrated nature of the search engine market and makes it probable that the market will continue to evolve down this path towards monopoly. We compared the welfare performance of a monopoly, measured by the quality of search provided, as compared to the benchmark of the socially optimal provision and showed that a monopolist was likely to provide an inefficiently low search quality (and engage in ‘distortion’ of its organic results). Given this, some form of oversight, possibly including formal regulation, will become increasingly necessary. One potential remedy is the division of the search engine market into ‘software’ and ‘service’ segment, analogous to divisions in other regulated industries such as electricity and telecommunications.

In conclusion, the structure of the search market, in particular its economies of scale and tendency to winner-takes-all competition, are likely to undermine the potential for vigorous market competition, especially in the long run. When monopoly, or near monopoly, does obtain the private interests of a search engine and the interests of society are likely to diverge. The power rapidly accumulating in the hands of a few major search providers is a great one and it is essential that policy-makers take steps to ensure it is used in ways that are for the benefit, not the detriment, of society as a whole.

References

- Armstrong, M. (2006): “Competition in Two-Sided Markets,” *RAND Journal of Economics*, 37, 668–691.
- Athey, S. and G. Ellison (2007): “Position Auctions with Consumer Search,” .
- Baye, M. R. and J. Morgan (2001): “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,” *American Economic Review*, 91, 454–474.
- Chen, Y. and C. He (2006): “Paid Placement: Advertising and Search on the Internet,” Working Papers 06-02, NET Institute.
- Dasgupta, P. and J. Stiglitz (1980): “Uncertainty, Industrial Structure, and the Speed of R&D,” *The Bell Journal of Economics*, 11, 1–28.
- Edelman, B. (2006): “Adverse Selection in Online ”Trust” Certifications,” .

- Edelman, B. (2007): "Assessing and Improving the Safety of Internet Search Engines," in *The Rising Power of Search Engines on the Internet*, 259–277.
- 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.
- Ellison, G. and S. F. Ellison (2004): "Search, Obfuscation, and Price Elasticities on the Internet," NBER Working Papers 10570, National Bureau of Economic Research, Inc.
- Evans, D. S. (2008): "The Economics of Online Advertising Industry," *Review of Network Economics*, 7, 1146–1154.
- Gandal, N. (2001): "The Dynamics of Competition in the Internet Search Engine Market," *International Journal of Industrial Organization*, 19, 1103–1117.
- Ghose, A. and S. Yang (2007): "An Empirical Analysis of Search Engine Advertising: Sponsored Search and Cross-Selling in Electronic Markets," Working Papers 07-35, NET Institute.
- Goldfarb, A. and C. Tucker (2007): "Search Engine Advertising: Pricing Ads to Context," Working Papers 07-23, NET Institute.
- Harris, C. and J. Vickers (1985): "Perfect Equilibrium in a Model of a Race," *Review of Economic Studies*, 52, 193–209.
- Jansen, B., M. Zhang, and Z. Ying (2007): "The Effect of Brand Awareness on the Evaluation of Search Engine Results," CHI 2007, April 28-May 3, 2007, San Jose, California, USA.
- Nocke, V., M. Peitz, and K. Stahl (2007): "Platform Ownership," *Journal of the European Economic Association*, 5, 1130–1160.
- Rochet, J.-C. and J. Tirole (2003): "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1, 990–1029.
- Shaked, A. and J. Sutton (1983): "Natural Oligopolies," *Econometrica*, 51, 1469–1483.
- Sutton, J. (1991): *Sunk Costs and Market Structure: Price Competition, Advertising, and the*, MIT Press.
- Sutton, J. (1998): *Technology and Market Structure: Theory and History*, The MIT Press.
- Telang, R., T. Mukhopadhyay, and U. Rajan (2004): "Market Structure for Internet Search Engines," *Journal of Management Information Systems*, 21, 137–160.
- Varian, H. R. (2007): "Position Auctions," *International Journal of Industrial Organization*, 25, 1163–1178.
- White, A. (2008): "Search Engines: Left Side Quality versus Right Side Profits," .