Is Google the next Microsoft?
Competition, Welfare and Regulation in Internet Search

Rufus Pollock

May 2009

CWPE 0921

IS GOOGLE THE NEXT MICROSOFT? COMPETITION, WELFARE AND REGULATION IN INTERNET SEARCH

RUFUS POLLOCK

UNIVERSITY OF CAMBRIDGE

APRIL, 2009

ABSTRACT. The rapid growth of online search and its centrality to the ecology of the Internet raise many questions for economists to answer: 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 (c) discuss the regulatory interventions that policy-makers could use to address these.

Keywords: Search Engine, Regulation, Competition, Antitrust, Technology

JEL Classification: L40 L10 L50

Emmanuel College and Faculty of Economics, University of Cambridge. Email: rp240@cam.ac.uk or rufus@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. I would also like to thank <http://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 Licence (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 the World-Wide-Web.¹ Beginning from practically nothing twelve 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.

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’.

It is also now 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 quarter two 2007, 35% of Google's total revenue (\$1.35 billion) 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 its 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. We begin in section 2 with a brief history of Internet search which provides the basic background for what follows. Those already familiar with this material, or who wish to get straight to the more 'economic' results, may safely skip this section. In section 3 we provide empirical evidence on the levels of concentration in the search engine market, both over time and across jurisdictions.

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 3 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 data clearly shows that the search engine market is indeed highly concentrated and has grown more so over time. Sections 4, 5 and 6 form the core of the paper. In section 4 we introduce a basic model of the search engine market and use it in section 5 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.

This motivates the second piece of theoretical analysis in section 6. 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. Among other things, 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

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 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.¹²

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).

⁹See Baye and Morgan (2001) for an early paper analyzing a 'yellow-pages' (information-gatekeeper) type model in an online environment. This type of work also connects directly into the much larger traditional literature on 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.

¹²There has been a very limited work on areas more directly related to ours, particularly on issues of market share. Gandal (2001), did (very) early empirical work which examined changes in market share in the late 1990s. Similarly, Telang, Mukhopadhyay, and Rajan (2004) also looked at market share, though from a theoretical perspective and try to explain the persistence of low-quality firms in a market where prices are zero.

2. A BRIEF HISTORY OF WEB SEARCH

The history of web search is inextricably bound up with the development of the world wide web. We therefore begin with a brief sketch that outlines the nature and history of the Web before turning to the question of search.

2.1. The World-Wide-Web. The World Wide Web is a hypertext system that has been adopted as the main method of information dissemination on the Internet. The element central to the Web - and any universal information system - was the creation of the Universal or Uniform Resource Identifier (URI), a method of uniquely assigning a name or address to any document or resource (for example a database) anywhere in the world. To most people this takes on the form of a URL, a Uniform Resource Locator, familiar as the ubiquitous `www.somename.com`. This in turn allowed at last a concrete implementation of hypertext, a method of inserting active links to other documents first conceived of by Vannevar Bush in the 1940s and elaborated by Ted Nelson in the form of his Xanadu project, and the feature which truly makes the Web a ‘web’. In many ways the great achievement of the Web has not been technical, but social: persuading a large number of different groups each with their own standards and interests to agree to the formation of this one universal system.¹³

The Web, though built upon much previous work, was at its initial stage largely the creation of one man: Tim Berners-Lee. Berners-Lee was born in London and educated at the Emanuel School and Queen’s College, Oxford, from which he graduated with a first in Physics in 1976. Both his parents were mathematicians and had worked on the team that programmed the world’s first commercial stored-program computer, the Manchester University ‘Mark 1’. While at Oxford he built his own computer using a soldering iron, an M6800 processor and an old television, and on leaving he became a software engineer, first with Plessey Communications, and then with D.G. Nash. In 1980 he went to CERN as a consultant for a six month period. While there in his spare time he wrote a program called Enquire in which he first used hypertext links to allow navigation.

¹³It is important to remember that when the Web first arrived it was only one among several competing alternatives and by no means the preeminent option. In particular, in the early 1990s, when the Web was launched both Gopher and WAIS performed similar functions and a work such as Ed Krol’s *The Whole Internet User’s Guide & Catalog* (published in 1992) clearly put these, more established protocols, above the newly arrived ‘World Wide Web’.

In 1984 he returned to CERN to work on data acquisition and control. As computers and the internet evolved during the 1980s Berners-Lee became ever more interested in developing a system that would allow information to be both created and shared in a universal format and would also support a hypertext system - a crucial aspect he believed of building a true 'web of knowledge'. In March 1989 he wrote his first proposal for a global hypertext system but it was not until May 1990 that he was authorized to pursue the project and settled upon the name World Wide Web (after considering others such as Information Mesh or The Information Mine). Writing initially for the NeXT system, Berners-Lee quickly produced a Web client or browser that would allow a user to create, edit or browse hypertext pages and which he named simply WorldWideWeb. He also produced a Web server which would store the pages and serve them up to the user as they were requested. In doing this he settled upon the basic standards which have continued to underlie the system to the present day, namely HTTP (HyperText Transfer Protocol), HTML (HyperText Markup Language) and URI (Universal Resource Identifier). By Christmas Day 1990 Berners-Lee and his colleague Cailliau had set up the first website named info.cern.ch and had transferred the first web pages over the internet.

Despite these advances, at this point there were no signs of the Web's huge future success but only difficulties. Paramount among these, was the problem of persuading users at CERN, each with their own computer system and way of doing things, to adopt the new approach. As Berners-Lee later wrote, "there was a constant background of of people promoting ideas for new software systems. CERN obviously couldn't tolerate everybody creating unique software for every function. Robert and I had to distinguish our idea as novel, and one that would allow CERN to leap forward. Rather than parade in with our new system for cosmic sharing of information, we decided to try to persuade people that we were offering them a way to extend their existing documentation system. This was a concrete and promising notion. We could later get them to sign on to the dream of global hypertext."¹⁴

Finding it difficult to persuade CERN of the importance of the new system Berners-Lee decided to release it outside CERN, and in August 1991 he released it on the Internet, posting notices in several Internet forums including alt.hypertext. Web sites began to

¹⁴Berners-Lee (1999), p. 47.

appear all over the world, and initially Berners-Lee would add a link to each one onto info.cern.ch. By measuring the number of ‘hits’ or page views of info.cern.ch Berners-Lee could monitor the early progress of the web. In July and August 1991 there were between 10 and 100 ‘hits’ a day. As Berners-Lee later wrote: “This was slow progress, but encouraging. I’ve compared the effort to launch the Web to that required to launch a bobsleigh: everyone has to push hard for a seemingly long time, but sooner or later the sleigh is off on its own momentum and everyone jumps in.”¹⁵

At this point progress began to take place more and more rapidly. First and foremost, browsers (software that would allow the use to access and view Web pages) were developed for many different platforms: Unix (Erwise 1991, ViolaWWW 1992 and Mosaic 1993), Apple (Samba 1992-3, Mosaic 1993) and later PC (Cello and Mosaic 1993). Since then most of these early ‘open’ browsers have been superseded by ‘free’ commercial products such as Navigator (Netscape), Internet Explorer (Microsoft), and Firefox (Netscape/Mozilla Foundation) though it is noteworthy that the creators of Navigator had also made Mosaic. Second, in a step that was to prove crucial to the long term direction and development of the Web, in April 1993 Berners-Lee persuaded CERN, who as his employers owned the intellectual property rights to his work, to put everything relating to the Web in the public domain. ‘Hits’ on info.cern.ch grew exponentially from the beginning. By summer 1992 the number had reached one thousand and by summer 1993 ten thousand: “I no longer had to push the bobsleigh. It was time to jump in and steer” said Berners-Lee.¹⁶

2.2. Web Search Engines. As the Web’s exponential growth commenced in the early 1990s users began to confront the issue of finding what they wanted in a rapidly expanding sea of information. At the very beginning it was feasible for Berners-Lee simply to add a link to a new ‘web-site’ to info.cern.ch but as the web grew such an approach rapidly became impractical – between 1993 and 1996 the web went from 130 sites to 600,000. Some way had to be found to *crawl*, *index* and *search* the web in a way that could cope with this exponential growth in material.¹⁷

¹⁵Berners-Lee (1999), p.54

¹⁶Berners-Lee (1999), p.81

¹⁷This chapter focuses on the web and therefore has excluded some earlier search engines such as ‘Archie’ and ‘Veronica’. ‘Archie’ was created by Alan Emtage, a McGill University student, in 1990 and indexed Internet ftp archives. ‘Veronica’ was created in 1993 by University of Nevada students and did the same thing as ‘Archie’ but for ‘Gopher’ archives.

Table 1 below, details the history of the main early Internet search engines. As it shows, much of the very early work, focused simply on ‘crawling’, compiling a list (the index) of the resources available. However, soon the quantity of material led to attention to the problem of presenting the material via querying and ranking of the index.

Date	System	Creator	Details
1993	WWW Wanderer	Matthew Gay (MIT)	Crawler and Indexer. Primary focus on charting growth but did have a basic search interface. Caused some problems because its roboting (initially) used a lot of bandwidth.
1993	ALIWEB	Martijn Koster	Created in response to the WWW Wanderer and provided Archie-like indexing for the web. Unlike the Wanderer, ALIWEB did not crawl but depended on user submission of sites.
1993	JumpStation / WWW / RBSE	-	All of these systems were active by December 1993 but only provided crawling and indexing with little attempt to sort the results for searchers. [2]
April 1994	WebCrawler	Brian Pinkerton (University of Washington / NextStep)	First system to index the full text of a web page. Used number of a links a web page had to determine its importance. Acquired in June 1995 by AOL for \$1M. [1] 41-42.
May 1994	Lycos	Michael Maudlin (CMU / DARPA)	Main advance was to add more sophisticated text analysis (including use of inbound link text) to the basic crawler in order to improve search results. Maudlin and Davis (CEO) bought out 80% for \$2M in June 1995. Backed by CMGI's @Ventures Lycos became a portal and was bought in May 2000 by Terra (Spanish Telecoms giant) for \$12.5 billion. It was later sold to a south Korean company for \$100M.
April 1994	Yahoo!	Yang and Filo (Stanford)	Yahoo! was not a search system but a directory. Starting out as a collection of the founders' favourite pages as the web grew it reorganized and provided search facilities but remained a hierarchical catalogue to which entries were hand-added. In late 1995 Yahoo! added ‘proper’ search to its directory via a partnership with Open Text. Later it would partner with other providers including Altavista and Google until acquiring its own search technology via purchase of Overture in 2003.
Aut. 1995 (1993)	Excite / Architext	6 Alumni of Stanford	Excite developed out of Architext which was created by 6 Stanford students in 1993. On its launch in 1995, it provided a web directory and full-text search engine using statistical analysis of word relationships. It quickly became a fully-fledged portal and acquired other search engine providers: Magellan (for around \$18M) and WebCrawler from AOL (\$4.3M). Excite itself was acquired by @Home (a broadband company) in January 1999 for \$6.5 billion (renamed to Excite@Home). It eventually ended up in Chapter 11 proceedings and was sold to Interactive Search Holdings (ISH) a small holding firm in 2002 for \$10 million. In March 2004 ISH was in turn acquired by Ask Jeeves.

December 15 1995 (1994)	Altavista	DEC (Louis Monier)	The first system to offer ‘full’ search and reasonable quality. Not only did it have the largest index of material (16M documents) on its release but it provided (virtually) unlimited bandwidth and natural language querying. Altavista was immediately successful racking up 300,000 visits on its first day and serving 4 billion queries in its first year. For a variety of reasons, including poor management, Altavista despite its strong start declined rapidly in the late 1990s and was eventually acquired by Overture for \$80M stock and \$60M cash on February 18, 2003 (Overture were in turn acquired by Yahoo! later that year). Nevertheless Altavista was one of the first major milestones in search development, providing, in the words of Gary Flake, “a significant improvement over the state of the art.” (quoted [1] p. 39).
May 20 1996	Hotbot / Inktomi	Berkeley	Pioneered the paid inclusion model in which sites would pay for inclusion in search results but this was never as effective as the pay per click model developed by Overture. Owned by Wired Digital, it was sold to Altavista and eventually in December 2003 to Yahoo! for approximately \$235 million. [2]
April 1997	Ask Jeeves (now Ask.com)	Ask Jeeves	Launched as a natural language search engine, Ask originally used human editors to match queries. For a while they were powered by DirectHit but in 2001 Ask Jeeves acquired Teoma to replace DirectHit. On March 21 2005, Ask Jeeves was acquired by Barry Diller’s IAC for \$1.85 billion with Ask Jeeves renamed to Ask in 2006.
1998 (1996)	Google (BackRub)	Page and Brin (Google)	Primary innovation was the development of the PageRank approach to site (and page) ranking when ordering search results. This method used not only the links to a site but the reputation (ranking) of those linking sites in determining the reputation of that site in a recursive process based on estimating an eigenvector of the adjacency matrix of the Web (an approach derived directly from the bibliometrics literature based on the analogy of links between websites and citations between papers). Google Inc was officially formed on September 7 1998 but the service had been up and publicly available for a substantial time previous to that (and had originally grown out of Page and Brin’s work on a research project named ‘BackRub’). [1] 73ff. By late 1998 when the company was formed, google.com was receiving around 10,000 queries a day (which though substantial was orders of magnitude less than most of the major existing search engines). [1] p.85.

Table 1: Development of the Major Search Engine Systems 1993-2000. Sources: [1]

Battelle (2005)[40 ff.] and [2] <http://www.searchenginehistory.com/#early-engines>

2.3. Making Money, Or, How to Sell Attention via Advertising or Paid Inclusion. It is clear that the developers of the early search engines had little idea how they would ever make money from what they were doing (in fact many of the early innovations

were developed in academia or in company labs where this question had secondary importance). By the mid-to-late nineties most observers, and most companies themselves, had moved towards the search-engine-as-portal model where the search-engine was seen as a simple way to generate traffic (and therefore ‘eyeballs’) which could then be converted into advertising revenue in the same way as any other attention.

There were two main respects in which this analysis proved to be wrong in the long run. First, and less importantly, viewing search engines as similar to any other ‘portal’ (or part of a portal) significantly underestimated the centrality and importance of search engines in the future Internet environment. Search engines are different from other sites because of their crucial ‘gateway’ role – a role whose importance has grown and grown as the exponential increase in information online has continued. This role guaranteed not only that traffic to search engines would continue to increase in line with (or even above) the general rate of usage of the web as a whole, but also that they would become an essential first-point-of-call for anyone venturing onto the Internet.¹⁸

Second, and more importantly, was the realization that taking search engine users as equivalent to the users of any other website underestimated significantly their value from an advertising perspective. Specifically, the user of search engine has provided an additional, and crucial, piece of information about themselves (or rather their intentions): their search query.¹⁹ This query immediately gives the operator of a search engine information as to what that user is looking for – for example if the user has queried for “shoes” we can be fairly certain the user is interested in shoes (and may even be interested in buying some). As a result, a search-engine can dramatically increase the relevance of the advertisements it displays – in a very similar way to the manner in which it uses the user’s query to select the ‘normal’ search results – and increased relevance of course means increased value to those who wish to advertise.

¹⁸Thus, for example, in 2008 a search engine (Google) would be nominated the top brand in the UK – and this despite being the only company in the top 50 to have been founded since 1990. See <http://news.bbc.co.uk/1/hi/business/7260107.stm>.

¹⁹The same is also true, though to a somewhat weaker extent, of a visitor to any website. After all, if the subject of the website is, say, football then the probability that a given visitor is interested in football-related information (including advertisements) is significantly higher than for a user drawn at random from the population. It is this feature that lies behind the massive rise in the online ad-broking business, that is where advertisements (especially text-advertisements) are provided to website operators by the very same companies that operate web search (Google, Yahoo! etc). We shall return to this issue in further detail below.

This idea in itself is fairly old. For example, its underlies all ‘Yellow Pages’, and almost all advertising will take some account of audience segmentation (after all advertisements in ‘Autocar’ are likely to reach a different set of people from those in ‘Vogue’).²⁰ Nevertheless, the realization of the particular value offered by search queries, and its general introduction to the online environment, is generally credited to Bill Gross and the company he started to exploit it: GoTo.com.²¹ Furthermore, while the idea may appear obvious in retrospect it was still slow to catch on in the late 1990s. For example, as Google’s founder’s started to worry about income generation in late 1999 and early 2000 their main thoughts were about ‘banner ads’ and the like – part of the old approach of simple ‘cost-per-impression’ (CPM) advertising in which the ad was (relatively) ‘untargeted’. However such thinking changed rapidly in the early 2000s as the market matured. In particular, in February 2002 Google launched an Adwords service that allowed firms to bid for higher ranking in the advertising listings found on their search results pages.²² These ‘keyword auctions’, in which advertisers bid for positioning on individual sets of keywords, are now used by almost all major search engines and form an area of active research both within these firms and the wider academic community.²³

²⁰And, of course, the use of general demographic information to target product information has not only continued in the digital, online world but grown dramatically – largely thanks to the increased ability to record and process information about users. It is this pool of highly specific information about users (including complete information about their friendship/acquaintance network) that makes social sites such as Facebook so attractive to advertisers and so (potentially) valuable to their operators.

²¹GoTo.com in turn became Overture in October 2001, which after acquiring Altavista early 2003, was itself finally bought by Yahoo! later that year.

²²Note that Google had launched simple ‘text-banner’ adverts at the top of their search results back at the end of 1999 – and were one of the last companies to do so (see Search Engine Watch’s Search Engine Report from Dec. 6 1999). However it was with the launch of the Adwords service that the possibilities of paid inclusion were first fully appreciated – and exploited. Regarding this service it is noteworthy that Overture (the successor of GoTo.com) alleged that it infringed on one of their patents (US patent 6269361) and began proceedings against Google in April 2002 (<http://www.news.com/2100-1023-876861.html>). This dispute was finally settled in 2004 (after Overture’s acquisition by Yahoo!) with Google issuing 2.7 million shares of common stock to Yahoo! (http://www.news.com/Google,-Yahoo!-bury-the-legal-hatchet/2100-1024_3-5302421.html).

²³Google, along with many other participants, use a form of generalized second-value auction. Apparently this was initially adopted simply for performance and usability reasons (a first-value auction would result in users continually ‘logging-in’ to check their position and shave their bid). However, research since then, see e.g. Varian (2007); Edelman, Ostrovsky, and Schwarz (2007), has demonstrated several attractive properties from a purely auction-theoretical perspective. Given the combination of active use (literally ‘billions of dollars’ worth of keywords auctioned) and the theoretical challenges in analysing what are, in effect, large dynamic auctions of multiple goods (consisting of both substitutes and complements) it is likely that this area will remain a fertile area of investigation for auction economists and others in the years to come.

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 2. 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.

3. 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 3 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 over 90% in all jurisdictions except Hong Kong.²⁴ 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.²⁵

Also interesting is the question of how market shares have evolved over time. Obtaining good (comparable) market share data over a reasonable stretch of time is not easy. 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

²⁴It 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>).

²⁵Perhaps 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.

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 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 2 datasets apparent in the 2004 values).²⁶ The basic story shown by the graph is a simple one: 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 market share over 50% to be indicative of a dominant position. There are two distinct issues in assessing whether there is cause for concern: first, the search market might still be competitive even in situations where one company (or a few companies together) has/have a very large market share. Second, 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, this might not be detrimental to overall social welfare – in fact the existence of a monopoly might even be welfare improving.²⁷ Clearly, neither of these questions can be adequately addressed without developing a more detailed analysis. And so it is to this task that we now turn.

²⁶This 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.

²⁷We shall discuss this point in more detail below so here we confine ourselves to pointing out that the 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.

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

FIGURE 2. The structure of the search engine business.

4. MODELLING THE SEARCH ENGINE MARKET

4.1. Introduction. The search engine market has certain distinctive features related to structure, costs and pricing, which must be central to any modelling exercise. We discuss each of these in turn.

The structure of the search engine market is displayed schematically in Figure 2. As can be seen it has a basic ‘3-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.

The first significant fact about pricing is that the primary ‘content’ input for search engines – the underlying information on the web – is provided for ‘free’. That is, because of the history and tradition of the Web (and the Internet), search engines have generally been permitted access to this content at no charge – after all most information posted publicly on the web is already free for anyone to look at, and in addition, search engines can help increase traffic to a website.²⁸

The next major fact, and equally important, is that search engines do not directly charge users for their service but supply it for ‘free’.²⁹ In our model below we shall take this as an assumption and we therefore think it worthwhile to discuss the likely reasons for this here, especially as unlike ‘content’, this outcome must be the result of conscious choice by search engines.

First, the use-value of a search engine (the value of a query) is likely to be very heterogeneous (both across users and time) and hence may be difficult to price ‘well’. Second, and more importantly search engines are essentially (meta-)information providers supplying users with information about where other information is located. Hence, charging for their service (i.e. charging users for queries) would suffer from all the classic Arrovian difficulties, most prominently that the value of a given query is often highly uncertain

²⁸Though like all other generalisations this is not completely true. First some websites have wished to restrict access to search engines, either because of concerns about caching and reuse or out of a desire to be remunerated by search engines (see e.g. Copiepress v. Google <http://www.groklaw.net/articlebasic.php?story=20070726152837334>).

It is also the case that search engines can impose very significant load burdens on websites when they ‘crawl’ them. This was particularly so in the early days of the web but even today search engines crawlers can easily account for a very substantial portion of total traffic – one of the authors has personal experience of this and has actually restricted search engine access to parts of a site he helps maintain precisely because the performance hit caused by search engine ‘crawls’.

²⁹Some search engines do in fact sell their search facility for use on corporate intranets etc but this provides a small percentage of their revenue.

before it is performed.³⁰ Third, and related to the previous two points, is that charging users would necessitate significant transaction costs on two main counts. First, in relation to administration of charges (processing and payment). Second in maintaining an effective exclusion regime which prevented those who had not paid for use for gaining access, directly or indirectly, to the search engine's search results. Fourthly, and finally, search engines have an alternative method to direct charging for generating revenue from users: selling users' 'attention' (and intentions), generated via the use of the search facility, to advertisers.

It is worth emphasizing that this last point is central to the explanation of zero user-charges for without this (extremely effective) alternative method for raising revenue there would have been no option but to charge users directly – whatever the drawbacks of this approach. It also rounds out the details of the charging structure by making clear that advertisers are the one group out of the three whose interaction with the search engine has a financial component.

In discussing why search engines do not charge users we have not explicitly excluded the possibility that the search engines could *pay* for users to use their search engine either directly or indirectly. For example, a search engine could pay to ensure they were the default search option in a web browser,³¹ or a search engine could even pay users directly for the searches they perform.³² However, the scope for paying just as for charging seems fairly limited, at least at the present. The reasons are similar to the question of charging – transaction costs related to monitoring and payment, uncertainty as to the (advertising) value of a user etc. Thus, rather than saying that search engines cannot charge it is perhaps more accurate to say that search engines cannot set prices (whether positive or negative) but rather are constrained to price at zero.

³⁰There is also the additional problem here that the value of a given query may well depend on the ability to perform other ones in complicated ways. Formally, some queries will be complements and other substitutes and this may vary across users in ways which are difficult to predict.

³¹This approach has, in fact, already been adopted with Google sharing ad-revenue with the Mozilla Foundation in exchange for Google being the default search engine in the Firefox browser (it is also reported that Yahoo! have entered into a similar deal with a competing Gecko-based browser named Flock).

³²Users might start auctioning their 'attention' to the highest bidder in the same way that search engines auction it to advertisers. Perhaps, more plausibly given how diffuse users are, one might imagine that intermediaries would enter (perhaps ISPs could take this role) who would locate themselves between users and the search engine and would charge search engines a fee for directing their user base to one search engine or another.

The last significant feature of the search engine market to mention relates to technology and costs. In particular, search engines are R&D intensive and the market generally displays high levels of innovation and obsolescence.³³ 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.

4.1.1. *Relevant Literature.* One might think that, given the three-sided nature of search, the obvious analytical tools to use would be those developed in the literature on two-sided, platform, markets (see e.g. Rochet and Tirole (2003); Armstrong (2006); Nocke, Peitz, and Stahl (2007)). However, at least under the present charging structure, the search engine business does not fit comfortably within this paradigm. In particular, 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 pay to participate.³⁴ 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.

³³For example, Microsoft claimed to be spending over \$1bn a year on its online services (including its search engine) in 2006 (http://www.cbonline.com/article_news.asp?guid=3D810B1B-BBE0-482D-A81C-DBE60BAB97C4).

³⁴Note here that ‘advertisers’ advertise on the search engine not on any content provider. Most search engine companies are also active in the ‘ad-brokerage’ market for reasons of economies of scope – selling advertising ‘space’ on search results also provides you with the tools (and customer base) to sell advertising ‘space’ on general sites. Furthermore, ‘ad-brokerage’ does fit well within the two-sided model since here the two sides (‘content providers’ and ‘advertisers’) do care about the size of the other group and the ‘ad-broker’ naturally takes a platform role. However, here we are going to focus exclusively on search engine provision and will ignore related (and significant) business activities, such as those related to large-scale ‘Ad-Brokerage’.

Instead it will be more useful to utilize the standard toolkit on 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. Having established this, in following sections we discuss the implications of monopoly for social welfare and regulation.

4.2. 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 basic assumptions which help simplify the analysis and allow us to concentrate on the key areas of concern:

- (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.
- (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,³⁵ but it also fits both 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’.

Formally, 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

³⁵There is some degree of evidence that users do use multiple search engines. For example Search Engine Watch 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.

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:³⁶

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

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 . Note also that the 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). Let q_u^i be the total user demand for search engine i . The user's outside option will be normalized to 0 and users use the search engine which delivers the highest utility. Thus, q_u^i equals the set of t such that $U_t^i \geq 0$ and $U_t^i > U_t^j$ for all other search engines j . Thus 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 . 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: $q_u^i = q_u^i(v^i, v^{-i}, p_u^i, p_u^{-i})$.

At this point we make a major assumption which reflects the current pricing structure of the search engine market and will help greatly simplify our analysis:

Assumption 1. Search engines do not charge users: $p_u^i = 0$.

Thus, utility becomes $U_t^i = u(t, v^i)$ and user demand for search engine i becomes $q_u^i(v^i, v^{-i})$. It will also be useful, for notational convenience, to drop the search engine index i except where it is absolutely necessary for clarity. Thus, $U_t = u(t, v)$ etc.

Having established the basic user model we now turn to 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

³⁶A 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 $\theta' > 0$).

of this revenue across all users of that search engine:

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

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’). We now make our second assumption that reflects our discussion in the introduction:

Assumption 2. ‘Core/user’ costs are primarily fixed. In particular the marginal cost of an additional user is approximately zero. Furthermore, the cost of supplying a given quality is (up to a point) independent of the number of users.³⁷

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

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

Before proceeding to the results of the next section it is worth making some observations. First, interpret q_u as a scalar which (taking other search engine qualities as constant) is a function of v .³⁸ 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:

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

Note that this now looks like a classic vertical product differentiation problem in which $\bar{p}$ now represents the price charged to a user (here it is the derived price of a user in terms of advertising revenue). However there are some major differences, in particular $\bar{p}(q_u)q_u$

³⁷Recall that quality has several components. Pure search results quality is essentially a nonrival good and therefore has absolutely zero marginal cost across users (the costs of producing algorithm to make the index and rank results are one-off). However the costs of maintaining the search service and keeping it responsive to users may have a greater marginal component – while 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.

³⁸Strictly speaking q_u is a set not a scalar. However, it could also be interpreted as the measure of this set (and hence a scalar) as long as we are careful – in particular by requiring that the increase in size arose from taking strict supersets (one may have the case of two sets of ‘users’ A, B with $|A| > |B|$ but because of its composition of B being more valuable).

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 first assumption 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).

5. MARKET STRUCTURE

In this section we formalize some 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.

Proposition 3. *User's 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 configuration are equally like which implies that each of these search engines has equal (expected) revenue. To avoid trivial cases we shall also make the following assumption:

Assumption 4 (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 5. *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 $v + \epsilon$ 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). Thus, firm must 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. QED. □

Remark 6. This problem is very similar to the problem of a R&D race with deterministic discovery functions.

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 7. *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. $\square$

5.1. 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,³⁹ 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.

5.1.1. 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)⁴⁰ (b) pure quality competition for users (i.e. zero prices and no user heterogeneity). In our view, any model which shares these basic features is likely to feature very high levels and a single dominant firm.

³⁹This is relevant to the empirical fact that today, though there may be one firm has very large proportion of the market, there are still other firms active.

⁴⁰Recall 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.

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’.⁴¹

This existing tendency to concentration is then 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 it is important to note the implicit assumption here would be that there is a single (overall) ‘quality’ attribute which all users value positively (and that this was 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 would be that search engines specialize in searching a particular kind of content.⁴² However, any such heterogeneities are likely 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.⁴³

5.1.2. *Contestability.* Thus, it is not surprising 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 behaves rather like it would under competition. Here, even though the fixed costs are

⁴¹Just 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.

⁴²For 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).

⁴³However, adding such ‘minor’ heterogeneities would allow the model to become more realistic by predicting the existence of several small, fringe firms.

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.⁴⁴

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.⁴⁵ 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.

Coming to the second point 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). Note here that 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,

⁴⁴For example, pursuing the analogy with the R&D literature, there are a variety of result (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.

⁴⁵This is also borne out by anecdotal evidence. At a 2007 round-table on search in Toulouse Francois Bourdoncle of Exalead stated that today that the core code for a search engine was around 3 million lines and would take \$20-100 million to develop – and of course this excludes the cost of actually running the service.

it is not clear that switching costs are 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 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 who are the leaders in almost every other jurisdiction, trail Baidu (the first-mover) in China despite significant efforts on Google’s part.⁴⁶ While such jurisdictional heterogeneity, particularly where it relates to first-mover advantage, does not necessarily imply high switching costs,⁴⁷ it does, at the very least, imply that there are significant factors affecting market share’s 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. It is also noteworthy, that 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

⁴⁶It is also worth noting that Google should be considered the original ‘first-mover’ in most of the jurisdictions in which it has a lead despite not being the first to enter formally (see Table reftable-search-history for details) because all of the other companies to pre-date it in the search market either were not focused on search itself (for example Yahoo! was a directory) or fell out of contention before the importance of search (qua search) was recognized (most prominently Altavista).

⁴⁷For 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.

or increase their value to advertisers, is clearly search engine specific. It therefore also leads to increased switching-costs. These points are obviously conjectural, however there is some empirical evidence that users display increasing ‘loyalty’ to search engines. For example a Jupiter Research study from 2006⁴⁸ 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% 4 years earlier in 2002). Of these 82% refined their query on their existing search engine and 18% switched engines whereas 4 years earlier only 68% stayed with their existing engine (and 32% switched).

5.1.3. *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 not likely to be very robust.

Thus, when examining the effect of monopoly it will seem reasonable to 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).⁴⁹ Thus, in the next section, a fair degree of latitude will be assumed for the monopolist in regard of pricing and quality provision,⁵⁰ and our

⁴⁸Reported at <http://searchenginewatch.com/showPage.html?page=3598011>.

⁴⁹This contrasts with a more two-sided model such as that found in the operating systems market where the barrier to entry for the monopolist (Microsoft) is related to the existing (and therefore expected future) installed base on the two sides of the market (application providers and consumers). That said, it is possible that the search market may develop in the direction of a more platform-like, two-sided model if content owners become more active (and therefore restrictive) with regard to search engine crawls and usage of material, particularly as search engine seek to expand the pool of material they cover (consider, for example, efforts such as Google Books). In this case, the analogies with the activities of participants in other two-sided markets may become more noticeable. For example, just as Microsoft have made significant investments downstream to integrate into the ‘applications/software’ side of the market for the twin purposes of promoting consumer demand and controlling porting (see Pollock (2007)) so we are likely to see increased activity by search engine firms to move into content ownership for analogous reasons (this is already partially occurring with Google’s acquisition of YouTube, Yahoo!’s acquisition of Flickr etc).

⁵⁰If 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.

attention will be on how the monopolist's choice of these variables affect consumer and societal welfare rather than on issues of market structure and market share.

6. MONOPOLY AND WELFARE

6.1. Monopoly. Having established the focus on the monopoly case we can make some simplifications. First, total demand q_u can be interpreted as a simple scalar equal to the measure of users whose utility is positive ($u(t, v) \geq 0$). It may be sometimes useful to invert the relationship $q_u(v)$ to obtain v as a function of q_u .⁵¹ Using this we then have:

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

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:

$$R' = R_v + R_q q' = c' + c'_A \quad (6.1)$$

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).

6.2. Welfare. The first step in analyzing welfare is to define a social welfare function W . We proceed as follows:

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

Note that, following on from the previous section, we assume there is only one search engine. We have also implicitly assumed that consumer surplus and producer surplus are accorded equal value in the social welfare function. Such an assumption is reasonably standard but one could argue that the widespread and diverse set of users and the relatively concentrated ownership of most search engine companies might merit explicit

⁵¹We will assume this relationship to be invertible so that we can write $v(q_u)$ though strictly all we have is q_u is non-decreasing in quality.

distributional weights. We have not pursued this possibility but note that it would be relatively easy to introduce an explicit weighting into the analysis.

Our next step is to observe that users' utility, search engine profits and users' utility must all be inter-related. After all, 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 would like to avoid specifying in detail the form of the advertising market and the equilibrium conditions and thus we take a reduced form approach to connect advertising, search and users. 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 search⁵² $U_S(v, q_u) = \int_{q_u} U_t$ 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,⁵³ 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.⁵⁴

⁵²As before all superscript i indices used to index the search engine will be omitted as there exists only one search engine.

⁵³See for example, the arguments in Athey and Ellison (2007).

⁵⁴See 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 (adverts) on search engines may be substantially poorer than the organic results. To take one example: in May 2006 the top sponsored link

‘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 takes the final form:

$$W = U_S(v, q_u(v)) - c - c_A$$

Since advertising does not now enter the formulation for W except via c_A it is immediate that we want $c_A = 0 \Leftrightarrow R_A = 0$.⁵⁵ Thus we reduce to:

$$W(C) = U_S(v, q(v)) - c(v)$$

Maximizing with respect to v we have the socially optimal level of quality v^W solves:

$$U'(v^W) = U_v + U_q q' = c'$$

6.3. How Optimal is Monopoly? The next step is to compare this search quality v , and usage q_u , with that under monopoly as defined in equation (6.1):

$$R'(v^M) = R_v + R_q q' = c' + c'_A$$

for ‘Skype’ was download-it-free.com who, despite their name, charged \$29 to download a copy Skype, a program that is supplied for free by its producer (skype.com – the first ‘organic link for this search’). He also discusses (see e.g. <http://www.benedelman.org/news/012606-1.html>) the possible incentives for search engines to behave in this way due to the large revenues that ‘bad’ sponsored links can generate.

⁵⁵This implicitly assumes the search engine could be directly funded by non-distortionary taxation. Clearly this is unlikely to be the case and therefore even a publicly provided search engine might want to use advertising if that were an efficient way to raise revenue (it might also be politically more palatable than raising taxes elsewhere). Nevertheless, the more general point that ‘society’ would choose a lower level of advertising than the search engine is likely to be robust. Furthermore, the fact that c_A is zero will have 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).

In particular we would like to know 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).

One might think this question is rather trivial. Consider the more traditional case where R denotes revenue arising from a traditional charging regime. In that case it would normal to assume certain specific relationships between utility U and revenue R . In particular when increasing quality v :

- The utility from an extra user t ($U_q q'$) would be larger than or equal to the revenue received by the monopolist ($R_q q'$).
- The effect on existing users would be greater for utility (U_v) than for revenue (R_v).

Furthermore, with a ‘normal’ assumption of diminishing returns to quality, these functions are decreasing in v . Together with the fact that:

$$R' = R_v + R_q q' = c' + c'_A \geq c' = U_v + U_q q' \quad (6.2)$$

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

However, here this need not be the case. In particular the very fact that search engines choose *not* to charge users implies that the value that a search engine extracts from an additional user in terms of advertising revenues is higher than the price it could charge that same user, and higher even perhaps than the value of that query to the user.⁵⁶ In this latter case R_q will be larger than U_q and hence one could have, depending on the relative magnitude’s of the direct effect of quality (R_v, U_v), that $v^M > v^W$.

Similar effects would also obtain if, as is possible or even likely, revenue displays increasing returns in the number of users.⁵⁷ Increasing returns could occur for two distinct reasons. First, and most obviously, economies of scale involved in advertising on a search

⁵⁶Of course this is modulo transaction cost issues and the question as to whether a search engine could price discriminate among users as effectively as it can among advertisers. If not, then of course the search engine has the classic problem of all monopolists that it has to charge the same price to all which, depending on the distribution of user values, may not be very attractive. Search engines also has the problem that it is selling information (the query result) whose value is highly uncertain in advance. As a result users may be unwilling to pay in advance and it is difficult to extract payment ex-post. Nevertheless, and in spite of these caveats, the basic point that the value a search engine extracts from advertisers per user may be more than that user’s query value still stands.

⁵⁷And even if charges are, say, per click-through or per-view.

engine, for example those that would arise from a fixed cost in generating or placing an advert, would lead to advertising demand increasing showing increasing returns to the number of search engine users. Second, 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, would also lead to advertising demand displaying increasing returns in the number of users.

In both cases, one has that revenue has increasing returns in the number of users. It further follows that, $R_q q'$ is, at least over some portion of its domain, increasing in quality rather than decreasing. This in turn means that, not only might revenue be increasing in v , but that (total) marginal revenue R' may be larger than total marginal utility, U' , and hence that the monopolist's quality v^M may be greater than the socially optimal level v^W .

Illustrations of these two basic cases are shown in Figure 3 and Figure 4. The two models are exactly the same except for the advertising revenue function. In the first case this is simply a linear function of users while in the second the revenue function is increasing ($R \sim q^2$) and then decreasing ($R \sim q^{1/2}$) in the number of users. In the first figure, as would be expected from the previous discussion, the monopoly quality is lower than the welfare maximizing quality. In the second figure, by contrast, the monopoly quality is higher than the welfare maximizing quality, illustrating the point that if advertising revenue displays increasing returns to users the monopolist may oversupply quality.⁵⁸

6.3.1. Total Utility and Quality. Having just explored how the way advertising revenue depends on users we should now turn to the other main functional associations: the relationship of total utility to quality and users ($U' = U_v + U_q q'$), and revenue's dependence on quality R_v .

Improving search quality increases both the utility of existing users ($U_v > 0$) and the demand for search (as new queries become feasible) $U_q q' > 0'$. In the case of utility, unlike for revenue, it seems reasonable that the demand effects display classic diminishing returns, i.e. $U_q q'$ is increasing in quality but at a decreasing rate. However, U_v may display increasing returns in v even when the direct effect (U_{vv}) is negative, if the result the cross

⁵⁸Note that in this case, unlike in the first one, whether oversupply occurs will depend on the parameters. That is, even with increasing returns if the increasing returns in users are weak (or demand diminishes sharply in quality) then it will still be the case that the monopoly quality is less than the socially optimal quality.

FIGURE 3. Welfare (solid line) and Profits (dashed line) as a function of quality when revenue and user utility both display decreasing returns to quality.

FIGURE 4. Welfare (solid line) and Profits (dashed line) as a function of quality when utility displays decreasing returns and revenue (initially) increasing returns to users (and thereby quality).

term, U_{vq} , is positive enough $-\frac{d}{dv}U_v = U_{vv} + U_{vq}q'$.⁵⁹ Thus, it is possible that U' itself displays increasing returns to quality, at least over some portion of the quality range.

⁵⁹It is also not impossible, though perhaps less likely, that U_{vv} itself is positive, at least for low levels of v .

To illustrate this, we provide two concrete examples in which search quality affects utility by reducing the cost of search (leaving the value of given search constant). Recall first that the individual user utility function is $u(t, v)$ and, that with definition of $q(v)$ as the solution of $u(q, v) = 0$, we have : $U(v, q) = \int_0^q u(t, v) dt$.

- $u(t, v) = k - \frac{t}{\sqrt{v+1}}$. Intuition: each user (query) has a constant value k but a variable cost of performance (to the user) which is increasing in t and diminishing in quality v . Demand also displays diminishing returns: $q(v) = k\sqrt{1+v}$. For a given user there are diminishing returns to quality and these are sufficiently strong that total utility still displays diminishing returns: $U(v, q) = kq - \frac{q^2}{2\sqrt{1+v}} = \frac{k}{2}\sqrt{1+v}$.
- $u(t, v) = 3t^2(k - \frac{t}{\sqrt{v+1}})$. Intuition: each user (query) has an increasing value but also increasing costs. Costs are reduced by search quality. One could think of this situation as a case where queries are of increasing complexity with both the value and costs increasing (quadratically and cubically here) in that complexity. Demand, $q(v)$ is the same as before but now $U(v, q) = kq^3 - \frac{3q^4}{4\sqrt{1+v}} = \frac{k^4}{4}(1+v)^{1.5}$, which displays increasing returns to quality.

As these examples show, there can be a fairly subtle interplay between demand and the value of quality, and the overall effect can be to make total user utility display either increasing or decreasing returns to quality. It should also be clear that there is no necessary reason for utility and revenue to be closely connected in any way. In particular, investments in quality which improve the net value of existing queries, for example by reducing their cost as in these examples, would be unlikely to have any effect on demand for advertising (if anything, and as we are about to discuss, they are likely to reduce ad-revenue). If this is so, then it increases the likelihood that the socially optimal level of quality exceeds that provided by the monopolist.

6.3.2. The Direct Effect of Quality on Revenue. Finally, we come to the question of quality's direct effect on the monopolist's revenue from advertising: R_v . As already discussed, it seems obvious that the indirect effect of quality on revenue, via an increase in the user base, is positive. However the direct effect, we would argue, is likely to be negative. This is for two reasons which we label 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,⁶⁰ if a search engine is able to display 'ads' 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. 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.⁶¹ However if the search engine does not present that information then it will necessary for companies to advertise, and, once again, an increase in search quality reduces advertising revenue (and vice-versa).⁶²

The second, 'antgaonism', 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*. For example, consider the following hypothetical example in which there is a query for 'vitamin supplements' which generates both 'organic' search results as well as advertisements to firms which supply such supplements. Suppose (this is hypothetical), 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 advertisement. As such by making this information prominent one reduces the amount of advertising revenue generated from that query.

⁶⁰See e.g. Athey and Ellison (2007) for explicit argument for this approach.

⁶¹If the company is good you are already in the list and if it is bad there is no point advertising as users will know, ipso facto, it is bad.

⁶²There 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, as just discussed, it also forces those commercial organizations to buy advertising. Even outside of the commercial sphere it appears this approach is taking hold: it was recently reported that a UK Government department ended up buying Google keyword advertisements as they found this was a more effective way of getting information to potential users than relying on the 'organic' search results themselves.

Together these two effects imply that the direct impact of quality on a search engine's revenue is negative: $R_v < 0$. Of course, search quality also has an indirect impact via increasing the number of users/queries. Furthermore, it is likely that this effect is larger than the direct one: $|R_q q'| > |R_v|$. Thus, overall it is still very likely that revenue is increasing in search quality, $R' = R_v + R_q q' > 0$, as the effect of better quality in increasing queries outweighs any effect of quality in reducing revenue per query. Nevertheless, this is in contrast with the case of social welfare and users' utility where quality's direct impact (U_v) is strongly positive, and this is therefore one major reason to suppose that a private search engine will under-provide quality.

We can summarize the above discussion in the following 'proposition':

Proposition 8. *It is more likely that the monopolist's under-supplies quality relative to the socially optimal level:*

- *The smaller the advertising revenue from new users (R_q) compared to the social value of new users (U_q) (this is just 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).*

Overall, it is likely, in our view, that the under-provision effect dominates. 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 likely to be positive (and substantial) while the direct 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 reduce 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 the effects, which has already been discussed at length, arises from the fact that, for those users a search engine already has, quality acts as substitute for advertising (which is what the search engine is ultimately concerned with). Note that

this second effect, unlike the first one, is not a general one but is likely to primarily affect quality in areas where it is more directly antagonistic to, or a substitute for, advertising. Thus, it is likely more to ‘distort’ quality rather than unilaterally reduce it and for this reason we term it the ‘distortion’ effect.

7. 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 adverts 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. Clearly one cannot address every single one of the concerns that have been voiced in a single paper such as this. However, we have been able to provide a parsimonious framework which allows us to address many of the main issues in a simple but rigorous way.

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

Once explicitly considering a monopoly (or close-to-monopoly) situation, our next step, and the more important from a regulatory point of view, was 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 allowed us to reduce welfare comparisons to a comparison of search engine quality, v . It was shown that monopoly could result in both over provision and under provision of quality. However, as discussed, it is likely that the under-provision effect dominates. This was primarily 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 firm etc. If this route were to be pursued one would probably need to focus on funding basic R&D which was then 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 – 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 costs 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 will be 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’.

Such an approach in which there was 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.⁶³ 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.

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.⁶⁴ 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.

⁶³We would move back towards the high-quality, zero-profits equilibrium.

⁶⁴At 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.

8. 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, and their history we turned to the main empirical and theoretical questions that animate our investigation: the current and future market 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.⁶⁵ The first step involved developing a basic model which captured the main features of the 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. Following on from this, we showed how the structure of the search engine market, in particular that users care about quality but are not charged, while advertisers care about users and are charged, explains the highly concentrated nature of the search engine market and make it probable that the market will continue to evolve down this path towards monopoly.

Given this, our next step was to investigate the welfare performance of a monopoly, measured by the quality of search provided, as compared to the benchmark of the socially optimal provision. It was shown that a monopolist would deviate from optimal provision in a variety of ways and was likely to provide an inefficiently low search quality (and engage in ‘distortion’ of its organic results).

Given this it seems likely that some form of oversight, possibly including formal regulation, will become increasingly necessary. Part of this effort could include taking steps to encourage a more diverse search environment. However, the structure of the search

⁶⁵By contrast there has already been substantial work on particular aspects of search engines such as their methods for auctioning advertising space (see references in the main text).

market, in particular its great economies of scale, may undermine the potential for, and benefits of, vigorous market competition, especially in the long run. When monopoly, or near monopoly, does obtain it was shown that there is no guarantee that the private interests of a search engine and the interests of society as whole will coincide – and good reasons to think otherwise.

It is therefore likely that search, if left entirely unregulated, will develop in ways that are not always to the benefit of society as a whole. For this reason it is important that policy-makers start now on the process of developing their strategy in relation to this key area of the knowledge economy. The power rapidly accumulating in the hands of a few major search providers is a great one. It behoves to ensure that it is used in a way that brings the greatest benefit to society as a whole.

REFERENCES

- Mark Armstrong. Competition in Two-Sided Markets. *RAND Journal of Economics*, 37 (3):668–691, Autumn 2006.
- Susan Athey and Glenn Ellison. Position Auctions with Consumer Search, October 2007.
- John Battelle. *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Portfolio, September 2005. ISBN 1591840880.
- Michael R. Baye and John Morgan. Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets. *American Economic Review*, 91 (3):454–474, June 2001.
- Tim Berners-Lee. *Weaving the Web: Origins and Future of the World Wide Web*. Orion Business, 1999.
- Yongmin Chen and Chuan He. Paid Placement: Advertising and Search on the Internet. Working Papers 06-02, NET Institute, September 2006.
- Ben Edelman. Adverse Selection in Online "Trust" Certifications, 2006.
- Ben Edelman. Assessing and Improving the Safety of Internet Search Engines. In *The Rising Power of Search Engines on the Internet*, pages 259–277, 2007.
- Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review*, 97(1):242–259, March 2007.

- Glenn Ellison and Sara Fisher Ellison. Search, Obfuscation, and Price Elasticities on the Internet. NBER Working Papers 10570, National Bureau of Economic Research, Inc, June 2004.
- Neil Gandal. The Dynamics of Competition in the Internet Search Engine Market. *International Journal of Industrial Organization*, 19(7):1103–1117, July 2001.
- Anindya Ghose and Sha Yang. An Empirical Analysis of Search Engine Advertising: Sponsored Search and Cross-Selling in Electronic Markets. Working Papers 07-35, NET Institute, September 2007.
- Avi Goldfarb and Catherine Tucker. Search Engine Advertising: Pricing Ads to Context. Working Papers 07-23, NET Institute, September 2007.
- Christopher Harris and John Vickers. Perfect Equilibrium in a Model of a Race. *Review of Economic Studies*, 52(2):193–209, April 1985.
- Bernard Jansen, Mimi Zhang, and Zhang Ying. The Effect of Brand Awareness on the Evaluation of Search Engine Results, 2007. CHI 2007, April 28-May 3, 2007, San Jose, California, USA.
- Volker Nocke, Martin Peitz, and Konrad Stahl. Platform Ownership. *Journal of the European Economic Association*, 5(6):1130–1160, December 2007.
- Rufus Pollock. The Control of Porting in Two-Sided Markets. Cambridge Working Papers in Economics 0754, Faculty of Economics, University of Cambridge, November 2007.
- Jean-Charles Rochet and Jean Tirole. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association*, 1(4):990–1029, 06 2003.
- Avner Shaked and John Sutton. Natural Oligopolies. *Econometrica*, 51(5):1469–1483, September 1983. ISSN 00129682.
- John Sutton. *Sunk Costs and Market Structure: Price Competition, Advertising, and the*. MIT Press, 1991. ISBN 0262193051.
- John Sutton. *Technology and Market Structure: Theory and History*. The MIT Press, 1998. ISBN 0585077533.
- Rahul Telang, Tridas Mukhopadhyay, and Uday Rajan. Market Structure for Internet Search Engines. *Journal of Management Information Systems*, 21(2):137–160, 2004.
- Hal R. Varian. Position Auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.

Alex White. Search Engines: Left Side Quality versus Right Side Profits, 2008.