Has Google Taken Over the Internet?

A Literature Review of Search Advertising: The Boom Market of the 21st Century

Michael Anthony

Queen's University

September 18, 2012

Abstract

This paper aims to familiarize the reader with the literature on search advertising, with a particular focus on issues pertaining to Google's dominance over the market and the antitrust issues this poses. Particular attention is paid to the mechanism by which advertisements are purchased from search engines (now known as the "position auction") and how an oligopolist might engage in anticompetitive activity in this market. The final section of the paper offers some remarks on the state of current antitrust cases against Google, and on the future outlook of the market for search advertising.

Acknowledgements

This paper was only possible with the assistance of Dr. Roger Ware, who was gracious enough to offer his supervision despite being on sabbatical at the time. I would also like to thank Tyler Longo, Eryn Fanjoy, and Jalna Meens for their support and assistance. I am grateful for the support of my other peers as well, who are too numerous and whose various assistances are too many to mention here. This paper would not have been possible without their help.

Contents

1 Introduction

In the past two decades, the internet has risen from a collection of computers sending kilobytes of data to each other via phone lines to arguably the largest source of information in the history of human civilization. Now a primary mode of commerce, communication, and entertainment, it seems almost quaint to recall that internet browser software once had to be purchased at retail stores, and numerous small search engines competed for several hundred thousand unique users. With the number of unique users now measurable in billions, many firms that were able to effectively monetize the internet have grown from tiny one room operations into industry juggernauts with thousands of employees and market capitalizations in the hundreds of billions of dollars.

There is perhaps no more poignant example of such a success story than Google, a firm so ubiquitous that it is listed as a verb in the Oxford English Dictionary ("let me Google that").[1] With 83% of all online searches worldwide according to one estimate (and closer to 90% in Europe and Canada),[2] Google makes up a massive proportion of search engine traffic, and by monetizing all these page views they have had enviable financial success. As Google offers almost no end user consumer retail products,[3] the bulk of their revenues come from various advertising programs: in 2011, for example, 96% of the company's revenue came from these sources (at approximately $36.5B).[4] One of the most important ways this is done is by "search advertising" or "search engine marketing," which is the practice of serving ads to users of a search engine, and tailoring these ads to the search query they input. For instance, if a user types in "garden tools" into a Google, Bing, or

[1] Google was officially added to the OED on the 6th of July, 2006; this was regarded as more like headline fodder than real news at the time, but it bears mentioning that neither "Yahoo" nor "Bing" are verbs in any major English language dictionary.
[2] See http://www.netmarketshare.com/search-engine-market-share.aspx?qprid=4 visited 2012-07-27
[3] Or rather, almost no end user consumer retail products that it directly monetizes.
[4] See http://investor.google.com/financial/tables.html visited 2012-07-27

other another search engine, they will receive a set of “organic” search results (those that are not sponsored by an advertising campaign) sorted by some algorithm designed to pick relevant search results, and they will also receive a set of paid advertisements that are targeted at the user based on their search query. This type of advertising is a rapidly growing market, and according to the Search Engine Marketing Professional Organization, it grows significantly from year to year. To highlight just how successful this market has been, consider its performance during the recent financial crisis (Table 1).

Table 1. The growth rates of search advertising in North America as compared to growth of the North American Economy during the aftermath of the financial crisis.

Year	North American SEM Spending (Billions)	Year over year growth**	Growth of US GDP***
2007	$12.2	30%	1.9%
2008	$13.5	11%	-0.4%
2009	$14.6	8%	-3.5%
2010	$16.6	14%	3%
2011	$19.3*	16%*	1.7%

*Projection. **SEMPO does not publically release more precise metrics, so numbers are rounded.
***According to the World Bank (http://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG) visited 2012-07-27.

While Yahoo and Bing have made some progress in gaining access to this market (including a deal in 2010 that saw them combine search distribution services) and Chinese search engine Baidu is rapidly gaining market share in the Asia-Pacific region, Google has become the undisputed champion in the battle over search advertising. In line with their control over internet search, Google is estimated to control as much as 75% of search advertising in the United States, and over 90% in Europe.[5]

This paper aims to cover part of the new but rapidly developing literature surrounding the rising phenomenon of search advertising and to discuss the antitrust challenges that

[5] See http://www.covario.com/phocadownload/design/q4-11_gpssa_fnl.pdf visited 2012-07-27

are posed by this unusual market. For the most part, this literature has focused squarely on two issues: the structure and implications of Google's pricing mechanism (which has now become the industry standard), and on Google's tremendous and potentially increasing lead over its competitors. Section 2 of this paper is focused on Google's pricing scheme, and attempts to explain the current literature on the "position auction" and how this enables Google to extract revenues from their customers. While this literature is new, it is already very expansive, so I focus on a few key papers that make up the canonical literature. Section 3 discusses competition issues in the industry, and covers some of the empirical justification for the theoretical assertions that have been made. In section 4, I make some modifications to a model by Frederico Etro that describe the potential for direct anticompetitive action by Google; these modifications rely on some more realistic assumptions on the market, though his result is preserved. I conclude with a brief future outlook of the search engine advertising market and a discussion of some of the regulatory challenges it faces. As much of the literature existed for some time as working papers prior to being published, the chronology is somewhat complex. Thus, rather than approaching the literature from an economic history standpoint, it is the overall aim of this paper to familiarize the reader with some of the antitrust challenges surrounding Google and search advertising in general.

2 The Position Auction

In first seeking to describe the sale of advertising slots to purchasers, Lahaie (2006), Iyengar and Kumar (2006), and Edelman *et al.* (2007) independently determined that Google and Yahoo were both using an auction mechanism that was not well described by previous literature. Previous efforts had either focused on one advertising slot, as in Liu and Chen (2006), or had been focused on other, more general auction systems, as in

Ellison *et al.* (2004). Further complicating matters, they noted that there are a few properties of search advertising auctions that make them unique: bidding is continuous and ongoing, and thus a bidder may change their bid at any time (this is actually observed, since bidders can change their bids from second to second through software automation); the marginal cost of serving an ad to a user who is already receiving search results is functionally zero, and thus a search engine that does not sell slots for a certain keyword is wasting potential revenue; and it is difficult to describe search advertising in terms of units. The advertiser, for instance, is interested in the cost per additional customer (cost per acquisition, or CPA), whereas the search engine is interested in the amount it collects per ad displayed (usually given as "CPM," or the cost per mille).[6] A pay system in which ads are priced on a CPA basis is an ideal option for an advertiser, since the risk for purchasing or bidding for ads is essentially zero (as they will collect a sale for each payment they make to a search engine), and a pay system in which ads are priced on a CPM basis is ideal for search engines, as they will earn stable revenue from any ad they sell. As a compromise between these two methods, search engines generally have auctions for search advertisements on a cost per click basis, or CPC, in which advertisers bid their willingness to pay for a single click on their ad.[7] In some senses, this makes search advertising an ideal topic for matching theoretical models to real world situations, but it also means that a novel *ad hoc* approach would be more suitable than attempting to fit an existing auction mechanism to this new problem.

Calling these new auctions generalized second price auctions, or GSP (though this model is also known as EOS, after its discoverers), Edelman, Ostrovsky, and Schwarz proposed

[6] The construction "cost per mille" is common in other forms of advertising, and is the cost per one thousand views. This is also known as "cost per impression" or CPI.
[7] CPM was the industry norm prior to the introduction of the CPC method in 1997 by Overture, which was later purchased by Yahoo. CPA has never seen wide use in internet advertising.

that that a search engine acts as an auctioneer, and it auctions off N advertising slots to K webpages seeking to attract traffic (hereafter referred to as "advertisers"). All the advertisement slots that are sold are displayed on a search page for that search query, with the owner of slot 1 appearing first, and so on down to the Nth slot. The slots are allocated according to the amount that an advertiser indicates as its maximum willingness to pay (this need not be their true MWTP), with the highest bid receiving slot 1 (which receives the most clicks due to its higher position on the page), the second highest bid receiving slot 2 (which receives the second most clicks), and so on. When a user clicks on the nth advertisement, the owner of the advertisement pays a sum of money to Google equal to the next highest bid (whose bidder occupies the $n+1$ position)[8], all the way down to $min\{N,K\}$. By winning the ith position in an auction, a bidder would receive an expected number of clicks a_i (where $a_i > a_{i+1}$ for every i), which bidder k values at s_k.

Edelman *et al.* note that the auction seems superficially similar to a Vickrey-Clark-Groves (VCG) mechanism, but a key difference is that in a GSP auction, truth telling is not a dominant strategy (by contrast, it is a well known property of VCG that truth telling is a dominant strategy). To see this, consider an example in which there are three slots on a page and four bidders. Suppose an advertisement in slot 1 receives 500 clicks per hour, while an advertisement in slot two receives 450, and an advertisement in slot 3 receives 100. Bidders 1, 2, 3, and 4 have valuations of $10, $7, $3, and $2, respectively. Consider this a continuous game in which bidders may revise their bid in response to another player at any point (this is an important feature of the way Google ad auctions work). In a VCG auction, winners of the auction pay the externalities they impose on the other

8 Edelman and Ostrovsky (2007) and Feng (2008) determined that when Overture used a more traditional bid scheme, strategic behavior by advertisers lead to a "price cutting war" similar to that seen between competing gas stations, and this caused major losses in revenue. The GSP mitigates that behavior considerably, although they determined that strategic behavior on the part of advertisers was still ongoing in this new mechanism.

bidders. That is, by entering the auction, the top bidder forces bidder 2 out of slot 1, bidder 3 out of slot 2, and bidder 4 out of slot 3. He imposes an externality of 50 clicks at $7 per click on bidder 2 ($350), 350 clicks at $3 per click on bidder 3 ($1050), and 100 clicks at $2 per click on bidder 4 ($200); thus bidder 1 must pay a total amount of $1,600 to have slot 1. This generates a payoff of $3,400. By the same logic, the second bidder pays $1,250, the third bidder pays $200, and the fourth bidder pays nothing. There is no possible benefit to be made by any one player by increasing or decreasing their bid. In a GSP auction, however, they pay the price set by the next highest bidder. If they are bidding truthfully, this means that bidder 1 pays $7 per click, bidder 2 pays $3, and bidder 3 pays $2, for payments to the auctioneer of $3,500, $1,350, and $200 (note that the revenue for the auctioneer is $5,050 under GSP in this example, as opposed to $3,050 under VCG). Now consider that bidder 1 earns a payoff of only $1,500 here. If he bids between $7 and $3, he pays $3 per click instead of $7 and earns a payoff of $3,250 ($4,500-$1,250). Note that if bidder 1 bids $3.01, the lowest possible bid to defeat bidder 3 and the amount that produces the lowest revenue for the auctioneer, a GSP auction generates revenues of $3,055 – higher than in a VCG auction.

Work by Agarwal *et al.* (2006) has attempted to develop an auction system that gives the same slot assignment as this GSP while still inducing truth telling. They noted that if click through rates were separable into a merchant specific factor and position specific factor (that is, click through rate is indexed by i, the merchant, and j, the slot), the VCG mechanism is applicable to the GSP problem. The setup of their model is similar to that of Edelman *et al.*, with N merchants bidding for K slots, and a click through rate (i.e. the rate at which an advertisement is clicked on) that is given by $CTR_{i,j}=0$ for $j>K$ and is nonzero otherwise. There is also some externally specified ranking function where merchants receive a weight w_i that is independent of their bid and causes them to be ranked by w_ib_i. A truthful auction can be achieved by using a pricing scheme that "ladders" prices by computing the price that a merchant would have paid for being a slot below their current slot and charging them only that price for the clicks they would have received anyway had they been in the lower slot. For clicks beyond those they would have received anyway, they are charged the amount necessary to remain in that slot in a GSP. The price is then given by the following.

$$p_i = \sum\left(\frac{CTR_{i,j} - CTR_{i,j+1}}{CTR_{i,i}}\right)\frac{w_{j+1}}{wi}b_{j+1}$$

They find that an auction using this pricing scheme will prove to be a profit maximizing truthful auction, and may lead to efficiency gains over a setup in which bidders strategically revise their bids in order to undercut opponents. This is known as a "laddered auction" as it recursively determines the prices for clicks. The difficulties here are that while it is feasible for a modelling exercise, it is intractable to specify the function by which click through rates are determined in the real world. The other difficulty is that while there is an obvious equilibrium for this laddered auction, the authors note that this generally earns less revenue for the auctioneer than the GSP auction.[9]

In fact, Edelman *et al.* point out that for any case in which advertisers bid the same amounts, the revenue for Google under GSP is weakly greater than revenue under VCG (which the laddered auction is similar to if click through rates are separable). Consider the advertiser in position N, the last slot. Under VCG they pay the externality of eliminating advertiser $K+1$ (i.e. $s_{k+1}*a_k$) who would have occupied that Nth slot otherwise, and under GSP they pay the bid of the player beneath them (which will be s_k, since truth telling is the dominant strategy in VCG, and we are concerned about the case where the

[9] In the provided example, three bidders with valuations of 500, 480, and 100 with $w_i = 1$ for all i bidding over two positions generate a revenue for the auctioneer of 54 in the laddered auction, and 111 in the GSP.

bids are the same). They will receive a payoff of a_k, and they have the same payment in both schemes. For an advertiser i where $i<min\{N,K\}$, the increment in payment amount is given by $(a_i-a_{i+1})b_{i+1}$, or what the advertiser in the next slot pays plus the externality imposed on them. This is strictly lesser than the payment increment under GSP, which is given by $a_ib_{i+1} - a_{i+1}b_{i+2}$ (there is a minimum increment over the next bidder to prevent the scenario in which all advertisers bid ε more than their nearest competitor). Thus GSP produces higher payments (and thus higher revenues for a search engine) for any given vector of bids. As it happens, in a static bidding game in which bidders pursue what is called a "locally envy-free equilibrium" (i.e. one in which no advertiser can increase profits by exchanging slots with the advertiser in the slot above them, such that for $a_is_i-p_i \geq a_{i-1}s_i-p_{i-1}$, for all $i \leq min\{N,K\}$) it is true that the worst possible equilibrium is equivalent to the equilibrium in a VCG auction. Though it is potentially social welfare decreasing, Google's use of GSP generates better revenues for them than VCG.

Varian (2007) describes another game theoretic setup of a related auction model and considers the maximum and minimum revenues in both Nash equilibria and symmetric Nash equilibria. The setup of the model is similar, as it is concerned with fitting a advertisers into s slots, where advertiser a's valuation of slot s is given by $u_{as}=x_sv_a$, where x_s is interpreted as a click through rate, and v_a is simply the value to advertiser a of each visit to its webpage. Each advertiser has its own v_a, and $x_1>x_2>...>x_n$. Advertisers bid b_a and are awarded a slot based on their bid. For each click on their link, they pay a price p_s, where $p_s=b_{s+1}$, and thus the owner of slot s gains profits given by $(v_a-p_s)x_s=(va-b_{s+1})x_s$.

This model treats the auction as a simultaneous move game with complete information, again reflecting the reality that advertisers can change their reported WTP at any time. In equilibrium, each agent will prefer its own slot to all other slots so that the profits for an advertiser to own slot s must be higher than the profits were that same advertiser to own any other slot. That is, no agent will want a slot lower or higher than the one they have, so that $(v_s-p_s)x_s \geq (v_s-p_t)x_t$ for $t>s$, and $(v_s-p_s)x_s \geq (v_s-p_{t-1})x_t$ for $t<s$. In a symmetric Nash equilibrium, this is modified to $(v_s-p_s)x_s \geq (v_s-p_t)x_t$ for all t, s.

Varian notes several matters of practical interest regarding position auctions. First, in the symmetric Nash equilibrium of a position auction, there is non-negative surplus, as $(v_s-p_s)x_s \geq (v_{S+1}-p_{S+1})x_{S+1}=0$, since $x_s=0$ for all $s>S$. Additionally, advertiser values are monotonic, so that advertisers with higher values are assigned to better slots, and thus the SNE is an efficient allocation. Prices are similarly monotonic, so that better slots cost more, and it is optimal for the first excluded bidder (i.e. the advertiser in position $S+1$) to bid their value. From this standpoint, the position auction approach seems ideal.

However, Varian concedes in his paper that some refinements are needed in order for this model to accurately describe Google's own auction system, though he does not discuss the revenue impacts of these refinements. Since 2002, Google has ranked advertisements (that is, it has allocated its slots on the search results page) according to the product of the advertiser bid and something called the "ad quality" or quality score. This is simply a measure of an ad's relevance to a search result, although the method by which is it determined is quite complex. It is determined by an algorithm that considers the past click through rate of the advertisement and the past click through rate of all the ads an advertiser has submitted using their AdSense or AdWords account. It also accounts for the quality of an advertisement's "landing page" (i.e. the webpage that a link directs to).[10] The precise method by which these scores are calculated is not public (both to avoid abuse by unscrupulous advertisers, and presumably to allow Google to extract information

[10] See http://support.google.com/adwords/bin/answer.py?hl=en&answer=2454010 visited 2012-07-28

rents), and Google's instructions for increasing the scores are not explicit, so it can be difficult for advertisers to increase their ad quality score. To model this, the click through rate of slot s is called z_s, which is given as the product of the so called "quality effect" e_s and the "position effect" x_s, such that $z_s = e_s x_s$ (note the similarity to the weighting mechanism proposed by Agarwal *et al.*). Since an advertiser's bid determines its position, they are effectively ordered by $e_s b_s$, and advertiser s pays the minimum necessary amount to stay in slot t, q_{st}. This makes the quality score incredibly important, as it can variously save or cost an advertiser large sums of money. It is illustrative to consider the following example of quality score that Google provides to its AdWords customers:[11]

There are three advertisers, labeled A, B, and C. They are each bidding over some arbitrary keyword, and provide bids (which are the maximum CPC they are willing to pay) of \$0.40, \$0.65, and \$0.25 respectively. In the simple position auctions without quality score considered above, this would result in B taking the highest slot, A taking the second, and C taking the third. However, the advertisements all have quality scores; advertisement A has a score of 1.8, B has a score of 1.0, and C has a score of 0.5. This causes the following allocation of slots (table 2).

The actual CPC is the amount that an advertiser pays to retain their spot. That is, advertiser A pays \$0.37 per click, as this is all that he would require to beat advertiser B's rank when using B's maximum bid, and likewise B only pays \$0.14 per click, as this would result in a rank score of 0.14, which is higher than advertiser C's. By contrast, advertiser C pays the minimum bid of \$0.20, which is set by Google itself (and is

[11] See http://services.google.com/awp/en_us/breeze/5310/index.html visited 2012-07-28

generally set to penalize low quality ads). To illustrate the importance of ad quality, if advertiser C increases their bid to 1.5, then the situation changes dramatically (table 3).[12]

Table 2. An example of how the Actual CPC is calculated.

Advertiser	Max CPC	Quality Score	Rank	Position	Min Bid*	Actual CPC
A	$0.40	1.8	0.72 (1.8x0.4)	1	0.04	$0.37
B	$0.65	1.0	0.65 (1.0x0.65)	2	0.05	$0.14
C	$0.25	0.5	0.13 (0.25x0.5)	3	0.20	$0.20

*This is determined by Google, and not stipulated by the advertiser, who provides only their maximum CPC as a bid.

Table 3. Another example of how the actual CPC is calculated.

Advertiser	Max CPC	Quality Score	Rank	Position	Min Bid	Actual CPC
A	$0.40	1.8	0.72 (1.8x0.4)	1	0.04	$0.37
B	$0.65	1.0	0.65 (1.0x0.65)	2	0.05	$0.39
C	$0.25	1.5	0.38 (0.25x1.5)	3	0.04*	$0.04

*This change is arbitrarily made in the example, and is not a consequence of a change in one of the other attributes.

Now advertiser B is actually spending more money than advertiser A to have an inferior slot. In fact, Varian notes that the additional difficulty in managing this quality score has spawned an entire industry of "Search Engine Marketers" (or SEMs) who offer a variety of products, including ad quality consultation, and placing ads on an advertiser's behalf.

[12] See http://services.google.com/awp/en_us/breeze/5310/index.html. Visited 2012-07-28.

Also note that the minimum bid quantity essentially represents the reserve price of advertisement sales for Google (although it is adjusted by Google based on the attributes of an advertiser's account history). Given that the minimum bid amount is usually quite low (unless it is increased to control potentially malicious advertisements, spam, or because an account is otherwise deemed suspicious), one might believe that reserve prices were not relevant in the market. To some degree this is true, especially for high value keywords where just bidding the minimum payment of $0.05 would mean that an ad would never be displayed. Moreover, minimum bids are determined based on an account holder's quality score, with high minimum bids assigned to account holders with lower quality scores. It then stands to reason that advertisers with high minimum bids must spend large sums of money anyway in order to beat competing advertisers with good quality scores, and thus the minimum bid would seldom be binding.[13] Nevertheless, given the structure of the GSP, Edelman and Schwarz (2010) show that increasing the reserve price has a large indirect effect on prices, as bidders bound by the minimum bid have their CPCs moved upwards, which in turn inflates the CPC for the next advertiser, and the next, and so on. As a result, the minimum bid can be used as a sophisticated way of increasing advertiser bids exogenously, while making the increase in price appear to be an endogenous feature of the market. In 2008, Google attempted to downplay the importance of the minimum bid by removing it as a displayed metric in the Google AdSense account statics, opting to replace it with a "first page bid" (or the price that a given user would need to bid to win the first slot).[14]

Athey and Ellison (2011) discuss the implications of the reserve price and note that Google has an incentive to have higher than optimal minimum prices. They also note

[13] See http://adwords.blogspot.ca/2006/01/common-adwords-misconception-explained.html. Visited 2012-08-13.
[14] See http://adwords.blogspot.ca/2008/08/quality-score-improvements.html. Visited 2012-08-13.

that much of the literature on position auctions is abstract and deals with "objects" rather than advertisements. Taking this into account, they have proposed a novel model for quality score weighted auctions. The setup of the model is quite different from any of the previous models discussed in this paper, and builds on work by Chen and He (2006) that introduces consumers as acting agents rather than an exogenous source of clicks.[15] In the model, a search engine displays M links, and a consumer j may click on them at a cost s_j, which follows some distribution G over [0,1]. There are N advertisers, and firm i has a probability q_i of meeting some need that consumers have, where q_i is drawn from a common distribution for all firms F over [0,1]. When an advertiser meets a need they get a benefit of 1, and when a consumer's need is met they get a benefit 1. The auction itself is modelled as an M stage game. In the first stage of the game, players submit their bids, and the N-M lowest bidders are eliminated. In the next round, all firms submit simultaneous bids, the lowest bid is assigned the lowest of the M slots, and then the remaining players bid over the remaining slots until they have all been assigned. The advertiser in slot i makes a payment given by b^i. For example, the advertiser in slot 3 makes a payment given by b^3.

Consumers engage in Bayesian updating, so if a link is clicked by a user and their need is not met, all advertisements in a lower slot are determined to have lower expected value; if they continue long enough without meeting their need, eventually the cost of clicking the next link will be greater than the expected value, and they will stop. This is an obvious justification for employing a quality score to match consumers to relevant ads: if consumers discover that they are able to meet their needs by clicking Google ads, they are more likely to click Google ads in the future (and thus produce revenues for Google).

[15] Varian (2006) claims that a conversion rate of 1 in 1000 is reasonable, which is orders of magnitude more effective than other forms of advertising.

Regardless of the other aims of the quality score and the GSP, it is true that rankings by auction provide information to consumers and this information increases welfare for all parties.

In order to incorporate quality scores in the model, firms are given a type (δ, q), where a firm of this type meets the needs of a fraction of consumers given by δq. Consumers can look at the advertisement and immediately determine the likelihood that the advertised site will meet their needs. A fraction of consumers, $1-\delta$, realizes the ad does not meet their needs, and a fraction δ realizes it may. If they click the ad, they incur a cost s, and there is a probability q that the site will meet their needs.

Here the effect is that a search engine acting as auctioneer can rank sites by $\delta_i b_i$ in the multistage auction game described above. In general, this results in a higher number of successful matches (and is thus a welfare improving result). However, they also note that this creates incentives for obfuscation: the total payment for firms ends up being $b^{k+1}\delta^{k+1}$, and so the incentive is to make the click through rate as large as possible, which will in turn lead to large revenues for the search engine itself. Thus any quality score mechanism of ranking websites effectively becomes a way of ranking sites by the amount of revenue they generate for a search engine.

3 Competition Issues in Search Advertising

While the position auction has been a major success for Google (and possibly its users, though its advertising customers and competitors would likely take issue with its informal motto: “Don’t be evil”), it is not the only one to use such an approach, and yet it is clearly the dominant firm. While the literature on the position auction represents a new foray into auction theory, the literature that has developed around competition between search engines is a more natural extension of the pre-existing literature on two-sided matching, and much of it owes a considerable debt to Rochet and Tirole (2003). What is notable about this literature is that because search engines are such a unique product, and because the current market structure is so topical, the literature that has developed is extremely cognisant of the particular tripoly structure that presently defines search advertising.

It seems at first glance that Google would be unable to exercise any sort of market power (and thus they are dominant simply by virtue of the fact that they are the best search engine), especially given that there are two major competitors that consumers can access with only a few key strokes. However, Argenton and Prufer (2012) have done work indicating that competition is not actually "only a free click away." While Argenton and Prufer are very explicit in insisting that they take no position on whether or not Google has committed any abuses of its market power, their model demonstrates that the search engine industry is prone to "market tipping" (that is, one large firm is able to use its size to grow faster than the other firms, causing the "balance" of the market to "tip" in its direction). They propose that intertemporal indirect network externalities result in a cascading feedback loop that makes a popular search engine increasingly popular. This is because search engines store logs of the results from a particular search query and note which links are most successful, then iterate on subsequent search results using an appropriate algorithm. As a result, there is a very large learning-by-doing element to conducting effective searches, and consequently, large search engines "doing" more (attracting more search traffic) are able to "learn" more (and thus match search results to their traffic better). Since they claim that 71% of search engine users believe that accurate search results are the most important feature of a search engine, this means that matching a user's query to an appropriate search result will generate additional traffic, thus making that search engine most attractive to advertisers.

They approach the market in an ad hoc fashion (though their model could be expanded to other scenarios), modelling a tripoly with firms 1, 2, and 3, and a unit mass of consumers who demand one query. Firms face some fixed cost F and compete for the unit mass of consumers by choosing a quality x_i, and pay some cost for this $C=x_i/N_i$, where N_i is simply the previous queries run on i. They earn revenue p, which is simply the advertising revenue associated with receiving an additional user. Firms then seek to solve the following problem:

$$\max_{xi}\ \pi_i = \frac{x_i}{x_1 + x_2 + x_3} p - \frac{x_i}{N_i} - F$$

This is modelled as a one shot game in which firm 1 chooses a quality level, then firm 2 chooses a quality level, and last, firm 3 chooses a quality level. Profit functions are given by the following:

$$\pi_1 = \frac{(N_2 - N_1(1 + N_2))^2}{(N_2 + N_1(1 + N_2))^2} p - F$$

$$\pi_2 = \frac{(N_2 + N_1(N_2 - 1))^2}{(N_2 + N_1(N_2 + 1))^2} p - F$$

$$\pi_3 = \frac{(N_2 + N_1(1 - N_2))^2}{(N_2 + N_1(1 + N_2))^2} p - F$$

What is somewhat surprising is that if N_i is interpreted in a dynamic way, as it increases, it causes firm 3 profits to fall until firm 3 exits. If the market is reinterpreted as a duopoly, then firm 2 profits fall until firm 2 exits. The market thus "tips" towards monopoly in the long run.

As a justification for Google taking the role of "firm 1", consider that at the time when

Google first introduced their new mechanism in 2002, the search engine market was considerably less concentrated than it is now. The serious participants in the search advertising markets were AltaVista, Ask Jeeves, Looksmart, Overture, Hotbot, Google itself, and Yahoo, which was still the largest search engine at this time.[16] Most of these used CPM pricing to display banner ads rather than the vastly more attractive CPC pricing (in fact, all of them used CPM except Overture, who famously sued Google over their adoption of the CPC method in 2002), and as such, although all these search engines were effectively incumbents, by introducing a quality score and CPC mechanism, Google was able to tailor its advertisements to its users. Thus obnoxious banner ads and popups gave way to relevant material that may have added to the attractiveness of using Google over other search engines. A justification for this comes from a study by Yang and Ghose (2010), in which they demonstrated significant interdependence between paid search results and organic search results (and by extension, the search engine). Though they determined organic search results had a 3.5 fold greater effect on user utility from clicks than paid search results, if one considers their increasingly sophisticated revision algorithms Google began to employ, it is reasonable to suggest that Google became a leader in a new market (and this market, rather than the one that existed prior to Google's technological leap, is prone to tipping).[17] The decrease of major global firms in the market from seven to three in spite of the incredible increase in the size of the market lends credence to the strength of the model. Pollock (2010) has corroborated the view that the market is progressing towards monopoly, as the competition is well characterized by winner-takes-all, and suggests that immediate regulatory action is critical. Analysis by

[16] See http://www.submitexpress.com/newsletters/jan_30_03.html visited 2012-08-03.
[17] Yahoo, one of the first to follow Google's example (and the only remaining competitor from the early 2000s), finally introduced a quality score equivalent with its new ranking algorithm Panama in 2007, five years after Google did in 2002. During this time, it hemorrhaged traffic, while the other competing search engines simply determined that it was not profitable to compete with Google and exited the market.

Motchenkova *et al.* (2012) suggests that a single large firm would have incentive to undertake actions that decrease consumer welfare even in the presence of several smaller competitors.

Spulber (2009), working along similar lines, corroborates this idea that information accumulation leads to dominance in the market. He identifies a "map of commerce," which follows the interplay between information revealed by advertisers and information revealed by users. Users reveal their preferences to the search engine by clicking on links in search results and avoiding others, which leads to a refinement of the advertisements and search results server, and advertisers reveal information such as their willingness to pay for advertising slots. By contrast, search engines, in spite of having FAQ pages for their advertising services, reveal comparatively little.[18] Search engines thus create strong information asymmetries which they can use to extract rents. Work by Borgers *et al.* (2007) suggests that these rents could be very large. Spulber notes that this effect depends strongly on the level of competition on the market (it is decreasing in the amount of competition). Since this level of competition may actually be very low (in spite of having at least three major firms in the market), this is a serious consideration.

An interesting point involving this information asymmetry that is raised by White (2008) and later by Taylor (2012) and Xu *et al.* (2010) is that while search engines attract users with accurate search results, they gain no benefit whatsoever when users do not click advertisements. The search engine would thus prefer for a link that is a paid advertisement to be clicked rather than an organic search result. It is thus in the interests of the search engine (and the advertiser) to make the matching between advertisement and search term as precise as possible. But this also leads to a conflict of

[18] A quick search for information on AdWords or AdSense, for example, will reveal countless internet forum posts devoted to understanding how the supposedly simple mechanism works.

interest for the search engine: if their organic results are sufficiently excellent, then users will tend to click on these very frequently. In fact, it has been empirically demonstrated by Hotchkiss *et al.* (2005) that over 77% of search engine users perceive organic links to be better sources of unbiased information than paid search results. C*eteris paribus*, if a search engine matches paid advertisements with search terms as well as they match organic search results with those terms, they will see less clicks on their advertisement links, and thus lower revenues: their success at organic search result matching causes them to "cannibalize" their own revenues.

White builds a model of a monopolistic search engine in which users use a search engine to determine their valuation of some good. This peculiarity is difficult to match intuitively to our notion of search, so as a motivating example, consider a user that is trying to accomplish some objective that can be met in multiple ways, such as selling a house. By typing "how to sell a house" into a search bar, they may discover the degree of difficulty involved in doing it themselves, and thereby determine their private valuation of hiring a realtor (who will probably have an advertisement for that keyword). In the model, a user i faces a cost $\rho(\theta_i,s)$ for using the search engine that is decreasing in the quality of search that it provides, s, and increasing in their own personal "query cost" of using the search engine, θ_i. Search engines face no cost for increasing their quality, but instead produce more relevant links on the left side of the page $l(s)$, which produce no revenue if clicked.[19] If a user i searches for a good and purchases it, they receive utility $u_i=v_i\text{-}p\text{-}\rho(\theta_i,s)$. If they search and do not purchase, they receive $u_i= -\rho(\theta_i,s)$, and if they do not search, they receive a utility of $u_i=0$. The game takes place in several stages, first beginning with search engines selecting an advertiser fee of A, and a quality of s, after

[19] White argues that these websites will always be competitors for those paying for advertisement space, as even if they began as only informational websites, they will naturally begin to offer sales services or advertise them themselves once they realize the amount of traffic they receive. As there is arguably only one major informational website that employs no advertising in Wikipedia.org, this logic seems sound.

which users search for a good which is advertised by $n=l+r$ firms (where r firms are paying advertisers) engaging in Cournot competition. It happens that in equilibrium, there is a considerable incentive to offer an inferior service even if it is costless.[20]

Taylor takes an even less rosy view, and notes that if consumers had search engine loyalty, they could be attracted to a search engine via an initially high quality (or a periodically high quality) and then served inferior search results since demand reacts smoothly and there are frictions in changing to another search engine. Thus there could be incentive to exploit users and advertisers for short periods of time before returning to the sort of search quality that attracts and maintains users. He also suggests that part of the notion behind the Yahoo or Google search bar for internet browsers is to induce users to choose one search engine over the other. It is probably true that marketing search bars is partly motivated by the fact that they lower the cost of using that particular search engine, but they are also marketed aggressive during installations for unrelated products, which would support Taylor's theory. This is also a possible purpose of Google's proprietary browser, Chrome, which is available as a free download, and whose address bar doubles as a Google search bar. As it happens, Jansen *et al.* (2009) empirically demonstrate that such loyalties and frictions do exist. They conducted an experiment with 32 participants measuring four search engine brands[21] and found that there was a 25% difference in average relevance rankings by participants despite the fact that the results were identical in content and presentation. Both Google and Yahoo demonstrated better than average performance, though Yahoo's was the best rated by a considerable margin, Google in second, MSN in third, and the fictional AI^{2}RS in last by a considerable

[20] This is obviously not the case in practice, but it only serves to strengthen the argument that a dominant search engine would underprovide search quality.

[21] These being Yahoo, Google, MSN (now known as Bing), and a search engine developed just for the study called AI^{2}RS.

margin. Based on this, it is reasonable to speculate that Yahoo was still thought of as the premier brand. Given the experiment took place in 2006, before Google had developed the same public image it now has, it is possible that if the experiment were repeated today with a new study cohort, Google would be perceived as more state of the art than its competitors. Whether or not this is true, Google likely has considerable brand loyalty at this point, and this would allow them to serve inferior search results compared to the user utility maximizing quality without causing user flight to other search engines.

Hagiu and Julien (2011) demonstrate that there is also an incentive to manipulate search results so that users who click on advertisements are sent to websites that do not match their need well. While this trades off against the desire to maintain users by ensuring that their search experience is accurate (and thus low cost), this also generates additional revenue when users click the back button to go click additional links. Given the somewhat enigmatic nature of quality score determinants, in practical terms this could be easily managed by Google in a fashion that would easily withstand anything but the most determined regulator scrutiny: simply artificially increase the quality score initially (via, for example, the landing page score) and the link will be placed higher on the page. This effect is inertial, since it starts a feedback loop in which the advertisement's higher position on the page leads to more clicks, which then lead to a higher quality score, and so on.

Agarwal, Hosanagar, and Smith (2011) empirically identify conditions under which an advertiser receives a higher click through rate on their ads when an organic link for their page places very high in organic search results, though curiously this results in a lower conversion rate in spite of the additional clicks. This is odd, as conversion rate is increasing in user perceived "website quality" (which, all things being equal, would also cause the organic and sponsored link rankings to increase). It is not yet clear what the implications of this are from an antitrust perspective or if this result is true for all keywords. If it were, it would indicate that a search engine has significant incentive to promote its advertisers' links in the organic search results as well, at the expense of the user who is looking for unbiased organic search content.

4 A Model of Google as a First Mover

As mentioned previously, one of the interesting features of the search advertising market (and by extension the search engine market) is the lack of entry, which seems surprising given that in its infancy, the internet was something of a "Wild West." Though brand loyalties, frictions in changing to another search engine, and other features of the market have been offered as explanations for why this could be, a search engine may be able to engage in more traditional anticompetitive activity to prevent entry. Antitrust issues in multisided markets have been explored by Evans (2003a,b), and Etro (2011) uses a similar approach to describe Google in particular. He proposes that one of the ways that Google limits the entry of competitors is through their quality score.[22] Having a high ad quality score for an advertiser's Google AdWords or AdSense account actually imbues all that advertiser's subsequent ads with an increased quality score just by virtue of coming from the same account. In order to have earned this high ad quality score for their account, the advertiser must have consistently produced high quality ads in the past (i.e. ads that were clicked on frequently and had good landing pages, among other less well defined features). Because of this, a high account ad quality score effectively represents a major sunk cost in that every advertisement that has ever been purchased by an account adds to

[22] The difficulty in transferring ad campaigns to other networks (resulting in de facto exclusivity) was actually a major component of the antitrust case against Google in Europe that saw Google and the EU settle in the summer of 2012.

the value of that account. Switching to another search engine effectively results in the abandonment of that sunk cost, and a new account ad quality score has to be built up from scratch. This means that advertisers that have spent lots of resources on one platform are effectively committed to it. Etro also notes that switching to another competitor also involves the additional sunk costs of expertise in managing any platform specific idiosyncrasies, such as training staff to use Yahoo's services instead of Google's, as well as the costs of transferring ad campaigns to a new system, and the associated risks of switching to a new platform.

In order to determine if Google has the potential to use its first mover advantage to stymy competition, he presents a model of two-sided competition in which platforms court consumers on one side, and advertisers on the other. This takes the form of a standard Cournot, where c_i is the marginal cost of a click for platform i, A_i is the number of ads available through platform i, Ci is the number of consumers reached by these ads, p_{ai} is the price charged to advertisers for an click, and p_{ci} is the cost per interaction to attract consumers. Note that p_{ci} can be more simply thought of as the general costs of running the search engine divided by the number of interactions that occur (the consumers have no real interest in the ads, just the various qualities of the search engine encompassed in the model by p_{ci}, which users themselves pay nothing for). Thus search engine profits can be written as $\pi^i=(p_{ai}-p_{ci}-c_i)A_i*C_i$, or, if A is a vector of advertisement choices so that $A=(A_1,A_2,\dots A_n)$, then search engine profits can be given as $\pi^i=(p_i(A)-p_{ci}-c_i)A_i*C_i$. Etro chooses to consider the number of platforms as exogenous, given the apparently large barriers to entry (as evidenced by the fact that the market in North America has consisted of Yahoo, Google, and Bing for some time).

In a simultaneous move game where platforms choose the number of advertisers and consumers they attract (intuitively this can be thought of as advertisers choosing the number of consumers via their p_{ci} decision, and the number of consumers subsequently attracts advertisers), the result is a symmetric Cournot equilibrium with the following first order conditions:

$$A_i\!: p^i(A) - p^i(C_i)c_i + \frac{dp^i(A)}{dA_i}A_i = 0$$

$$C_i\!: p^i(A) - p^i(C_i) - c_i - \frac{dp^i(C_i)}{dC_i}C_i = 0$$

However, if one search engine were able to act as a leader the results change markedly. As Google was a first mover in introducing a quality score metric for search advertising, and given that the ostensible reason for a quality score metric is to improve matching consumers to advertisements they would be interested in – something it admittedly achieves much better than a pure ranking on the basis of WTP – having a leader in this model is a better match for the actual market as it exists. By using the click weighted auction (i.e. incorporating quality scores into their auctions), Google can manipulate the WTP of its advertisers, which amounts to an effective increase in the price of advertisement slots on their web searches. Etro indexes these effects as a term ϖ which can be thought of as the amount a search engine invests into developing a sophisticated mechanism (thus increasing its ability to match advertisements with appropriate web searches), or more simply as the accuracy of the mechanism. The profit function for the leading platform in this case is very similar to the auction without quality score mechanisms:

$$\Pi^i = [p^i(A,\varpi) - p^i(C_i) - c_i]A_iC_i$$

Etro's profit function incorporates costs associated with acquiring additional advertisers and increasing the sophistication of the auction mechanism into the c_i term, and does not directly account for the fixed costs incurred by choosing ϖ. These costs are likely to be significant, so I expand his analysis slightly by including them in the profit function as a fixed cost. The marginal cost term c_i is omitted, as the marginal cost of each additional display is very small (much less than one cent per ad served)[23] and the significant costs associated with search advertising are attracting consumers to the search engine, the research and design of the system, and the physical and human capital costs required to maintain it. As additional costs involved in increasing the capacity of advertisers using any system are likely to also be significant (for example, Google cannot automate assigning a score to the landing page of a sponsored link, and thus they must have employees evaluate each one), the fixed cost of the ranking mechanism is also modelled as a function of A_i. For tractability, C is also considered as exogenous, as it is heavily dependent on marketing efforts and the effectiveness of the search engine's organic search. Based on this, it is not an enormous stretch of the truth to presume that C is unaffected by a search engine's decisions of ϖ and A. The profit function of a search engine in this model is given by the following:

$$\Pi^i = \left[p^i(A,\varpi) - pc^i(C_i)\right]A_iC_i - F_i(\varpi, A_i)$$

Given Etro's decision to model the number of platforms exogenously, this game is a natural extension of Fudenberg and Tirole's (1984) animal taxonomy nomenclature. By checking the equilibrium marginal cross-effects of ϖ and A, it is possible to determine the strategic incentives of the leader:

[23] Google serves 180 billion impressions per month. Given that cost of revenues of $4B for 2011, this would suggest a marginal cost of at most 0.19 cents per ad served. If less favorable assumptions were used, the cost would be much lower. See http://www.wordstream.com/blog/ws/2012/05/15/ipo-facebook-vs-google-display-advertising and http://investor.google.com/financial/tables.html visited 2012-08-21

$$\frac{\mathrm{d}^2\Pi^1}{dA_1 d\varpi} = \left[\frac{\mathrm{d}^2\mathrm{p}^1(A,\varpi)}{dA_1 d\varpi}A_1 + \frac{\mathrm{dp}^1(A,\varpi)}{d\varpi}\right]C_1 - \frac{\mathrm{d}^2 F_1(A,\varpi)}{dA_1 d\varpi}$$

By the definition of ϖ, higher ϖ increases the price of each interaction, and thus $dp^i/d\varpi > 0$ for all ϖ. Etro makes the assumption that this effect is likely to be diminished by a greater number of advertisers, so the advertiser elasticity of $dp^i/d\varpi$ is less than unit ϖ, and thus while the first term in brackets is negative, the bracketed first term as a whole is positive. This comes from the idea that $A_i=A^i(p_a)$, where p_a is the vector of all platform prices; the rationale for this is that platforms set prices and then advertisers make their advertising decisions with each platform based on the prices set by those platforms so that $dp^i(A)/dA_i < 0$ for any i. This does not reflect the actual behavior of the market: in practice, the price is set by the number of competing advertisers (remembering that the only price a search engine stipulates is the minimum bid, and this creates an indirect effect that is not diminishing in the number of advertisers), and certain extremely valuable keywords that produce a high number of conversions at a relatively high price (for instance, those relating to computer software or insurance) are much more expensive to bid on than those that are less popular (for instance, those relating to landline telephones).[24]

Here it is taken that $dp^i(A)/dA_i > 0$, implying that both of the bracket terms are positive. This is a better match with our intuitive understanding of the search advertising market, given that consumers on each platform are exogenous in this game (increasing the number of advertisers does not decrease the payoff to the advertisers who win the position auction except for the increased cost in acquiring clicks), and also given that these auctions are a

[24] A click for keywords including “software” (for example, “marketing software” or “technical lab software”) in 2011 had an average CPC of $35.29; a click for “landline” might cost about $1.25. See http://www.wordstream.com/blog/ws/2011/07/18/most-expensive-google-adwords-keywords and http://www.spyfu.com/term.aspx?t=879607 visited 2012-08-23.

mechanism to increase willingness to pay (thus driving up the price). This would imply that increasing ϖ for any amount A_i of advertisers would increase the willingness to pay of those advertisers, and, vice versa, that increasing the number of advertisers bidding for slots for any given mechanism ϖ would weakly increase price. Taking the bracketed terms to both be non-negative, the sign of the cross effect is then determined by the fixed cost term. This term is clearly non-positive, as adding additional advertiser capacity to a complex mechanism (higher ϖ) would be more complicated (and thus costlier) than adding it to a simpler mechanism(lower ϖ), and increasing the sophistication of a mechanism that served a few advertisers would be costlier than increasing the sophistication of a mechanism that served many. Thus the value of the equilibrium cross effect is positive or negative dependent on the following inequality:

$$\left[\frac{\mathrm{d}^2\mathrm{p}^1(A,\varpi)}{dA_1 d\varpi}A_1 + \frac{\mathrm{dp}^1(A,\varpi)}{d\varpi}\right]C_1 \lesseqgtr \frac{\mathrm{d}^2 F_1(A,\varpi)}{dA_1 d\varpi}$$

This is somewhat ambiguous depending on the particular forms that the functions take. In the event that the magnitude of the cross effect on the RHS is large, this would imply a lean and hungry look from Fudenburg and Tirole's taxonomy. The leader then has an incentive to limit the services they offer so that they appear aggressive to competitors. In the context of the organization as a whole, this would entail limiting services external to their AdWords and AdSense services to appear as though they had committed to advertising; in the context of the model, this would likely entail under developing their mechanism to appear tough to competitors.[25] By contrast, if the LHS is larger, then the cross effect is positive, and this is an example of a top dog scenario. In the context of the organization as a whole, the search engine has an incentive to invest in many different

[25] Based on the fact that this is manifestly not the approach Google has taken in either case, we can safely assume that the LHS is large, which is logical given that it implies a large consumer base, or that Google's mechanism is effective at extracting rents at the equilibrium.

products and services; in the context of the model, a search engine would have incentive to overinvest in the mechanism itself.[26] The ways this limits competition are myriad. An example of this is that businesses that primarily focus on online sales depend on search advertising for a significant part of their total sales. If a sophisticated auction mechanism caused online retailers to have high advertising budgets, it is unlikely that they could afford to spend similar sums to advertise with a smaller entrant. Equally, a sophisticated auction mechanism would have a higher conversion rate of clicks to purchases than a less sophisticated one as the matching algorithm would be superior, so advertisers would be unwilling to switch to a less sophisticated entrant that produced fewer matches. In either case, a leader in the new market would have an incentive to blockade entry and engage in actions that reduce consumer welfare.

5 Conclusion

The fact that so much of the literature on the search advertising market is at least partially concerned with monopolization is not coincidental. Over the past several years, Google has faced investigations from a number of major regulatory bodies. These have ranged from accusations of copyright violations by Google's Google Books service,[27] to the ongoing Samsung vs. Apple tablet saga (which Google has considerable interest in, as they are the developers of the operating system for Samsung's tablet), and, of course, alleged abuses of its search neutrality and AdSense. The most noteworthy example of this is the major antitrust action between the EU competition bureau and Google, in which it was alleged that: Google displays its own vertical search services (these are specialty services,

[26] This, by contrast, seems to be representative of the state of Google's present strategic position.

[27] A lawsuit launched in 2005 by the Authors Guild and another by the Association of American Publisher alleged that by allowing digital copies of copyrighted materials to be accessed online, Google was committing copyright infringement. Google announced in October 2008 that it settled for $125 million. See http://www.authorsguild.org/advocacy/articles/settlement-resources.attachment/authors-guild-v-google/Authors%20Guild%20v%20Google%2009202005.pdf visited 2012-08-26.

such as “Google News”), effectively leading to an increase in traffic to websites integrated with these vertical search services and a decrease in traffic for those who are not;[28] that Google indexes reviews from other websites and displays them on its own services, denying those websites the resulting traffic;[29] that AdSense traffic acquisition agreements prevent Google’s competitors from displaying ads on websites that publish Google advertisements; and finally, that AdSense is designed in such a way that makes ad campaigns non portable, and thus means that organizations facing limited advertising budgets focus all of their efforts on Google. At present, Google and the EU are set to settle (much to the chagrin of competition lawyers who would otherwise be looking forward to issuing a stack of large invoices), at least on some of the issues that have been raised, and the terms of the settlement are currently being negotiated over.[30]

Of course, provided that it can continue to attract users to its search engine, an organization like Google is almost certainly sophisticated enough to find ways of meeting any requirements issued by an antitrust body without actually resolving the abuses they allege. Some of these have been discussed previously in this paper (for example, tying quality score to an advertiser’s account history), and given the complexity of the mechanisms involved in search advertising, there are countless others still waiting for some enterprising engineer at Google to develop (new products created by Google are likely to affect competition, perhaps even inadvertently).

One of the important implications of this complexity is that regulators may not be properly equipped to deal with these problems.[31] A case in Brazil launched by a shopping

[28] Work by Tarantino (2011) strongly suggests that this is indeed the case.
[29] For example, Google aggregates restaurant reviews and displays them for its Google Places service.
[30] See http://searchengineland.com/europe-offers-google-thorny-olive-branch-finds-market-power-abuse-but-offers-to-settle-quickly-121943 visited 2012-08-27.
[31] The 2008 financial crisis provides a very recent example of how sophisticated products can lead regulators to totally misjudge the state of the market.

comparison website called Buscapé made similar allegations, and was recently ended by a summary judgment.[32] It is the opinion of the author that the judges' assertions are dubious. For example, it was determined that Google was not a monopoly, since the plaintiff can be found by methods other than Google. This is an unconvincing definition for a monopoly, especially given Google's 90% market share in Brazil. It was also alleged that Google Shopping is not a shopping comparison site, and therefore not a competing product (it is instead just a thematic search option within the generic search available to Google). While it may be true that Google Shopping is a "thematic search option," suggesting it is not a competitor (even though it performs the exact same function in the exact same fashion) is analogous to insisting that a $20 discount for an energy efficient appliance is not a "subsidy" because the word "rebate" is written on the receipt. Services like those offered by Google are not well defined by classical notions of competition,[33] and there is a tendency for competition officials to view these through a more traditional lens, such as might be appropriate for a physical goods market.

One interesting possibility is that as major consumer technology companies expand into new markets, it is conceivable that Apple might create some search engine technology of their own (their recent transition from using Google Maps on their iPhone product line to the proprietary service Apple Maps is compelling evidence that they are attempting to limit their dependence on Google's services). Their Siri software (a combination voice recognition and digital personal assistance software package) is, after all, just a particular sort of search engine devoted to answering certain queries. Given their technological sophistication, Apple could produce a service that was of comparable quality to Google's.

---

[32] See http://www.scribd.com/doc/105502055/BUSCAPE%CC%81-vs-Google-Summary-Judgment-ruling visited 2012-08-27.

[33] Devine (2008) notes that search engines do not compete on price so much as they attempt to "displace one another's products entirely."

In addition, their brand is so strong that users of Apple's mobile and computer products would probably switch over to a new Apple search engine in droves. There is perhaps no other firm who is better poised to be a competitor for Google in the market for Search Advertising, should they choose to enter.

For the moment, however, it seems that Google will face nothing like the public scrutiny that Microsoft did during its period of operating system industry dominance, and that its dominance is likely to continue. What eventual outcome the settlement between Google and the EU will result in is unclear, but this case serves to highlight the importance of the ongoing research on this industry. Through a better understanding of its complexities, it will be possible to meet the impending regulatory challenges of the future, and ensure a fair and efficient market going forward.

6 References

1. Agarwal, G., Goel, A., Motwani, R. (2006). Truthful Auctions for Pricing Search Keywords, *Proceedings of the 7th ACM Conference on Electronic Commerce.*
2. Agarwal, G., Hosanagar, K., Smith, M. (2011). Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets. *Journal of Marketing Research,* 48:1057-1073.
3. Argenton, C., Prufer, J. (2012). Search Engine Competition with Network Externalities. *Journal of Competition Law and Economics*, 8(1): 73-105.
4. Athey, S., Ellison, G. (2011). Position Auctions with Consumer Search. *The Quarterly Journal of Economics,* 126(3): 1213-70.
5. Borgers, T., Cox, I., Pesendorfer, M., Petricek, V. (2007) Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence. Working Paper.
6. Chen, Y., He, Y. (2006). Paid Placement: Advertising Search on the Internet. *Working Paper.*
7. Devine, K. (2008). Preserving Competition in Multi-Sided Innovative Markets: How Do You Solve a Problem Like Google? *North Carolina Journal of Law & Technology.* 10(59): 59-118.
8. Edelman, B., M, Schwarz, M. (2010). Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions. *Working Paper.*
9. Edelman, B., Ostrovsky, M. (2007). Strategic Bidder Behavior in Sponsored Search Auctions. *Decision Support Systems*, 43(1): 192-198.
10. Edelman, B., Ostrovsky, M, Schwarz, M. (2007). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *The American Economic Review,* 97(1): 242-59.
11. Ellison, G., Fudenberg, D., Mobius, M. (2004). Competing Auctions. *Econometrica,* 77:427-52.
12. Etro, F. (2011). Google Economics: A Model of Leadership in Search Advertising. *Working Paper.*

13. Evans, D. (2003). Some Empirical Aspects of Multi-sided Platform Industries. *Review of Network Economics*, 2(3): 191-209.
14. Evans, D. (2003). The Antitrust Economics of Multi-Sided Platform Markets, *Yale Journal of Regulation*, 20: 325-81
15. Feng, J. (2008). Optimal Mechanism for Selling A Set of Commonly Ranked Objects. *Marketing Science*, 27(3): 501-12.
16. Fudenberg, D., Tirole, J. (1984). The Fat-Cat Effect, the Puppy-Dog Ploy, and the Lean and Hungry Look. *American Economic Review.* 74(2): 361-66.
17. Hagiu, A.. Wright, J. (2011). Multi-Sided Platforms. *Working Paper.*
18. Hotchkiss, G., Garrison, M., and Jensen, S (2005). Search Engine Usage in North America. *An Enquiro Research Initiative.*
19. Iyengar, G., Kumar, A. (2006). Characterizing Optimal Keyword Auctions. *Proceedings of the 2nd Workshop on Sponsored Search Auctions.*
20. Jansen, B., Zhang, M., Schultz, C. (2009). Search Engine Brand and the Effect of User Perception of Searching Performance. *Journal of the American Society for Information Sciences and Technology.* 60(8): 1572-1595.
21. Lahaie, S. (2006). An Analysis of Alternative Slot Auction Designs for Sponsored Search. *Proceedings of the 7th ACM Conference on Electronic Commerce.*
22. Liu, D., Chen, J. (2006).Designing Online Auctions with Past Performance Information. Designing Online Auctions with Past Performance Information. *Decision Support Systems,* 42: 1307-1320.
23. Motchenkova, E., Lionas, I., Bartelsman, E. (2012). Market Dominance and Quality of Search Results in the Search Engine Market: Analysis of Exploitative and Exclusionary Abuses. *Working Paper.*
24. Pollock, R. (2010) .Is Google the Next Microsoft? Competition, Welfare and Regulation
25. in Online Search. *Review of Network Economics*, 9(4): 1-29.
26. Rochet, J., Tirole, J. (2003). Platform Competition in Two-Sided Markets. *Journal of the European Economic Association,* 1(4): 990-1029.

27. Spulber, D. (2009) "The Map of Commerce: Internet Search, Competition, and the
28. Circular Flow of Information". *Journal of Competition Law and Economics*, 5(4): 633-682.
29. Taylor, G. (2012): .Search Quality and Revenue Cannibalisation by Competing Search Engines. *Journal of Economics & Management Strategy*, forthcoming.
30. Tarantino, E. (2011). Vertical Search Engines Foreclosure. *Working Paper.*
31. White, A. (2008). Search Engines: Left Side Quality versus Right Side Profits. *Working Paper*.
32. Varian, H. (2006). The Economics of Internet Search. *Rivista di Politica Economica*, lecture.
33. Varian, H. (2007). Position Auctions. *International Journal of Industrial Organization*, 25:1163-1178.
34. Xu, L., Chen, J., Whinston, A. (2009). Too Organic for Organic Listing? Interplay between Organic and Sponsored Listing in Search Advertising. *Information Systems Research*, forthcoming.
35. Yang, S., Ghose, G. (2010). Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence?. Marketing Science, 29(4): 602-623.
36. Covario Industry Report (2012). Fourth Quarter Provides Strong Finish to 2011 Paid Search Trends.
37. Econsultancy. (2008). State of Search Marketing Report 2007. SEMPO: New York
38. Econsultancy. (2008). State of Search Marketing Report 2008. SEMPO: New York
39. Econsultancy. (2010). State of Search Marketing Report 2009. SEMPO: New York
40. Econsultancy. (2011). State of Search Marketing Report 2010. SEMPO: New York