First International Conference on Internet Technologies and Applications, Wales, September 2005, to appear

COOPERATIVE STRATEGIES FOR SEARCH AUCTIONS

Brendan Kitts, Parameshvyas Laxminarayan, Benjamin LeBlanc

iProspect, 311 Arsenal Street,
Watertown, MA. 02472-2602. USA.
bkitts@excite.com

ABSTRACT

Pay per click (PPC) search auctions are a method for selling positions on search engines that were the largest recipient of advertising spending on the internet in 2004. Little formal analysis has been conducted into this auction model. In this article we examine the potential for collusion to occur. We show that price collusion can occur between participants when their Nash equilibrium prices are adjacent. We use this to develop a coordination algorithm that optimally distributes profit on the auction between participating players.

KEYWORDS

Pay per click, placement, search engine, Overture, auction, collusion, cooperation, Google

1. INTRODUCTION

A Pay Per Click (PPC / Pay For Placement / Pay For Position) auction allows participants to bid on sponsored positions in search engines. The two largest PPC auctions in 2004 were *Google* Adwords and *Yahoo* PrecisionMatch, each of which generated over a billion dollars in revenue.

Pay Per Click (PPC) auctions consumed 40% of all internet advertising in 2004, making them bigger than any other form of internet advertising [16][17]. Even if other forms of internet advertising become popular in the future, the size and success of this auction is a major reason for the academic study of this internet application.

Various authors have reported that prices have been increasing on these auctions. Price increase estimates of about 20% per year are common [10][15]. Our own tracking data indicates that prices are increasing at about 35% per year [11], and SEC filings from the search engines show a similar trend (figure 1). Given these price increases, it is rational for companies to turn to collusive practices to try to keep prices down and profit margins intact.

Search engines too, may benefit from questionable practices. The revenues involved in search engine auctions are so high, that an increase in fraudulent clicks of just a few percent, for example, could have a dramatic impact on real revenue.

In this article we will discuss so-called "deviant" strategies for pay per click auctions. We will describe different types of search engine and bidder collusion, and under what conditions collusion may be possible. In particular we focus upon a method for optimally re-pricing bids such that payoff for players are maximized. We conclude by discussing the wider issues surrounding cooperation, including the impact of collusion on the search engines and public at large.

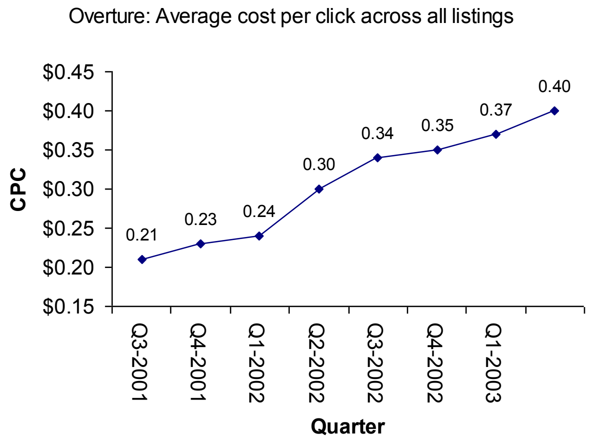

Figure 1: Average auction prices as reported by Overture's Quarterly SEC filings.

2. RELATED WORK

A considerable literature exists on auction collusion. Collusive strategies such as so-called Bidding Rings and Bidder signaling have been discussed [19]. A "ring" is a group of individuals who agree to forgo bidding to allow one member to win the product, and then divide up the proceeds afterwards. Theoretical aspects of collusion have been dealt with by [13]. [6] describe statistical approaches for detecting collusion, such as tests for independence between bidding on several auctions in which other participants feature. [9] describe the effects of collusion in the recent 1996 FCC Spectrum auctions. Bidders had a strong incentive to cooperate because they were attempting to acquire a portfolio of licenses in different geographic areas. Because direct communication was disallowed, collusion was accomplished by bidders using the floating point part of their bid price to signal to what extent they were willing to escalate their bidding war.

3. DEVIANT SEARCH ENGINE STRATEGIES

Search engines may cheat on auctions in a variety of ways.

Shills: Shill bidders are bidders who serve to drive up the price of keywords for the benefit of the search engines. Shill bidders are probably not viable in pay per click auctions because bidders can run a search and verify their rank and the legitimacy of the other bidders on the auction. However, a similar role is played, unintentionally, by affiliates. Affiliate marketers are independent arbitrage bidders who bid on behalf of a company and deliver traffic to that company. The company in turn pays the affiliate for the traffic and/or conversions. The problem emerges when the affiliates bid more aggressively than the parent company, in effect increasing all bid prices. *Google* recently made it illegal for more than one advertisement to be displayed for the same destination site, thereby solving the "user experience" problem of having redundant listings shown on search auctions. This in effect forces the affiliates and parent into a private auction, the winner of which is displayed on the search auction. However this does not address the fundamental problem of affiliates bidding higher than the parent.

Disclosure of Private Bidder Information: Keyword discovery is a major task in pay per click bidding and, once discovered, the keywords are treated like trade secrets. However, search engines may disclose this information to competitors to encourage them to enter new auctions and drive up prices.

Lax Enforcement of Click Fraud: Click Fraud is the practice of artificially generating clicks on advertisements, thereby costing the advertiser more but not leading to any additional conversions. If the search engine does not invest in resources to detect and remove fraudulent clicks, then the fraudulent clicks will inflate the volume of clicks to the advertiser, and so the fees the advertiser must pay. For instance, if 50% of clicks were fraudulent, and bidders did not respond and change their prices, the Search Engine would generate twice the revenue. *Google* and *Yahoo* are both at pains to assure advertisers that they aggressively use systems to detect and remove fraudulent clicks prior to reaching the advertiser. We put this to the test in Figure 3. This figure shows the clicks recorded by our independent tracking system on PPC advertisements (we counted the number of times our server fielded a request from one of our keywords) versus the clicks reported by the *Yahoo* search auction on the same. The search engine appears to under-report clicks by approximately 15% over a long period of time, which is consistent with the hypothesis that the search engine is not charging for some clicks.

Scattershot Matching: Search engines have a strong incentive to display advertisements as much as possible. Every time a search is conducted and no pay per click advertisements are shown, the search engine loses money. Under so-called "Broad-match" keyword settings, the keyword that an advertiser buys can match phrases that use the keyword in any order; for instance, "under ground" may match "ground under", "metropolitan under ground", "under water" [2]. What most people do not know is that the keyword matching conducted by the search engines is far more aggressive than even this example suggests. If a bidder buys a keyword such as "Acura accessory", this keyword will actually attract traffic from search queries that do not include any of these words (figure 2). Perhaps the search engines are using co-occurrence analysis between searches to find queries that are related [4]. Some search engines are better at this than others - one anonymous search engine matches on average 16 unique phrases per keyword purchased, where-as another on the same set of keywords matches only 2.8 unique phrases. We believe this significantly affects the ability of the search engine to generate revenue from their auction. The use of such advanced methods to connect advertisers with traffic, however, opens the possibility for abuse. The search engine could charge for queries, for which the advertiser with their choice of keywords, did not want to pay.

Access Quotas: Some search engines have introduced access quotas to limit the number of times per day that bidders update their bids, and so preserve communications and server bandwidth. This is a reasonable action to take when bidders are using machine-mediated methods to change bids - as anyone who has written a web spider knows (myself among them!), it is easy to accidentally saturate a server. However, in some cases rapid updates are necessary in order to exploit lucrative daytime consumer patterns such as lunch-time and early-evening shopping behavior (see for example [12]). If bidders do not have enough quota to exploit these patterns, bidders will be forced to bid low throughout the day. We therefore believe it is in the search engine's interest to allow for rapid bidder updates, and believe that access could simply be regulated by imposing a nominal cost per access. This would only make it economical to execute fast updates if there is a real profitability case for taking this action, rather than sloppy bidding agent design.

Other Areas: Because search engines are responsible for the clicks and prices, there are a lot of areas in which they may potentially cheat. Clicks can be audited by bidders fairly easily by setting up their tracking URL to re-direct through a tracking server. However, prices are more difficult to audit. For instance, search engines could pass the price and position of the click to the destination URL in the query parameters, thereby allowing the advertiser to validate the price they have been charged.

Keyword for which Advertisement to be Shown	**Queries for which the advertisement was displayed**
acura accessories	acura rsx interior accessories
acura accessories	where can I get stuff for my acura
acura accessories	accura accessories
acura accessories	acura assessories
acura accessories	acura grills

acura accessories	engine upgrades for integras
acura accessories	integra aftermarket
acura accessories	rsx armrest
acura accessories	mdx hitch
acura accessories	replacement mirrors

Figure 2: Search engine user queries that led to the advertisement associated with keyword "Acura accessories" being shown.

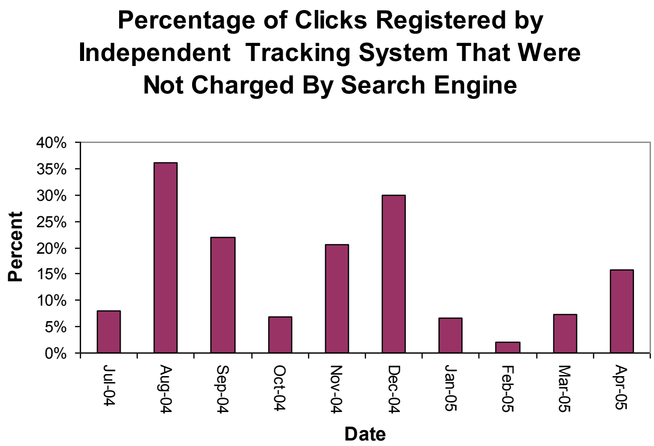

Figure 3: Over a 10 month period the above search engine under-charged us by an average of 15%. We believe this is evidence that the search engine is identifying fraudulent clicks and not charging customers for those clicks.

We show elsewhere via a relationship we call *Ryan's Theorem* [11] that if bidders set bids so as to ensure that Cost Per Acquisition (CPA) remains fixed, then strategies that attempt to increase charges, yet erode conversion rates or keep them the same, such as Click Fraud, Scattershot Matching or Price fudging, would be ineffective - bidders would change their bids to keep their CPA constant and the engine fees would not increase. As efficiency decreases, bidders would divert budget away from these engines. This leads to the interesting result that for many types of potential fraud, search engines may be able to be self-regulated.

4. DEVIANT BIDDER STRATEGIES

Many types of implicit and explicit cooperation are possible for bidders on Pay per click auctions.

(a) Bidding Rings: Rings are not known to exist on keyword auctions at this time. A bidding ring could be formed by a single party "the ring leader" (similar to what in the industry are called "click aggregators") bidding on the auction. They would receive the traffic and then re-direct the traffic to "ring members". However search engines employ human editors to verify that landing pages conform to certain editorial requirements, such as being free from pop-ups. As such, it is likely that these redirection rings would be discovered by editors.

(b) Information Sharing: There are many ways in which information revelation could benefit cooperating companies. Companies could share information about "best advertising practices". Alternatively, auction information could decrease exploration costs. In sealed bid auctions such as *Google Adwords*™, the relationship between price and position is not disclosed to the bidder. In order to bid rationally, agents need to expend resources by placing "spectulative bids" in order to explore the prices of surrounding positions. By sharing exploration, companies would simply make available their information gleaned on price-positions, and in return each company would need to make very few exploratory bids, and many more optimized bids.

(c) Display: It may be possible for companies to coordinate their advertising creative, colors, branding, in order to attract the eye. For example, imagine a block of three competitors that have coordinated their bids so that they all occupy positions 1 to 3. By moving together and displaying consistent coloring, it is possible that these companies could attract a greater number of clicks than either would have attracted alone (figure 3).

(d) Self-sacrifice: It is possible for a company to subvert its own utility by dropping out of a profitable position to increase the profit for other companies. Although this might seem undesirable, this strategy would be plausible when multiple divisions of the same company are competing on the same auction.

(e) Price cooperation: It may also be possible to increase the utility of all parties by coordinated bidding so as to avoid unnecessary price escalation. This last form of collusion we will refer to as "price cooperation". This is interesting because none of the utilities of the participating players are subverted, and all participating players can potentially increase their utilities. In the remaining part of this article we focus on price cooperation between bidders.

5. PRICE COOPERATION

5.1 Formal Definition of PPC Auction

PPC auctions are second price, perpetual auctions run by search engines. Each competitor enters an offering bid b_k which is the most they are willing to pay should a customer who typed phrase k click on their advertisement in the search results. The search engine sorts the prices and awards position p to the p-th highest bid. We will write $b_k^{(p)}$ for the bid at position p. The clearing price, $b_k^{(p)+}$ which the participant pays to the search engine, is one cent above the bid of the next competitor immediately below; $b_k^{(p)+}=b_k^{(p+1)} + o$.

Utility functions for the participants are usually a function of the click traffic that reach their site minus advertising costs. For example, a simple profit utility π_k would equal expected margin per click r_k minus the advertising costs b_k, multiplied by the number of clicks received at the awarded position; $\pi_k =(r_k- b_k^{(p)+})\, c_k$

5.2 Nash Equilibrium

Let an action profile $b_k^{(1..P)}$ be a vector of prices paid by competitors in positions 1..P. This action profile will form a Nash equilibrium if and only if every player's bid is a best response to the other players' actions. A best response b_i is a bid which generates the highest utility for competitor i given that competitor j will choose their own utility maximizing action.[20]. In terms of a PPC auction, a Nash equilibrium is a vector of prices such that no competitor i can increase their utility by changing their price, ie. $\pi_{k,j}^{(p)} \geq \pi_{k,j}^{(i)}\ \forall i,j$. Proofs showing the conditions under which bidder prices converge to a Nash equilibrium can be found elsewhere [11][14]. $\wp$

5.3 Theorem 1: Conditions Necessary for Price Cooperation

If the profit generated by all cooperating companies must be greater than or equal to the profit generated by the same companies at their Nash equilibrium prices, then it is necessary for the cooperating companies to have adjacent Nash equilibrium prices.

Proof

Suppose at Nash equilibrium we have non-cooperators n_i interspersed between friendly companies c_1 and c_2 with each agent bidding the exact minimum o cents over the competitor immediately below. Thus the auction looks like $b^{(1..N)}=(n_1, \ldots, n_j, c_1, n_{i+1}, c_2, n_{i+2}, \ldots, n_C)$ where n_i are non-cooperators, and c_1,c_2 are companies willing to cooperate. If c_1 were to move to another position above or below; either n_{i+1}-0.01 or n_i+0.01, it would generate the same profit. If c_1 were to change its price between n_{j+1}+0.01 and n_j-0.01 it would generate less profit by definition (4.11). Hence if either company were to move away from its present position it would generate less profit and this would contradict the requirement for generating more profit. As a result, it is impossible to generate higher profit for both companies. ℘

5.4 Cooperative Pricing Algorithm

We can now sketch an algorithm for finding mutualistic prices. First assume we are controlling a set C of cooperating companies. Set all prices equal to the price of competitor above them minus one cent; ie. $b_k^{(p)*} = b_k^{(p-1)}-o$. This is referred to as "gap jamming" and is an example of an anti-social bidding method [8]. For each company i, move that company into its optimum price. Repeat this process until none of the companies change their price. At this point the Nash equilibrium will have been reached. Next find any companies which have positions that are adjacent. If such competitors are found, then repeat the following. Find the lowest priced cooperator p. Set that price to the price of one penny above the competitor below it; $b_k^{(p)*} = b_k^{(p+1)} + o$. Find the next lowest priced cooperator p-1. Reduce its price to one penny above the price below it. And so on. After these adjustments have been completed, all competitors will be in Nash equilibrium positions, but will be saving the maximum amount.

5.5 Example

Let there be three cooperating companies, A, B, and C, as well as one non-cooperating D. Prices for A,B,C,D are initialized as follows:

$0.23, $0.22, $0.21, $0.20

We successively optimize each of the company prices, with respect to the prices chosen by the other cooperating companies (they aren't cooperating right now, but they will later). Each company, A, B, C therefore repeatedly leapfrogs each other, steadily increasing the price. For instance, after one iteration, the prices for A,B,C,D may be:

$0.24,$0.23,$0.22,$0.20

Eventually, prices will be increased to a point at which no further profitable change is possible, for example

$1.92, $1.91, $1.90, $0.20

Due to the gapjammer rule, the third company at $1.90 keeps prices propped up. At this point we have found the Nash equilibrium prices for each competitor. Now we repeatedly shift the smallest bid cooperative company down in price as far as it will go. For instance, after one iteration we will have

$1.92, $1.91, $0.21, $0.20

After second iteration we have:

$1.92, $0.22, $0.21, $0.20

After a third iteration we will have

$0.23, $0.22, $0.21, $0.20

We have now arrived at a set of prices that are as cheap as possible but which are in the exact same positions that the companies achieved at their respective Nash equilibrium.

The theoretical cost savings from cooperative price setting are substantial. Company A and B are each saving $1.70 per click. If the number of clicks per month in positions 1..3 were (1000,800,500) respectively, then company A would save $1,700, company B would save $1,360 and company C would generate the same amount of profit as before.

6. EXPERIMENT

From November 7-14 2003 we ran a live auction with 265 listings divided amongst three accounts on a *Yahoo* auction. Our client company had requested the listings be divided between three distinct accounts, accounts 6, 7 and 8. The separation of these accounts introduced the problem of unwanted competition. As a result, we wanted to test the Cooperative Pricing Algorithm to eliminate unnecessary competition between "friendly" listings. From the 7th-14th November 2003 we placed bids in a purely competitive manner, and captured data on the auction.

59 keywords (22%) were bid on by two or more cooperating bidders. 33 of these were excluded because their profit models were not deemed to be reliable, for example because they had positive slope. Of the remaining 26 auctions, we found that 3 that had optimum prices that were identical. The remainder of our analysis is concerned with these 3 cases in which met the Theorem 1 requirements for cooperation. The three cooperative listings are shown in figure 4.

Auction	Account 6 optimum bid	Account 6 optimum position	Account 7 optimum bid	Account 7 optimum position
adult education	$0.71	15	$0.71	15
graduate degree	$0.51	5	$0.51	5
master degree	$2.97	3	$2.97	3

Figure 4: Three listings for which accounts 6 and 7 have the potential to coordinate their prices.

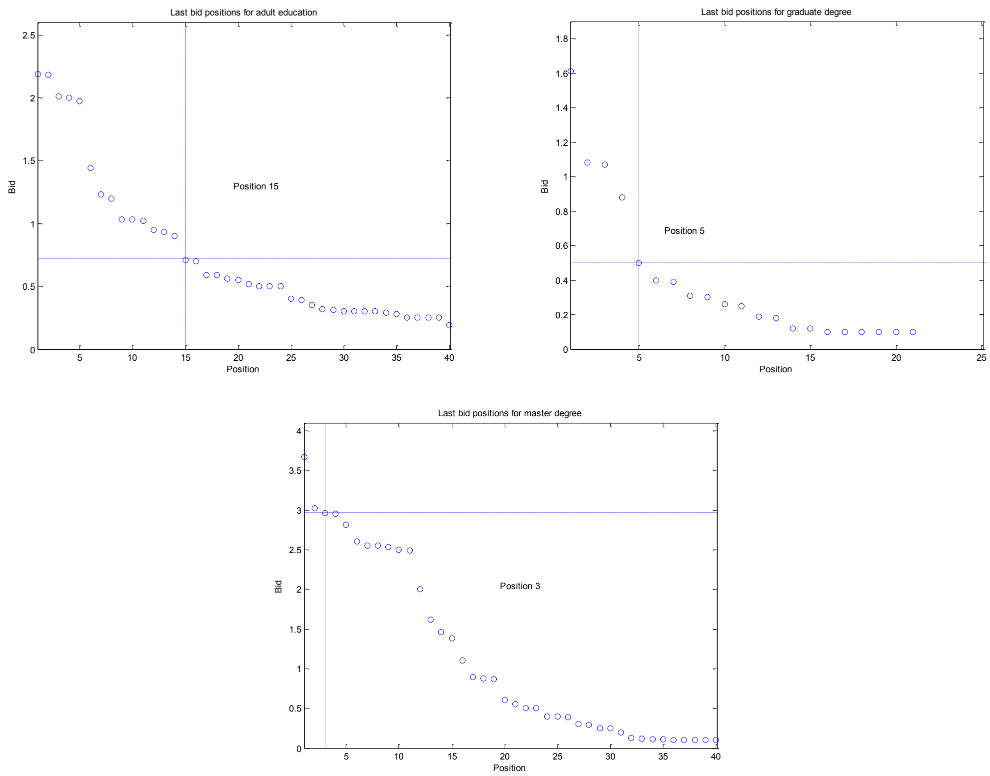

Figure 5: Price versus position for the three auctions, with optimal bidding prices for account 6 and 7 shown. Each circle is a position on the auction. The dotted cross-hairs are the position that is optimal. Out of 265 listings, only these auctions had identical optimal prices for account 6 and 7.

Figure 5 shows the auction prices for all three auctions, plotted against position. The optimal positions are indicated as a cross-hair. In the auction for "Adult education", both account 6 and 7 are seeking position 5, at a cost of \$0.51. This would normally spark a bidding war. Figure 6 shows the outcome. Account 6 would be prepared to escalate until a price of \$0.81. Account 7 would escalate until \$0.73. Thus, 6 would receive top position. Following the Cooperative Pricing Algorithm,, we drop account 7 to one cent above the competitor below, \$0.71, and account 6 to \$0.72. The positions are the same, but account 6 now saves \$0.02 per click. If we were to follow this strategy across all cooperating keywords, account 6 would save \$6.51, and 7 would save \$0.29 per month. (Figure 6)

SearchTerm	Account6 MaxPrice	Account7 MaxPrice	Price Below	Price Above	Account6 Coop price	Account7 Coop price
adult education	<u>0.81</u>	0.73	0.70	0.90	0.72	0.71
graduate degree	<u>0.64</u>	0.54	0.50	0.88	0.52	0.51
master degree	2.98	<u>2.99</u>	2.96	3.02	2.97	2.98

Account6 Nash price	Account7 Nash price	Account6 savings per click	Account7 savings per click	Account6 clicks per hour	Account7 clicks per hour	Account6 savings per hour	Account7 savings per hour	Account6 savings per month	Account7 savings per month
0.74	0.73	0.02	-	0.4234	0.153958	\$ 0.0085	-	\$ 6.30	-
0.55	0.54	0.03	-	0.0094	0.011338	\$ 0.0003	-	\$ 0.21	-
2.98	2.99	-	0.01	0.0947	0.039482	-	\$ 0.0004	-	\$ 0.29

	Total	$ 6.51	$ 0.29

Figure 6: Maximum prices, Nash prices, Cooperative prices for account 6 and 7. The maximum prices are the highest prices that account 6 and 7 would pay for the same position. Nash prices are the prices that would result after a bidding war in which both agents acted in their self-interest. Cooperative prices are the prices that could result if each party behaved so as not to exploit the other, with the object of increasing profits of both parties.

7. PREVALENCE OF ADVERTISER COLLUSION

Neither *Yahoo* nor *Google* make any mention of bidder collusion in their Terms Of Services Agreements [1][3]. Both auctions allow any bid price above a minimum bid. Therefore, will cooperative bidding become widespread?

Both auctions are growing rapidly as new advertisers enter the auction, and this leads to the problem of “cheaters” repeatedly disrupting coalitions by out-bidding them [5]. In addition PPC auctions are highly diversified with over 200,000 individual advertisers on *Google* at the end of 2004. These numbers make it more difficult for coalitions to grow large enough to significantly affect prices.

In the future these conditions may change however. One would expect the growth of advertisers to slow down and become indexed to growth in population. The difficulty of bidding on these auctions has also led to an increasing number of companies that out-source their bidding to advertising agencies. Obviously, agencies have no interest in increasing their own click prices (ie. hurting the clients they represent). Thus, a slowing in growth and an increase in concentration of bidders should lead to conditions that encourage collusion.

8. IS ADVERTISER COOPERATION IN THE PUBLIC INTEREST?

A second question is whether cooperation is ethical. Cooperative strategies obviously benefit advertisers at the expense of search engines, since they decrease the revenue captured by them. They also obscure the true values of different keywords. But is this necessarily a bad thing?

In the case of the FCC Spectrum auctions, collusion was a serious matter since the public was paying for the difference between the true value and the bid. Yet search engines are for-profit entities - the only entity being harmed is a private company, which is already making a large amount of money. Why should we worry if *Google* is making a little less money?

On the other hand, if keyword prices are allowed to migrate to their true valuations, it is less likely that poorly targeted advertising creative would be seen by internet users, since the razor-thin profit margins force advertisers to make their advertising selective. This is one of the most fascinating aspects of pay per click auctions - click prices force advertisers to “evolve” their advertisements so that they are seen by fewer and more targeted people. Every person that sees a pay per click advertisement, clicks on it, and does not convert, is recorded as “cost” in the profit equation for advertisers. Profit increases when cost decreases. Thus, for purely selfish, utility-maximizing reasons, pay per click bidders have economic incentives to refine the words they buy to a smaller and smaller number of customers who are actually interested in the product or service. This kind of targeting is achieved by (a) targeting words with greater specificity, such as 4 and 5-word collocations such as “Helicopter Tours Las Vegas under 100 dollars”, and (b) creating long lists of “negative match” keywords that will stop the advertisement being shown - for instance, “grand canyon”, “hoover dam” if we wanted to only find customers looking for helicopter tours around the city. Higher click prices will accelerate this trend, and make it even more costly to have badly targeted advertisements. Thus, elimination of collusion should benefit consumers.

We believe that eventually search engines will erect prohibitions against collusion, since this is in their interest. Such a prohibition would increase the quality of advertising creative, yet at the expense of making it more difficult for advertisers. If such an action is taken, methods should be considered for improving the accountability of search engines as well.

REFERENCES

[1] *Google AdWords Terms of Service Agreement*, http://www.google.com/apis/adwords/terms.html, Google Corporation, February 2005.

[2] *Google AdWords FAQ*: Account Setup, https://adwords.google.com/select/faq/account.html#8, May 2005.

[3] *Yahoo PrecisionMatch Terms of Use Agreement*, http://www.content.overture.com/d/USm/legal/tou.jhtml, Yahoo Corporation, February 2005.

[4] Agrawal, R. and Srikant, R. (1994) Fast algorithms for mining association rules, *Proceedings of the 20th International Conference on Very Large Databases*, Santiago, Chile.

[5] Axelrod, R (1984), *The Evolution of Cooperation*. Basic Books

[6] Bajari, P. and Summers, G. (2002), Detecting Collusion in Procurement Auctions: A Selective Survey of Recent Research, Working Papers in Economics 01-014, Stanford University.

[7] Blank, C. (2005), Study: Keyword Prices Get Cheaper After Holidays, *DMNews*, Feb 8, 2005. http://www.dmnews.com/cgi-bin/artprevbot.cgi?article_id=31793

[8] Brandt, F. and Weib, G. (2001), Vicious Strategies for Vickrey Auctions, *Proceedings of the Fifth International Conference on Autonomous Agents*, Montreal, Canada.

[9] Cramton, P. and Schwartz, J. (1999), Collusive Bidding: Lessons from the FCC Spectrum Auctions, Working Papers in Economics, University of Maryland.

[10] Elliot, N. and Scevak, N. (2004), Paid Search Through 2009: Rising Cost Per Click Drives Market, Forces Advertiser Efficiency, *Jupiter Research Report*, Jupitermedia Corporation, August 2004.

[11] Kitts, B., Laxminarayan, P., LeBlanc, B. and Meech, R. (2005), A Formal Analysis of Search Auctions Including Predictions on Bidding Tactics and Click Fraud, *First Workshop on Sponsored Search at the Sixth ACM Conference on E-Commerce*, Vancouver, BC.

[12] Kitts, B. (2005b), Being in the Right Time at the Right Place: The Effect of Time on Pay Per Click Auctions, *Search Engine Advisor Newsletter*, iProspect, http://www.iprospect.com/media/newsletter_january_kitts.htm

[13] Krishna, V. (2002), *Auction Theory*, Elsevier Science, CA.

[14] Kursad, A. (2004) Dynamics of Bidding in Search Engine Auctions: An Analytical Investigation. University of Alberta School of Business Working Paper.

[15] McGann, R. (2005), Valentine's Day Sparks Keyword Bidding Frenzy, *ClickZ Stats*, February 11, 2005, http://www.clickz.com/stats/sectors/retailing/article.php/3482701

[16] Morrissey, B. (2003), Online Ad Rebound Underway, *ClickZ Stats,* June 12 2003, http://www.clickz.com/stats/sectors/advertising/article.php/5941_2221201

[17] *IAB Internet Advertising Revenue Report*, 2003 Second Quarter and First Six-Month Results, Interactive Advertising Bureau, Price Waterhouse Coopers. http://www.iab.net/resources/ad_revenue.asp

[18] Rashtchy, S. and Avilio, J. (2003), The Golden Search, US Bancorp Piper Jaffray Equity Research Report, Minneapolis, MN, March. Report available upon request from www.piperjaffray.com

[19] Reynolds, K. (1996), Collusion and Tricks, Agorics web page, http://www.agorics.com/library/auctions/auctions11.html

[20] Osborne, M. (2003), *An Introduction to Game Theory*, Oxford University Press.

[21] Vise, D. (2004), Advertisers Bid, You Click, They Pay: Keyword Auction Nets Google Most of Sales, *Washington Post*, Thursday, May 13, 2004; Page E01 http://www.washingtonpost.com/ac2/wp-dyn/A22435-2004May12?language=printer