Experimental Evaluation of Sponsored Search Auction Mechanisms

Full Papers

Di Shang

Long Island University Brooklyn

Brooklyn, New York, United States

di.shang@liu.edu

Roumen Vragov

Mount Saint Mary College

Newburgh, New York, United States

rumvra@yahoo.com

Abstract

The theory of sponsored search has been developing rapidly although with disagreement in scientific circles on answers to some basic questions about sponsored search. This study focuses on two of these questions, namely, if a search engine seeks to maximize profits, 1) what should its pricing policy be and 2) what should its ranking policy be. This paper uses experiments with economically motivated human subjects to address these questions. We evaluate six different sponsored search auction formats with two different pricing policies (Pay-per-transaction & Pay-per-click) and three different ranking policies (Rank by relevance, Rank by click-through rate, & Rank by both relevance and click-through rate). Our results suggest that Pay-per-click is superior and the reason behind its superiority is behavioral in nature whereas the ranking policy has significant effect on search engine revenue and advertiser profit.

Introduction

Sponsored search advertising is a multi-billion dollar market; just Google reported ad revenues of $59 billion for 2014[1]. Revenue from sponsored search is a source of income not only for search engines but also for many other companies and individuals who partner with search engines to display advertisements on their own sites. The success and popularity of sponsored search as a new way to advertise and attract consumers has stimulated interesting theoretical debates in academia. There currently are regularly scheduled interdisciplinary workshops and conferences dedicated to issues related to sponsored search. Academic research has also been encouraged and supported by the main players in the sponsored search auction market: Google, Yahoo, Bing, and AOL among others.

The sponsored search auction market consists of three main players: the search engine, advertisers, and consumers. The search engine wields important power in this electronic market because it decides on the rules of the market mechanism and thus controls the way in which price and product information is conveyed from advertisers to consumers. Search engine managers possess an array of important strategic auction market variables that they can manipulate in pursuit of maximizing the search engine's profits. Two of the most important variables are the pricing policy and the ranking policy of the sponsored search mechanism.

The pricing policy describes how advertisers must pay to participate in the sponsored search auction. There are three main pricing policies that have established themselves in the online advertising community: pay-per-impression (PPI), pay-per-click (PPC), and pay-per-transaction (PPT). PPI is a relatively old policy that has been applied to traditional media: radio, TV, newspapers, magazines, posters, etc. According to this policy advertisers have to pay depending on the number of consumers that are expected to see the ad. PPI is based around the idea that the advertising media is generally not interactive, so there is no way to measure precisely how many consumers acted because they observed the advertising message. PPC & PPT have become popular together with the Internet and they are both based on the idea of interactive advertising. Under PPC advertisers pay only if a consumer clicks on the

---

[1] See http://investor.google.com/financial/tables.html

advertising message, and under PPT advertisers pay only if a consumer buys the advertised product. PPC & PPT are currently considered the two main options of choice among sponsored search auction managers.

The ranking policy describes how advertisers are ranked in the list of sponsored search links usually shown above (or to the right of) the organic search results. It makes sense that the first criterion according to which advertisers should be ordered is how much they are willing to pay as a fee to the search engine (or the advertisers' bids). However, because the search engine's revenue also depends on how many consumers actually use sponsored links, search engine managers deem it necessary to include some measure of consumer preference in the ranking policy. Thus there are three main ranking policies that have been used in sponsored search in addition to ranking by the advertiser's bid: rank by relevance (R), rank by click-through-rate (CTR), or rank by some combination of relevance and click-through rate (R x CTR). Relevance reflects how useful a link is to a consumer given the search terms that a consumer used. For example if a consumer searches for "refrigerator", the most relevant advertisers shown are the ones who bid on the keyword "refrigerator". This is rarely a perfect measure of the actual relevance of the sponsored links because often a consumer cannot express perfectly through keywords what s/he desires. In addition, relevance does not reflect how good a match is for the search consumer, i.e. a less relevant product might be relatively cheaper than a more relevant product. One way to measure indirectly how good a match is to consider how many times consumers themselves consider it a good match. This is where the click-through rate comes in. The actual click-through rate is the expected number of times that a sponsored link will be clicked by consumers during a certain period of time. This rate is also hard to predict but past click-through rates can be measured accurately and used as proxies.

So far theory has offered conflicting conclusions on the effect of these different policies on the profitability of the sponsored search auction mechanism. Some theoretical models favor PPC over PPT as the better pricing policy and others favor PPT over PPC. The same applies to the different ranking policies. Testing these alternative theoretical models using data from naturally occurring sponsored search auctions has also been extremely challenging because consumers' and advertisers' costs and preferences are strictly private and often unobservable, irrelevant effects are hard to filter out, and, since not all alternative policies or combination of policies have been implemented, some important and relevant data has never been generated. Under such circumstances laboratory experiments with economically-motivated human subjects can be used not only as a way to provide some needed practical guidelines to sponsored search mechanism designers but also as a way to evaluate the usefulness of several competing theories of sponsored search.

In this study we aim to provide experimental evidence to address the following specific research questions:

R1: Which payment policy (PPC or PPT) generates the highest revenue for the search engine?

R2: Which ranking policy (R, CTR, or RxCTR) generates the highest revenue for the search engine?

To answer these questions we design an experiment using the methodology of experimental economics (see Smith, 2003). We use a standard 2x3 factorial design in order to control for possible interaction effects between the pricing and ranking policies. The need of laboratory experiments in e-commerce research has been recognized and discussed in Kauffman and Wood (2008). Under their methodological framework, laboratory experiments allow the researcher to study relevant phenomena with complicating environmental context-related conditions removed. Such experiments enable precise measurements of relevant effects while necessarily sacrificing some generality and realism.

The methodology of experimental economics is suitable for our investigation especially because consumer and advertiser preferences and costs which are private and unobservable in practice can be induced and are directly available to the researcher in the laboratory (Smith, 1976). Through laboratory experiments we can relax many indubitably unrealistic assumptions made by most theoretical models of sponsored search in order to reduce complexity including but not limited to the presence of only one or two products being advertised, the absence of correlated product values, the presence of only homogeneous and rational search consumers and advertisers on the market. Laboratory experiments also allow us to gather data from pricing and ranking policy combinations that have never existed in practice.

Our results suggest that the search engine should use Pay-per-click as the pricing policy. The superiority of Pay-per-click is behavioral in nature that suggests that future models of sponsored search should incorporate behavioral variables. We do find significant profit effects due to ranking policy. Since there are some traces of dependencies, we suggest ways in which our experimental design can be extended to further probe the effect of ranking policies.

Literature Review and Hypothesis Development

Sponsored search has emerged as a topic of interest for many scientific disciplines including economics, information systems, computer science, and marketing. For an overview of sponsored search see Jensen & Mullen (2008); for a review of relevant economic theories and models see Athey & Ellison (2008) and Vragov (2009). Varian (2005) & Edelman et al. (2007) provide a summary of relevant mechanism design literature, and Szymanski & Lee (2006) provide a comparison of different sponsored search auction mechanisms (GFP, GSP, VCG). There are also some empirical studies, for example, for panel data study, Ghose and Yang (2008) show relationship between different sponsored search metrics such as click-through rates, conversion rates, cost-per-click, and ranking of advertisements using a 6 month panel dataset. Here we focus our literature review on several studies that have attempted to answer the questions we posed.

Hu et al. (2010) study performance-based pricing models in online advertising and finds that if click-through rate is measured with greater precision, pay per click payment is better off than pay per transaction because the method exposes both publishers and advertisers low risk in the presence of uncertainty of user click-through.

Other researchers have generated different conclusions. Blumrosen et al. (2008) conduct some interesting simulations using a theoretical model and observations from real advertising data. They find that pay per transaction is better than pay per click for the search engine. According to them the reason for this result is the non-uniform conversion rates among advertisers. Athey & Ellison (2008) argue that differences in method of payment for advertisers are related to the informative content of the displayed ads. This implies that in absence of informative content there should not be differences in surplus or profits that depend on the method of payment.

Dellarocas and Viswanathan (2008) provide a theoretical model of the interaction between advertisers, search engine, and consumers and find that both pay-per-transaction and pay-per-click lead to higher prices and decrease in pay-offs for all because advertisers are risk-averse but the authors do not rank the two methods in terms of efficiency. The only difference between the two pricing policies in the theoretical model discussed in Vragov (2009) is that a pure-strategy Nash equilibrium exists under pay-per-transaction and such does not exist under pay-per-click. A ranking of the two pricing policies is not provided.

Current research is also inconclusive about the effects of ranking mechanism on search engine revenue. Feng et al. (2007) suggest that performance of sponsored search mechanisms depends on the degree of correlation between advertisers' bid and their relevance to the search term, and that ranking based on click through rate and bid price performs well. They do not compare this with other ranking policies. Weber and Zheng (2007)'s study also reports that the revenue-maximizing search engine ranking policy is the one that uses a weighted average of click-through rate and bid amount. However, the authors also point out that the weighted ranking reduces overall consumer surplus compared to the socially optimal design of ranking purely on product performance. Using computational equilibrium analyses, Thompson and Leyton (2008) report that weighted pay-per-click auctions outperform un-weighted pay-per-click auctions because weighted pay-per-click auctions find the optimal allocation, in term of social welfare, with extremely high probability.

Athey & Ellison (2008) suggest that click through weighting of bids does not cause differences in surplus in the limit. Furthermore, when the number of firms is small, click-through weighting causes inefficiencies in the set of listed firms and also inefficiencies in the ordering of listed firms. In addition, rank by click-through rate may result in welfare losses when asymmetries in the click-through weights make the ordering of the ads less informative (relevant) about quality. Vragov (2009) argues that usage of past click-through rates for ranking decreases the competition in the market for ads which leads to lower revenue for the search engine.

The literature review above shows us that results on the efficiency of sponsored search mechanisms are often controversial, which indicate that results are not robust to small changes in assumptions or settings. We decided to use laboratory experiments with human subjects to bring theoretical models mentioned above closer to reality and to investigate our research questions for the first time with one more method of scientific exploration. Comparing results from different methodologies can provide a fuller picture of the complex environments we are studying. In the experimental design described below we make no assumptions about consumers and advertisers (they are randomly picked from the undergraduate population of a large urban business school) and we use a very general product space that contains both substitutes and complements. The strictest assumption we make is that the search engine has a perfect knowledge of ad relevance based on the search terms consumers type and that consumers know what to type in order to indicate how relevant an ad is to them. As discussed in the last section we plan to relax these assumptions as well in future designs.

The literature review above shows us that theoretical models of sponsored search differ in terms of their predictions concerning the effect of the various pricing and ranking policies on the revenue of the search engine. One can imagine that, given these contradictions, it would be hard to choose which models to use as a basis to derive testable hypotheses. In such cases and since we plan to use a standard ANOVA with repeated measurements, it seems sensible to hypothesize initially that no differences between policies exists. In fact hypothesizing that the pricing policy has no impact on the search engine's revenue makes a lot of sense. If an advertiser is willing to pay a certain amount of money to get a certain number of sales, then the advertiser will change his/her bid depending on the pricing policy. Since there are usually more clicks than transactions, the advertiser will bid less if s/he is charged per click and more if s/he is charged per transaction. Overall we would expect the advertiser's total expenditure in the sponsored search auction to be roughly the same in equilibrium, which should translate to similar revenues for the search engine. We should note, however, that at least until advertisers discover the correct conversion rate, pricing per click is riskier for advertisers. Advertisers always know how much they make per transaction because they are setting the product price and thus know what their bidding limit is. Finding out how many clicks on average they need before a transaction happens can take time during which negative profits could occur if advertisers mistakenly overbid. This, however, should not have any noticeable effect on the search engine revenue. Thus our first main hypothesis is the following:

H1: The pricing policy has no impact on the search engine's revenue from sponsored search auctions.

We can build a similar argument for the effect of the ranking policy. Both ranking by relevance and ranking by (past) click-through rate are ways to make the list of sponsored links match the consumer preferences better. The higher the relevance, the more likely it is that consumers are willing to buy the product that is being advertised. Both methods are, however, imperfect. Advertisers determine the relevance by choosing the keywords or combinations of keywords they bid on. On the other hand, if the (past) click-through rate is chosen as a proxy for relevance, then the consumers themselves are the ones that determine how useful a link is to them by the act of clicking on it. Note that a consumer's gain from buying a product depends not only on relevance but also on the product's price. A less relevant product might be listed at a much lower price and might turn out to be a better deal than a more relevant but disproportionately more expensive product. Thus it seems that ranking ads by (past) click-through rate should produce a better overall match between consumers and products than ranking by relevance because it is reasonable to assume that consumers know their own preferences better than either the search engine or the advertiser. Ranking by a combination between relevance and click-through rate should be even better for the market because information from all market participants is aggregated when the ranking is performed. Nevertheless, it is not clear which market participant will reap this gain in market efficiency. Gains will accrue to consumers and advertisers for sure but whether the search engine's revenue will be affected is still an open question. Thus we can formulate our second hypothesis in the following way:

H2: The search engine's revenue from sponsored search auctions is not affected by ranking policy.

Experiment Design and Procedure

Experiment Design

The experiment has a standard 2x3 design, and each session consists of 30 rounds (Table 1).

	Rank by relevance (R)	Rank by click-through (CTR)	Rank by relevance x click-through (R x CTR)
Pay per transaction (PPT)	Treatment 1 5 sessions 60 subjects	Treatment 3 5 sessions 60 subjects	Treatment 5 5 sessions 60 subjects
Pay per click (PPC)	Treatment 2 5 sessions 60 subjects	Treatment 4 5 sessions 60 subjects	Treatment 6 5 sessions 60 subjects

Table 1. Our standard 2x3 experimental design

Our experimental environment is symmetric because it consists of 6 search consumers and 6 advertisers. Clearly using only one or two consumers/advertisers creates problems for markets because of concentration of market power. We chose to use six advertisers to avoid this problem since past experiments have demonstrated that this power tends to disappear when there are more than five market participants. We also decided not use more than six advertisers because the market environment would have become too complex given that we are simulating substitute/complement relationship between the different products. Complex environments require much more dedication by experimental subjects, whose attention is increasingly hard to maintain after 2 hours in the laboratory. Then we fixed the number of consumers to six as well again to avoid concentration of market power and also to keep the environment symmetric.

Every advertiser produces a unique product and every advertiser can sell up to 6 units of the product at constant marginal cost of production. This was done to insure that each advertiser has enough production potential to satisfy all consumers. There are two product categories: A and B. There are three products in category A: A1, A2, A3, and three products in category B: B1, B2, and B3. Each product is uniquely assigned to be produced by only one advertiser. We drew three random numbers between 0 and 100 and rounded them to the closest 5 to determine the cost of products A1, A2, and A3. We then drew another 3 random numbers between 0 and 50 to determine the costs of the remaining products in descending order. The resulting set of experimental product costs is shown in Table 2. The costs are shown in experimental dollars.

Advertiser	Product	Cost per unit	Production Capacity
1	A1	55	6
2	A2	50	6
3	A3	40	6
4	B1	35	6
5	B2	25	6
6	B3	45	6

Table 2. Supply features

On the demand side we have 6 consumers with combinatorial preferences. We chose a random number from the interval [100, 200] and rounded it to the closest 5 to get the maximum product value for a consumer. We then drew a random number between [0, 5] to determine the demand slope of one group of consumers. We also introduced two different click costs. Consumers were divided in three equal groups (2 subjects each) depending on which product from Category A they like best. The purpose of this was to create variant relevance for the three products in category A among the consumers. The values for the products in Category B are the same for all consumers. The bonus that consumers receive when they buy one product from Category A and one product from Category B is also the same for all subjects. The values and costs for each consumer are shown in Table 3.

	Group 1		Group 2		Group 3	
Product	Consumer 1	Consumer 2	Consumer 3	Consumer 4	Consumer 5	Consumer 6
A1	120	120	110	110	90	70
A2	100	100	120	120	105	95
A3	80	80	115	100	120	120
B1	65	65	65	65	65	65
B2	80	80	80	80	80	80
B3	75	75	75	75	75	75
Bonus	30	30	30	30	30	30
Click Cost	5	15	5	5	5	15

Table 3. Consumers' values, bonus, and click costs.

Because of the way values are assigned to consumers according to Table 3, consumers view products in Category A as imperfect substitutes within their own categories and products in Category B as complements in between the two different categories.

Experimental procedure

The experiment proceeds in the following fashion. First subjects are recruited from the undergraduate student population of a large urban business school. When they arrive at the laboratory, subjects are randomly assigned to be consumers or advertisers in the experiment. Then subjects are assigned to a computer terminal and given computerized instructions. After reading the instructions subjects sign the informed consent form and participate in a practice round, which lasts 4 minutes. Then the actual experimental session starts. The experimental sessions lasts usually around 70-85 minutes. Subjects can ask questions any time during the experimental session. At the end subjects are paid in private the earnings that they received during the experimental session converted into US dollars plus a $10 show up fee. The average subject earnings were approximately $65.00.

When an experimental round starts first it is the advertisers' turn to make their decisions. They have to decide how much to charge consumers who wish to buy their product and how much to pay to the search engine for being displayed on the consumers' screen. Advertisers bid for exposure in our experimental design. The three advertisers in Category B each have to submit a price for their product and a bid to the search engine. The three advertisers in category A have more options. Since consumers differ in terms of their preferences for products in Category A, advertisers are allowed to submit bids separately for consumers who prefer A1, consumers who prefer A2, and consumers who prefer A3. Thus advertisers who produce products in Category A can submit a price for their product and up to three bids: one bid for product A1, A2, and A3. If they submit the same bid for all products in category A, this means that they are bidding on one more general search phrase that describes all products in category A. If the three bids they submit are different, then advertisers chose to bid on three more specific search phrases that more specifically describe each of the products in the A category.

After advertisers have made their decisions, the search engine collects all bids and prices and decides separately which products to display to consumers who prefer respectively A1, A2, and A3 and also in what order to display them. The search engine picks the advertisers with the top 4 bids for each of the three product groups and displays their products to the consumers.

The order in which the top 4 bids are shown to each consumer varies by treatment. The display order is the second treatment variable. The search engine always picks to show the products with the highest 4 bids, however, the ordering of the products is different. The products are ranked and displayed by

Relevance (R), by past Click-through Rate (CTR) or by the product of the two (R x CTR). Ties in rank or bids (for advertisers) are broken at random.

When the ranking is done, consumers can proceed to make their decisions. Each consumer can see that four of the products are available for sale. A consumer can click on a product to check its price, which results in a click cost as shown in Table 2. After checking some or all of the prices, the consumer can proceed to make purchase decisions. If a consumer decides to purchase a product, the consumer receives the value of the product from Table 2 as revenue and has to pay the price that the advertiser of that product indicated.

After all consumers make their decisions, the experimental software displays the round results to the advertisers. The way in which advertiser profit is calculated varies by treatment. We provide some examples in Appendix to explain the ranking of bids and how profit is calculated.

RESULTS

All the results reported below are based on a standard two-factor ANOVA with replication and 5 statistically independent observations in each of the 6 cells. If the optimal allocation is implemented, the total surplus is 840 experimental dollars per round or 25,200 per session. The average efficiency attained during the experiment was 75% (630 experimental dollars per round or 18,900 per session).

Our hypothesis of no effect of pricing policy on search engine revenue (H1) cannot be rejected at the significance level of 0.05 with p-value equals to 0.126. However, ranking policy shows significant impact on search engine revenue, and our hypothesis of no effect of ranking policy on search engine revenue (H2) is rejected at the significant level of 0.05 with p-value less than 0.03. For search engine, rank by the combination of Click-through and Relevance generates the highest revenue because this ranking method results in the highest average bidding prices from advertisers (Figure 1).

There currently are several search engines in the market for sponsored search with Google having the largest market share. When deciding on the appropriate pricing and ranking policy, a search engine cannot focus solely on maximizing its own revenue but also has to evaluate the effect of its policies on its users: the advertisers and the consumers. Policies which maximize the search engine's revenue but hurt consumers and advertisers might not be viable in practice because of competition among search engines. Therefore, we are also interested in examining whether it hurts search consumers and/or advertisers if a search engine focuses solely on maximizing its revenue given the available alternatives.

From the results, we see that there is indeed conflict of interest of market participants. Prices and bids grow faster and are higher in PPT than in PPC. Product prices (p-value = 0.047) and advertiser bids (p-value = 0.040) are significant higher in PPT if we consider only the last 10 rounds of the experiment. Consumer surplus is also clearly better under PPC than under PPT if we consider only the last 10 rounds (p-value = 0.049). We see that there are much fewer advertisers in pay per transaction treatment (only two of 2700 advertiser/round) who incur negative profit than in pay per click treatment (431 times out of 2700 advertiser/round). However, the average total advertiser profits are not significantly different between the two payment mechanisms. Total market surplus, which is measured by the sum of search engine revenue, advertiser profits, and consumer surplus is significantly higher under PPC than under PPT (p-value = 0.018). This was caused by a significant increase in market trading volume (p-value = 0.023) mostly due to increasing purchases of products A3 (p-value = 0.031) and B2 (p-value = 0.017), which have the lowest production cost in their corresponding categories.

For consumers, though no significant difference of profit, rank by Relevance appears to be a good choice for them because of the lowest product price and the highest number of transaction. Advertisers obtain highest profit in rank by Click-through rate. However, consumers in this treatment incur highest search cost (total click cost).
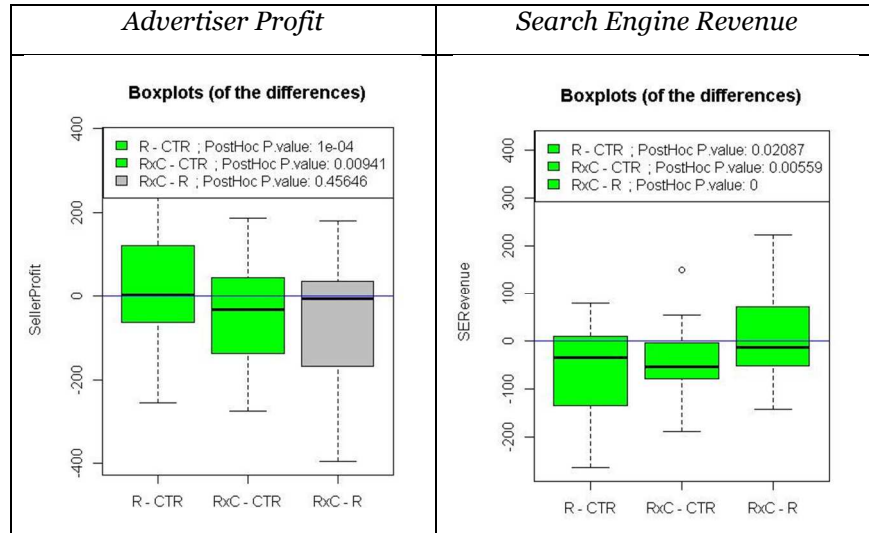
Figure 1 Advertiser Profit and Search Engine Revenue cross Treatments

Conclusion

Our results suggest that the superiority of the pay-per-click policy might be behavioral in nature. When advertisers pay per transaction they can rarely incur losses because they know quite well their per unit revenue in advance except for some complementarity among products, which could only increase this per unit revenue. Under PPT advertisers are much more aggressive in their bidding for the top slots in the list of sponsored links as indicated by the higher bids. More aggressive bidding leads to bid wars among advertisers that cause higher prices for consumers and worse product matches. The same does not happen under the PPC policy. This policy has an in-built uncertainty about the conversion rate or the expected per-unit revenue. Since it is easy to end up in a negative territory, advertisers are much more careful in their bidding and bidding wars similar to these under PPT do not occur. Both prices and bids under PPC seem to stop increasing and converge to a more stable level towards the end of the session.

The main conclusion from our results so far is that the sponsored search mechanism performs better in general when advertisers are charged by pay-per-click. Search engine revenue is affected by ranking policy, and there is apparent conflict of interest between the search engine and advertisers and between search engine and consumers. We have successfully established an experimental baseline and a design that can be used in future experiments, and we have provided directions on how some of the results obtained here could be bolstered or refuted. We have shown that none of the theoretical models mentioned in the literature review expected the important behavioral effect of more aggressive bidding under PPT than under PPC. Whether this effect disappears in the long run or with experience is an open question. Some of the issues we raise here can be tested using data from real sponsored search auctions. Especially beneficial could be a comparison of Google's AdWords program in which advertisers pay per click and Bing's ad program, in which member advertisers pay per transaction. We hope that our study has also provided some guidance to practitioners and policy-makers regarding the effect of two important policies on search engine revenue and social welfare in general.

REFERENCES

Athey, S. and Ellison, G. (2011). "Position Auctions with Consumer Search," *The Quarterly Journal of Economics* (126: 3), pp. 1213-1270.

Bhargava, H. K. and Feng, J. (2008). A Model of Sponsored Results in Intelligent Recommenders and Search Engines. UC Davis Graduate School of Management Research Paper No. 04-08. Available at SSRN: http://ssrn.com/abstract=1264018.

Blumrosen, L, Hartline, J. and Nong, S. (2008). "Position Auctions and Non-uniform Conversion Rates," *Workshop on Ad Auctions '08,* Chicago, Illinois, USA.

Chen, J., Liu, D. and Whinston, A.B. (2009). "Auctioning Keywords in Online Search," *Journal of Marketing* (73), pp. 125-141.

Dellarocas, C. and Visawanathan, S. (2008). The Holy Grail of Advertising? Allocative Efficiency and Revenue Implications of 'Pay-per-Action' Advertising in Environments with Quality Uncertainty, working paper, University of Maryland.

Dhar, V. and Ghose, A. (2010). "Sponsored Search and Market Efficiency," *Information Systems Research* (21:4), pp. 760-722.

Dickinger, A. and Zorn, S. (2007). "Compensation Models for Interactive Advertising," in *Proceedings of 5th International Conference on Knowledge Management (I-KNOW),* Graz, Austria, pp. 101-108.

Edelman, B. and Ostrovsky, M. (2007). "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review* (97:1), pp. 242-259.

Edelman, B. and Ostrovsky, M. (2007). "Strategic bidder behavior in sponsored search auctions," *Decision Support Systems* (43:1), pp. 192-198.

Fan, W., Gordon, M.D. and Pathak, P. (2005). "Genetic Programming-Based Discovery of Ranking Functions for Effective Web Search," *Journal of Management Information Systems (*21:4), pp. 37-56.

Feng, J., Bhargava, H. K. and Pennock, D.M. (2007). "Implementing Sponsored Search in Web Search Engines: Computational Evaluation of Alternative Mechanisms," *INFORMS Journal on Computing (*19:1), pp. 137-148.

Ghose, A. and Yang, S. (2009). "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets," *Management Science* (55:10), pp. 1605-1622.

Goldfarb, A. and Tucker, C. (2008). "Economic and Business Dimensions: Search Engine Advertising," *Communication of the ACM* (51:11), pp 22-24.

Hu, Y.J., Shin, J. and Tang, Z. (2010). "Pricing of Online Advertising: Cost-per-Click-through vs. Cost-per-Action," in *Proceedings of the 43rd Hawaii International Conference on System Sciences*, 5-8 January 2010, Koloa, Kauai, HI, USA.

Jansen, B. J. and Resnick, M. (2006). "An examination of searcher's perceptions of nonsponsored and sponsored links during ecommerce Web searching," *Journal of the American Society for Information Science & Technology (*57:14), pp. 1949-1961.

Jansen, B. J., Brown, A. and Resnick, M. (2007). "Factors relating to the decision to click on a sponsored link," *Decision Support Systems* (44:1), pp. 46-59.

Kumar, S., Dawande, M. and Mookerjee, V.S. (2007). "Optimal Scheduling and Placement of Internet Banner Advertisements," *IEEE Transactions on Knowledge and Data Engineering* (19:11), pp. 1571-1584.

Kuruzovich, J., Viswanathan, S., Agarwal, R., Gosain, S., and Weitzman, S. (2008). "Marketspace or Marketplace? Online Information Search and Channel Outcomes in Auto Retailing," *Information Systems Research* (19:2), pp. 182-201.

Lamberti, F., Sanna, A. and Demartini, C. (2009). "A Relation-Based Page Rank Algorithm for Semantic Web Search Engines," *IEEE Transactions on Knowledge & Data Engineering* (21:1), pp. 123-136.

Smith, V. L. (1976). "Experimental Economics: Induced Value Theory," *American Economic Review* (66:2), pp. 274-279.

Smith, V. L. (2003). "Constructivist and Ecological Rationality in Economics," *American Economic Review* (93:3), pp. 465-508.

Storey, V. C., Burton-Jones, A., Sugumaran, V. and Purao, S. (2008). "CONQUER: A Methodology for Context-Aware Query Processing on the World Wide Web," *Information Systems Research* (19:1), pp. 3-25.

Thompson, D. and Leyton-Brown, K. (2008). "Tractable Computational Methods for Finding Nash Equilibria of Perfect-Information Position Auctions," *Workshop on Ad Auctions '08,* Chicago, Illinois USA

Varian, H.R. (2009). "Online Ad Auctions," *American Economic Review* (99:2), pp. 430-434.

Vragov, R. (2009). "Sponsored Search as a Strategic E-Service," *Journal of E-Services and Mobile Applications* (1:1), pp. 21-37.

Vragov, R. and Levine, I. (2007). "A Pricing Algorithm for Mixtures of TV Programming and Commercials," *IEEE Broadcast Technology Society Newsletter*.

Weber, T.A. and Zheng, Z.E. (2007). "A Model of Search Intermediaries and Paid Referrals," *Information Systems Research* (18:4), pp. 414-436.

Appendix: Examples of Bids, Ranking and Profit Calculation

Example: Given the advertisers' bids displayed in Table 3 we will demonstrate how the search engine decides which products to show to the various groups of consumers. Let us first run the procedure for Group 1 (Consumer 1 and Consumer 2). Consumers 1 and 2 prefer A1. We take the bids from advertisers 4, 5, and 6 and also the bids for A1 from advertisers 1, 2, and 3 and align them from highest to lowest. The resulting array of bids is shown in Table A1, column 2. A similar procedure is performed for the remaining two groups of consumers (see columns 3 and 4).

Advertiser	Bid for Consumers in Group	Bids	Prices
1	1	11	90
	2	5	
	3	1	
2	1	3	105
	2	7	
	3	9	
3	1	4	99
	2	7	
	3	8	
4	Same for all	10	60
5	Same for all	8	62
6	Same for all	3	53

Table A1. Example of advertisers' bids and prices

Example: Suppose that products are displayed in order of their relevance (R) to the consumers. For Consumer 1 the top four products to be shown are A1, B1, B2, A3 (see Table 4, column 2). These products will be displayed from top to bottom in the order of their value to Consumer 1, which is A1 first (value - 120), then A3 (value - 80), then B2 (value - 80), then B1 (value - 65) (See Table A2).

	Consumers in Group 1	**Consumers in Group 2**	**Consumers in Group 3**
Shown	A1 – 11	B1 – 10	B1 – 10
Shown	B1 – 10	B2 – 8	A2 – 9
Shown	B2 – 8	A2 – 7	A3 – 8
Shown (Lowest Accepted Bid)	A3 – 4	A3 – 7	B2 – 8
Not shown	B3 – 3	A1 – 5	B3 – 3
Not shown	A2 – 3	B3 – 3	A1 – 1

Table A2. Ranking of bids from Table A1.

Example: Suppose that Consumer 1 checks the prices of three of the four products displayed to her: A1, B1, A3 (see Table A2, column 2). The prices of these products are respectively: 90, 60, 99 as submitted by the advertisers (see Table A1, column 5). Given these prices, Consumer 1 decides to buy A1 for 90 and B1 for 60. Using the values and costs in Table 3, we can calculate Consumer 1's profit in this round. It is equal to 120 + 65 + 30 − 90 − 60 − 3 x 5 = 50. Note that Consumer 1 receives a bonus of 30 because she bought

a product from category A and a product from category B. Consumer 1 also incurs a click cost of 15 because she clicked three times to check the prices of three products.

Search engine revenue is calculated differently depending on the treatment. If advertisers pay per transaction, then search engine revenue is just the sum of the corresponding lowest accepted bids for all transactions. If advertisers pay per click, then search engine revenue is the sum of the corresponding lowest accepted bids for all clicks.