

Aggregators and Contextual Effects in Search Ad Markets

Asela Gunawardana
Microsoft Research
One Microsoft Way
Redmond, WA 98052, U.S.A.
aselag@microsoft.com

Christopher Meek
Microsoft Research
One Microsoft Way
Redmond, WA 98052, U.S.A.
meek@microsoft.com

ABSTRACT

Allocation and pricing of ad slots in search ad markets is typically done through a Generalized Second Price (GSP) auction, which assumes that the click-through rate (CTR) of an ad displayed in a particular position depends only on the identity of the ad and the position it is displayed in. In particular, it is assumed that there are no contextual effects, where the CTR of an ad depends on the other ads displayed with it.

We argue that such effects do exist. In particular, we discuss how the economics of ad aggregation may lead to ad aggregators imposing a negative contextual effect on the CTR of advertisers listed below them. We perform an experiment where we intervene in the auction to vary the position of selected aggregators relative to other advertisers. We then measure CTR on ads controlling for ad identity and position to quantify the magnitude of the contextual effect of being listed below an aggregator. In addition we describe the result of the intervention on search engine revenue, which is consistent with an increase in advertiser social welfare.

Categories and Subject Descriptors

H.4.0 [Information Systems Applications]: General; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Economics, Human Factors

Keywords

Search Advertising, Ad Aggregation, Contextual Effects

1. INTRODUCTION

Search advertising is a massive industry, with Google alone reporting a revenue of over US\$10 billion in 2007 from its websites [8]. In search advertising, advertisers bid to have short text ads shown on the search engine results page (SERP) for a particular query, and clicked by search engine users. In the dominant pay-per-click (PPC) model, advertisers only pay the search engine when their ad is clicked. The ads that are displayed on the SERP and their ranking are computed using an auction as follows. The advertisers submit bids representing the most they are willing to pay per click.

Copyright is held by the author/owner(s).
WWW2008, April 21–25, 2008, Beijing, China.

The design of this auction assumes that the click-through rate (CTR) of an ad can be factored into a component that depends only on the position of the ad on the page, and another component that depends only on the ad itself. Under this assumption, the ads are selected and ranked by sorting in decreasing order of their corresponding bid times the ad-dependent component of CTR. Under a suitable pricing rule, this ranking can be shown to maximize advertiser social welfare [2, 4, 10].

However, if the CTR of an ad does not factorize as assumed above, this is no longer true [2]. In particular, if CTR depends only on ad identity and position, but does not factorize, finding the optimal allocation could not be done by simply sorting. Instead, it would require the solution of a weighted bipartite matching problem, for which efficient polynomial-time algorithms exist [6]. In contrast, if CTR on ads depends significantly on *contextual effects*, i.e. what other ads are displayed in what positions on the page, the search for the optimal assignment may need to be carried out over all sets of assignments, which is exponential in the number of positions available on the page. This could be computationally expensive in practice, since even though the set of positions in which ads can be displayed is typically small (less than ten), the set of ads that can be displayed is quite large (in the thousands, if not greater).

Thus, allocation by ranking by bid times ad CTR may be suboptimal in the presence of contextual effects. The contributions of this paper are to suggest a mechanism by which such effects may arise, and to provide empirical results that quantify how large these contextual effects can be. In particular, we argue that ad aggregators can have significant contextual effects on the CTR of other advertisers.

The economic incentives of aggregators are different from those of merchants. While merchants have an incentive to provide goods and services that users arguably value, ad aggregators have an incentive to attract users and induce them to click on as many ad links as possible, *whether this provides user value or not*. Thus, the behavior of a user after viewing or clicking on such an aggregator ad is likely to be different from her behavior after clicking on a merchant ad.

In order to measure and quantify the contextual effects that aggregator ads impose on the CTR of other advertisers, we intervene in a real-life search ad market to vary the selection and ranking of the ads displayed. This allows us to control for position and advertiser identity when measuring the contextual effect aggregator ads have on other ads. To do this, we first identify several aggregators by analyzing user behavior on advertiser landing pages. We give

empirical results that show that advertisers that most induce users to click on links that redirect to other sites are ad aggregators. These redirections are a measure of how often an advertiser induces users to click on ads, as opposed to selling them goods or services. Thus, we hypothesize that these aggregators are likely to trigger different user behavior than merchants, and therefore cause contextual effects on the CTRs of advertisers they share the SERP with.

We select several of these aggregators and intervene in the auctions for 1000 randomly chosen search queries that they compete for in a real-life search ad market, thereby varying their position in the SERPs returned for these keywords. We build a model to predict the CTR of ads based only on search query, position, and advertiser identity in impression and click data collected both before and after the intervention. We then build a new model that attempts to reduce the remaining uncertainty in CTR by considering whether or not an aggregator ad appeared above the ad in question. This analysis shows that appearing below an aggregator ad lowers CTR by about 10%. This suggests that ranking by bid times ad CTR is suboptimal. We show that search engine revenue increases when aggregator position is depressed, which is also consistent with the initial ranking being suboptimal.

The work of Abrams and Schwarz [1] examines the design of an auction that takes the per-click “hidden costs” an advertiser’s choice of landing page imposes on the search engine. This work assumes that such a hidden cost imposed by an ad exists, that it is known to the search engine, and that it depends only on the ad. In contrast, the effects we describe in the present paper are contextual, and depend not only on the ad imposing the cost on the market, but on the other ads displayed, and their positions. The related work of Engel and Chickering [5] examine the design of an auction that optimizes a publisher utility function that depends on the user’s utility from clicking on an ad and the short-term revenue from advertisers. Once again, these results do not apply in the presence of the contextual effects we describe. The work of Becker *et al* [3] shows that an unfactored CTR model predicts actual click-through better than a factored model, but was unable to show a substantial contextual effect. We suspect that this is due to the lack of variability in the positions that advertisers are displayed in, because of the stability of their bids and CTRs. In the present paper, we use a treatment to induce such variability, and are able to show substantial contextual effects.

2. AD AUCTIONS AND CONTEXTUAL EFFECTS

In this section, we review the ad selection problem and how the dependency of ad CTR on context affects the difficulty of assigning ad slots optimally.

The problem of ad auction design is to decide which of N ads $i = 1, \dots, N$ to display in K available slots $j = 1, \dots, K$ on a SERP. For simplicity, we do not address the question of how many ads should be displayed on each SERP. The advertisers submit bids b_i indicating the maximum amount they are willing to pay per click. We assume an advertiser’s value per click is independent of the position in which the ad was displayed and what other ads were displayed and clicked on on the SERP. An auction can maximize the total value (social welfare) delivered to advertisers if it provides incentives for advertisers to bid truthfully, revealing how

much they each value a click at, and if it allocates the slots to the ads appropriately.

An allocation of slots to ads is an injection $\mathbf{a} : 1, \dots, K \rightarrow 1, \dots, N$ from slots to ads. We note that there are $\frac{N!}{(N-K)!}$ such allocations. The vector of binary variables indicating whether the ad in each of the slots $j = 1, \dots, K$ was clicked will be represented by c . The value of a SERP to advertisers is given by

$$\begin{aligned} V(\mathbf{a}) &= \sum_c p(c|\mathbf{a}) \sum_{j=1}^K c_j b_{\mathbf{a}(j)} \\ &= \sum_{j=1}^K b_{\mathbf{a}(j)} p(c_j = 1|\mathbf{a}) \end{aligned}$$

when bids are truthful. Choosing the assignment to maximize the social welfare of advertisers and then setting prices according to a Vickrey-Clark-Groves (VCG) mechanism incentivizes the advertisers to bid truthfully.

It is usually assumed that $p(c_j = 1|\mathbf{a})$ satisfies

$$p(c_j = 1|\mathbf{a}) = f_{\mathbf{a}(j),j} \quad (1)$$

$$= g_{\mathbf{a}(j)} h_j. \quad (2)$$

That is, it is assumed that the probability that the j th ad on the page will be clicked depends only on the position j and the identity $\mathbf{a}(j)$ of the ad assigned to that position, and that it factorizes into a component that depends only on the position and a component that depends only on the identity of the ad. In this case, the optimal assignment is to set $\mathbf{a}(1), \dots, \mathbf{a}(K)$ to the top K ads in decreasing order of $b_i g_i$ [4, 10].

In the case that the conditional independence assumption (1) holds but the separability assumption (2) does not, the total value to advertisers can be rewritten as

$$V(\mathbf{a}) = \sum_{i=1}^N \sum_{j=1}^K 1(\mathbf{a}(j) = i) b_i f_{i,j}$$

which can be recognized as an instance of the standard assignment or weighted bipartite matching problem. This implies that in this case, finding the assignment that maximizes advertiser social welfare can be done using efficient polynomial time algorithms [6].

Since in the general case there are an exponential number of allocations, the search for optimal allocations can be computationally expensive when dependencies which violate the conditional independence assumption (1) exist. We refer to such dependencies as contextual effects, since in these cases the probability of a click on an ad depends on its context in addition to its position and its identity. While we only discuss contextual effects of other ads shown on the SERP, we note in passing that other contextual effects such as dependencies on the organic search results and on the user’s identity and his web browsing history may also apply.

3. AD AGGREGATORS

In this section, we describe the economics of ad aggregation, and how this can lead to contextual effects that violate the conditional independence assumption (1). Ad aggregators are businesses who derive revenue from users clicking on syndicated online ads on their web pages. These syndicated ads are provided by an ad network, which charges

advertisers for clicks on their ads. Google’s adSense network is an example of such a network. When a user clicks on syndicated ads on an aggregator’s site, the ad network passes on a portion of the advertiser’s payment to the aggregator. Ad aggregators typically draw traffic to their sites by bidding to display their own pay-per-click ads on search engines. They are profitable when their per-click payments to the search engine are less than their revenue from syndicated ad clicks on their site. Shopping aggregators are a variant of ad aggregators that display price comparisons for goods and services. This information is often provided through a feed similar to that of a syndicated PPC ad network, and the aggregator is typically paid for clicks on these listings.

Since aggregators earn revenue from clicks, their revenue is not directly dependent on providing user value. This is in contrast to merchants, whose revenue depends on selling goods and services to users who presumably value the goods and services they buy. Thus, aggregators have an incentive to induce users to click on their syndicated ads through various means such as

- Providing price comparisons in the case of shopping aggregators.
- Popups—clicking on a syndicated ad on the aggregator page causes the corresponding advertiser’s landing page to open in a popup window on top of the aggregator page. Closing this window causes the user to see the syndicated ads again, perhaps leading to another click.
- Redirect traps—the aggregator’s landing page silently redirects to another aggregator page, usually via an HTML Meta Refresh tag. The timing of the redirect makes it difficult for the user to use the browser’s back button to return to the SERP.
- Deceptive UI. E.g. making the aggregator page look like the results page of a search engine so that users are unaware that they are clicking on syndicated ads.

Many of these lower the probability that users return to the SERP after visiting an aggregator page, since aggregators attempt to induce the user to remain on the aggregator page and click on syndicated ads instead of returning to the SERP. In addition, these techniques can frustrate users, making it more likely that users abandon their search, further reducing the probability that users return to the SERP. This in turn reduces the probability that the user will click on other ads on the SERP. We argue that ads that appear below aggregators are less likely to be clicked on than those that appear below merchants. This means that the CTR of an ad can depend on its context—in particular, on whether it appears below an aggregator or not. In the remainder of the paper, we attempt to measure and quantify this effect.

4. MEASURING CONTEXTUAL EFFECTS

In order to quantify the contextual effect caused by aggregator ads, we identify a set of aggregators that are likely to cause such an effect, intervene in the auction to vary their position, and then measure this effect. It is necessary to intervene in the auction because advertisers tend to occupy fairly stable positions in the auctions they participate in. This makes it difficult to tease apart position, advertiser,

and contextual effects in the measured click-through rates. In Section 4.1 we describe how we detect aggregators. We then describe how we intervened in a real search ad market in Section 4.2. In Section 4.3 we describe how we quantified the contextual effects of ad aggregators on CTR, and then report on the effect of the intervention on search engine revenue in Section 4.4

4.1 Detecting Ad Aggregators

In order to measure and quantify the contextual effects caused by aggregator ads, we must first identify a set of aggregator ads to study. In this section, we summarize a technique [9] for detecting aggregators, and give an empirical measure that is likely to be correlated with how large a contextual effect they cause.

As described above, aggregators derive revenue from clicks on syndicated ads. Since their payment comes from the ad syndication network, they cannot link directly to advertisers. Instead, they need to link to the ad network, which then redirects the click to the advertiser, so that clicks can be accounted and billed for. This means that clicks on syndicated ads can be distinguished from other clicks, since clicks on syndicated ads redirect through an ad syndication network. We will refer to clicks that are redirected through a third-party as redirections.

We propose the use of the ratio of the number of redirections from an advertiser’s page to the number of incoming clicks that the advertiser bought from the search engine to generate those redirections as a measure of how much an advertiser engages in ad aggregation. This quantity can be measured using data sources such as browser toolbars that track user click behavior or from web proxy logs. For this paper, we present results from using logs from Windows Live Toolbar, a browser add-in that records anonymized user clicks with the users’ permission. In particular, we examined the browsing behavior from about 2700 toolbar users on about 4,500 ad clicks on Windows Live SERPs.

We flagged all advertisers with non-negligible redirection rate, and hand verified that they were indeed aggregators. Figure 1 shows the redirection rates for the top redirectors we identified. All were verified to be aggregators.

4.2 The Intervention

We now describe an empirical study for measuring the contextual effect of aggregators on the CTR of other advertisers. Using the method described above, we selected several aggregators that were active in the market and had large redirection rates for inclusion in the study. We selected aggregators with higher redirection rates because high redirection rates indicate that users stayed on their sites to click on more syndicated ads rather than returning to the SERP to click on ads from other advertisers.

We then randomly selected 1000 search queries in whose auctions these aggregators were active, and intervened in these auctions on Windows Live Search to vary the positions in which aggregators and non-aggregators occur. We then examined click-through on over 2,000,000 SERPs for the selected queries from before and after the intervention. Figure 2 shows how placement of aggregator and non-aggregator ads varied from before to after the intervention. Note that the number of non-aggregator ads in the top (mainline) slots increased while the number of aggregator ads in these slots decreased. Thus, the number of ads that appeared below an




Figure 1: The number of redirections per ad click for the top advertisers by redirection rate. All were verified to be aggregators.

aggregator ad decreased.

4.3 Quantifying Contextual Effects on CTR

In order to quantify the contextual effect of aggregators on CTR, we first build a baseline model that satisfies equation (1). Thus the baseline model predicts clicks based on advertiser identity and position. This baseline model is a logistic regression similar to that proposed in [3], and was trained using the Sequential Conditional Generalized Iterative Scaling algorithm [7]. Due to data sparseness, it is impractical to build separate models for use in predicting CTR for the auctions corresponding to each search query, as implied by equation (1). Therefore, we build a single model for use across all queries, but condition on the search keyword as well.

We then train a model that attempts to improve this baseline by adding a binary contextual feature that fires when an aggregator ad appears above the position in question. Note that this simple feature violates the conditional independence assumption of equation (1). This model is trained using minimum divergence estimation to improve upon the baseline. Because of this, differences between the models are effects that can be explained by the new feature, but not by the old ones. Thus, the difference between the click probabilities of the two models a measure of the inaccuracy of a CTR estimate that ignores contextual effects.

In order to evaluate the reliability of the effect we measure, we randomly divide the 1000 keywords into partitions of 100 and do the analysis separately on the partitions. The results indicate that being listed below an aggregator reduces CTR. This effect is statistically significant, passing a sign test with $p = 0.1\%$ (i.e. all ten partitions gave this result). Figure 3 illustrates how strongly CTR is affected by this contextual effect. It can be seen that being listed below an aggregator lowers CTR by about 10%.

4.4 The Effect On Revenue

We also present results on the effect of intervening to demote aggregators on revenue. Because the average number of ads per SERP in each position has a strong influence

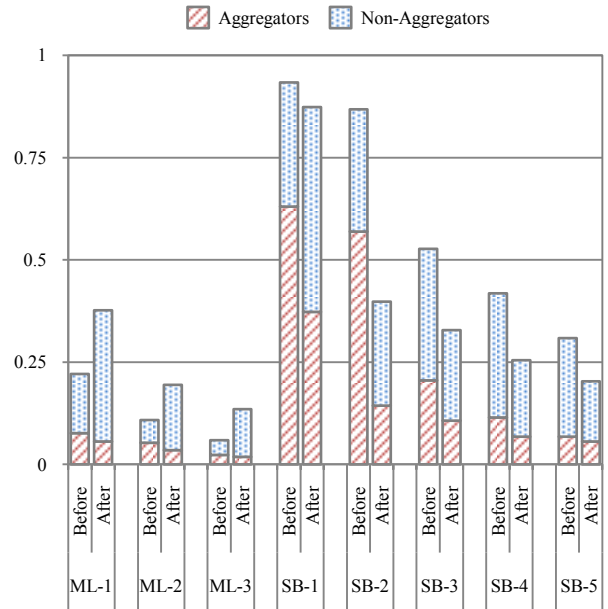


Figure 2: The average number of aggregator and non-aggregator ads per SERP in each ad slot, both before and after the intervention. ML-1 through ML-2 denote mainline slots above the organic results, while SB-1 through SB-5 denote sidebar slots to the right of the organic results.




Figure 3: The change in CTR because of the contextual effect, vs. the baseline CTR. The broken lines indicate +/- one standard deviation.




Figure 4: Revenue per SERP before and after the intervention, averaged across keyword. Keywords were weighted by number of queries. Revenue is broken out for aggregators and non-aggregators, and revenue from non-aggregators is further broken out by whether the advertiser lowered their bid, kept it constant, or raised their bid from before to after the intervention.

on revenue, we controlled for this by selecting 45 search queries for which this did not change. Figure 4 shows revenue per SERP before and after the intervention, averaged across these search queries. The revenue from aggregators and non-aggregators is shown separately, and the revenue from aggregators is broken out by whether the advertiser lowered their bid, kept it constant, or raised their bid from before to after the intervention. It can be seen that revenue per SERP after the intervention was 35% higher. The number of impressions per SERP in each position and the distribution over search queries is controlled for, and it can be seen that the increase is due mainly to non-aggregators that did not raise their bids. Therefore, we argue that the increase was caused by demoting aggregator ads. This increase in revenue suggests that total advertiser value per SERP also increased.

5. CONCLUSION

We have demonstrated that contextual effects do influence the CTR of search ads. In particular, we have shown that being listed below an ad aggregator can lower CTR by as much as 10%, and that this effect is not explained by the position and identity of the ad being modeled. We have also shown that lowering the ranking of ad aggregators relative to the ranking assigned by the standard algorithm can improve search engine revenue by as much as 35%. Thus, our results suggest that the use of click-through models that take contextual effects into account could have a positive impact on revenue. However, the use of models that incorporate such effects can significantly increase the computational complexity of finding the optimal ranking. How best to take contextual effects into account while keeping the slot allocation problem tractable is an area for future research.

6. REFERENCES

- [1] Z. Abrams and M. Schwarz. Ad auction design and user experience. In *WINE*, pages 529–534, 2007.
- [2] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *EC*, pages 1–7, 2006.
- [3] H. Becker, C. Meek, and D. M. Chickering. Modeling contextual factors of click rates. In *AAAI*, pages 1310–1315, 2007.
- [4] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *Amer. Econ. Rev.*, 97(1):242–259, 2007.
- [5] Y. Engel and D. M. Chickering. Incorporating user utility into sponsored-search auctions. In *AAMAS*, 2008. To appear.
- [6] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the Association for Computing Machinery*, 34(2):596–615, July 1987.
- [7] J. Goodman. Sequential conditional generalized iterative scaling. In *ACL*, 2002.
- [8] Google, Inc. Google announces fourth quarter and fiscal year 2007 results, 2007.
- [9] A. Gunawardana, C. Meek, and J. Biggs. A quality-based auction for search ad markets with aggregators. In *EC*, 2008. In review.
- [10] H. R. Varian. Position auctions. *Int. J. Indust. Org.*, 2008. To appear.