# Position Auctions with Externalities

Patrick Hummel[1] and R. Preston McAfee[2]

[1] Google Inc.
phummel@google.com
[2] Microsoft Corp.
preston@mcafee.cc

**Abstract.** This paper presents models for predicted click-through rates in position auctions that take into account the externalities ads shown in other positions may impose on the probability that an ad in a particular position receives a click. We present a general axiomatic methodology for how click probabilities are affected by the qualities of the ads in the other positions, and illustrate that using these axioms will increase revenue as long as higher quality ads tend to be ranked ahead of lower quality ads. We also present appropriate algorithms for selecting the optimal allocation of ads when predicted click-through rates are governed by a natural special case of this axiomatic model of externalities.

## 1 Introduction

In sponsored search auctions, advertisements appear alongside search results in a variety of positions on the page, some of which are more prominent and thus more likely to be clicked than others. In both academic work and practice, it is standard to model each position as having some quality score that reflects the relative probability that an ad will receive a click in that position and then ranking the ads by a product of their bid, the maximum amount the advertisers will pay per click, and a quality score, which reflects the probability an ad will receive a click if the advertiser is shown in the top position.

Although it is almost universal to assume that an ad's click probability is a product of the ad's quality score and a quality score of a position, this formulation may be suboptimal. The formulation implicitly assumes that the probability an ad receives a click in a given position is independent of the identities of the other ads on the page. However, this assumption is unlikely to hold in practice. [3] notes empirically that many consumers are likely to search for what they are looking for by beginning their search at the top and ceasing to search after they have found what they are looking for, and [11] presents evidence that ads impose large negative externalities on other ads by virtue of the fact that the ads can be substitutes for one another. Consequently, placing higher quality ads at the top of the page decreases the probability that a user clicks on other ads.

How then, might one incorporate this possibility into sponsored search auctions to choose a more efficient allocation of ads? This paper presents methods for achieving this goal. We begin by presenting a general axiomatic model of

T.-Y. Liu et al. (Eds.): WINE 2014, LNCS 8877, pp. 417–422, 2014.

predicted click-through rates when the probability an ad receives a click may depend not only on the quality score of the ad and the position in question, but also on the quality scores of the other ads that are shown in the other positions. We analyze the properties of this axiomatic formulation, and illustrate that as long as higher quality ads are typically ranked ahead of lower ads, then moving towards this new axiomatic model will increase revenue in expectation.

A drawback of the most general possible formulation is that computing the optimal allocation of ads is unlikely to be computationally feasible because one would likely need to try each possible configuration of ads in order to choose the optimal configuration, and this is likely to be too slow to be useful in practice. For this reason, we also develop a second formulation that is a special case of our most general methodology that has the advantage of admitting a rapidly converging algorithm for computing the optimal allocation of ads.

While a few papers have presented theoretical analyses of circumstances where the click-through rates are not equal to a product of the quality score of an ad and the quality score of a position, these papers differ significantly from our paper. [1] and [12] consider models in which users search from the top to the bottom that lead to non-separable click probabilities and tractable algorithms for choosing the allocation of ads, but do not consider more general models, as we do in this paper. [7] and [8] further analyze the equilibrium and efficiency properties of such a model, but [11] finds empirical evidence that models in which users search from the top of the page to the bottom do not fully match the data. Other papers that consider different models of externalities (*e.g.* [2], [5], [6], and [10]) also do not present algorithms for choosing the allocation of ads.

## 2 Model of Externalities

There is an auction for $s$ advertising positions on a page. Each advertising position $k$ has a quality score $n_k$, where we assume without loss of generality that $n_k$ is non-increasing in $k$. There are also $m$ advertisers. Each advertiser $i$ has a quality score $q_i$ reflecting the relative clickability of the ad and makes a bid $b_i$ reflecting the maximum amount that this advertiser will pay per click.

In this setting, a standard model of position auctions such as [4] or [13] would assume that the probability advertiser $i$ receives a click in position $k$ is $n_k q_i$. We instead allow the probability an advertiser receives a click to depend on these quality scores in a more nuanced way. Let $p_{(j)} = f_j(q_{(1)}, \ldots, q_{(s)}; n_1, \ldots, n_s)$ denote the probability that the advertiser in the $j^{th}$ position receives a click as a function of the quality scores of the ads in the first $s$ positions as well as quality scores of the $s$ positions. In addition to requiring this probability to be increasing in the underlying quality scores of the ad and the position, $q_{(j)}$ and $n_j$, we stipulate that this probability should satisfy these axioms:

(A) $f_j(q_{(1)}, \ldots, q_{(s)}; n_1, \ldots, n_s)$ is non-increasing in $q_{(k)}$ for all $k \neq j$.

(B) Increasing the quality score of an ad in a higher quality position decreases the click-through rates of ads in other positions by more than increasing the quality score of an ad in a lower quality position. Formally, let $\boldsymbol{q} \equiv (q_{(1)}, \ldots, q_{(s)})$

denote a vector of qualities for which $q_{(i)} = q_{(k)} = q^*$ for some particular $i$ and $k$ satisfying $n_i > n_k$. Also let $\boldsymbol{q}_{(i)}$ denote the vector of qualities that would result from replacing $q_{(i)} = q^*$ with $q_{(i)} = \hat{q}$ for some $\hat{q} \neq q^*$, and let $\boldsymbol{q}_{(k)}$ denote the vector of qualities that would result from replacing $q_{(k)} = q^*$ with $q_{(k)} = \hat{q}$ for the same $\hat{q}$. Then $|f_j(\boldsymbol{q}_{(i)}; n_1, \ldots, n_s) - f_j(\boldsymbol{q}; n_1, \ldots, n_s)| \geq |f_j(\boldsymbol{q}_{(k)}; n_1, \ldots, n_s) - f_j(\boldsymbol{q}; n_1, \ldots, n_s)|$ for all $j \notin \{i, k\}$.

Axiom (A) simply reflects the possibility that when a higher quality ad assumes a particular position, the ad is likely to decrease the probability that ads in other positions receive a click. This axiom is plausible because if the quality of an ad in a particular position increases, users are relatively more likely to click on this ad, which in turn draws their attention from the other ads.

Similarly, axiom (B) reflects the fact that increasing the quality of an ad in a higher quality position does more to increase the probability that users will click on that ad, so increasing the quality of an ad in a higher quality position also draws more user attention from other ads than increasing the quality of an ad in a lower quality position. Thus both axioms (A) and (B) reflect sensible properties on how changing the qualities of ads in other positions is likely to affect the probabilities that other ads receive a click.

Throughout our analysis, we focus on mechanisms in which the auctioneer seeks to maximize total expected welfare with respect to the bids of the advertisers. That is, the auctioneer maximizes $\sum_{j=1}^{s} b_{(j)} p_{(j)}$, where $b_{(j)}$ denotes the cost per click bid of the advertiser in the $j^{th}$ position and $p_{(j)}$ denotes the probability that the advertiser in the $j^{th}$ position receives a click. We also focus on a generalization of the generalized second price auction in which the advertiser in the $j^{th}$ position is charged a cost per click $c_{(j)}$ that represents the smallest bid that this advertiser could make while still maintaining the $j^{th}$ position when the allocation of ads is chosen using the above algorithm.

## 3 General Results

We first derive some general results on how using an alternative model of predicted click-through rates meeting the axioms given in the previous section would affect revenue from online auctions. To do this, we compare two otherwise identical methods for predicting the click-through rates of ads in position auctions. The first method is one in which the predicted click-through rates of the ads in slots $j \notin \{k, k+1\}$ are independent of the quality scores of the ads in positions $k$ and $k+1$, as in a standard model. The other method we consider is one in which the predicted click-through rates of the ads in positions $j \notin \{k, k+1\}$ may depend on the quality scores of the ads in positions $k$ and $k+1$ in a manner that satisfies the axioms (A) and (B) presented in the previous section.

There are two different ways that incorporating the possibility that the quality scores of ads may affect the click-through rates of ads in other positions could affect revenue. First there is the possibility that this could affect the allocation of ads that is shown in the auction. In this case, if the revised model of predicted click-through rates is more accurate, then one would choose a more efficient allocation of ads, and thereby typically achieve higher revenue.

However, in a substantial percentage of auctions, allowing for the possibility that an ad's predicted click-through rate may depend on the quality scores of the other ads will not change the allocation of ads but will affect the pricing. It is thus important to assess how the prices that the advertisers pay would be affected by the changed model of predicted click-through rates even if this does not affect the allocation of ads. This is addressed in the following theorem:

**Theorem 1.** *Consider two different models of predicted click-through rates for position auctions that are identical except for the following:*

*(1) For the first model, the predicted click-through rates of ads in slots* $j \notin \{k, k+1\}$ *are independent of the quality scores of the ads in slots* $k$ *and* $k+1$*.*

*(2) For the second model, the predicted click-through rates of ads in slots* $j \notin \{k, k+1\}$ *depend on the quality scores of the ads in slots* $k$ *and* $k+1$ *in a manner that satisfies axioms (A) and (B).*

*Then if the allocation of ads that is selected by the two models of predicted click-through rates is identical, the advertiser in position* $k$ *pays more per click under the second model if and only if* $q_{(k)} > q_{(k+1)}$.

All proofs are in the appendix of the full paper [9]. Theorem 1 indicates that if we take into account the externalities that the ads in positions $k$ and $k+1$ impose on the other ads, then the advertiser in position $k$ will pay more per click if and only if this advertiser has a higher quality ad. Since Theorem 1 applies to all slots $k$, repeatedly applying Theorem 1 to every slot suggests that if a model with externalities has no effect on the allocation of ads, then this model will typically increase the cost per click paid by an advertiser if and only if this advertiser's quality score exceeds that of the advertiser just below him.

The results of this section suggest that if one can more accurately describe click probabilities by using a model of the form in Section 2, then one should be able to increase revenue. Typically higher quality ads will be ranked higher than lower quality ads, so the result in Theorem 1 suggests that even if this model does not change the allocation of ads, revenue should still increase. And if one is able to choose a more efficient allocation, then one would also expect revenue to increase. Thus revenue is likely to increase from using predicted click-through rates of the form in Section 2 as long as such a model is more accurate.

## 4 Practical Formulation

For general models of the form in Section 2, it may be difficult to select the efficiency-maximizing configuration because there are an exponentially large number of feasible configurations and it is not obvious how one can rule out different configurations as dominated by others. Thus it is important to use a model where one can select the efficiency-maximizing configuration in a computationally tractable way. In this section we present a specific formulation of the model in Section 2 that permits such a practical implementation.

In particular, in this section we consider a model in which the predicted click-through rates for the ads in position $i$ are of the form

$$p_i = \frac{\nu n_i q_i}{1 + \lambda \sum_{j=1}^{s} n_j q_j}$$

where $\lambda$ and $\nu$ are positive constants, $n_i$ denotes the quality score of the $i^{th}$ position, and $q_i$ denotes the quality score of the ad in the $i^{th}$ position. This formulation is sensible because one would expect the percentage decrease in an ad's click-through rate due to the negative externalities imposed by the other ads to be proportional to the total click-through rates of these ads, meaning an ad's click-through rate is likely to be decreased by a factor proportional to $1 + \lambda \sum_{j=1}^{s} n_j q_j$. Also note that in this formulation, setting $\lambda = 0$ and $\nu = 1$ would recover the standard formulation of predicted click-through rates, so optimally choosing these parameters can never result in less accurate predicted click-through rates than the standard formulation. Further note that changing the value of $\nu$ would never change the optimal allocation of ads. Here $\nu$ is a term that only serves to make the predicted click-through rates unbiased on average.

Now define $S$ to be the expected social welfare from a given ranking of ads when $\nu = 1$. We first begin our analysis of this formulation by noting when using a non-zero value of $\lambda$ would result in changing the allocation of ads:

**Theorem 2.** *Welfare is enhanced by switching the order of the ads in positions $k$ and $m$ where $k < m$ if and only if $b_m q_m - \lambda q_m S > b_k q_k - \lambda q_k S$.*

Theorem 2 suggests that if one can obtain a good estimate of the social welfare $S$ that will result in the efficiency-maximizing configuration, then it may be feasible to rank the ads on the basis of scores of the form $b_m q_m - \lambda q_m S$ to achieve the efficiency maximizing allocation. We now exploit this insight to derive a computationally efficient way of selecting the optimal ordering of ads.

The algorithm proceeds by selecting a value $S_L$ that is lower than the social welfare $S$ that will result in the efficiency-maximizing configuration and another value $S_H$ that is higher than this social welfare $S$. The algorithm then repeatedly replaces either $S_L$ or $S_H$ with $\hat{S} \equiv \frac{1}{2}(S_L + S_H)$ until it finds some such value of $\hat{S}$ that is guaranteed to result in the efficiency-maximizing allocation when ranking ads by the scores $b_m q_m - \lambda q_m S$ when $S = \hat{S}$. Such an algorithm will typically require very few steps in practice because after $n$ passes, $\hat{S}$ will be within a factor of $2^{-n}$ of the true social welfare $S$ corresponding to the efficiency-maximizing configuration. The detailed steps for the algorithm are as follows:

(1) Define $S_L$ to be the expected social welfare that would result if the ads were ranked by the scores $b_m q_m$.

(2) Define $S_H \equiv \sum_{m=1}^{s} n_m b_m q_m$ when the ads are ranked by the scores $b_m q_m$.

(3) Calculate the rankings of the ads when the ads are ranked by the scores $b_m q_m - \lambda q_m S$ for $S = S_L$ and $S = S_H$.

(4) If the rankings of the ads in step (3) are the same for both $S = S_L$ and $S = S_H$, then choose this ranking of the ads.

(5) If these rankings are different, let $\hat{S} \equiv \frac{1}{2}(S_L + S_H)$ and calculate the ranking of the ads when the ads are ranked by the scores $b_m q_m - \lambda q_m \hat{S}$.

(6) Let $\phi(\hat{S}) \equiv \sum_{m=1}^{s} n_m(b_m q_m - \lambda q_m \hat{S})$ when the ads are ranked by the scores $b_m q_m - \lambda q_m \hat{S}$. If $\phi(\hat{S}) < \hat{S}$, then let $S_H = \hat{S}$. Otherwise let $S_L = \hat{S}$.

(7) Repeat steps (3)-(6) until the rankings in step (4) are the same for both $S = S_L$ and $S = S_H$, and choose the resulting ranking of ads.

This algorithm indeed results in the efficiency-maximizing allocation:

**Theorem 3.** *The ranking of ads that results from the algorithm considered above is the efficiency-maximization allocation.*

This completes our results for the model of position auctions with externalities. In the full version of the paper [9], we also present an additional model of position auctions that takes into account the fact that ads from well-known brands are less adversely affected by being shown in a lower position than ads from better-known brands [10]. In this model of brand effects, we again present appropriate algorithms for selecting the optimal allocation of ads and show that a purely greedy approach of ranking the ads will potentially cost as much as half of the total possible social welfare. We refer the reader to [9] for more details.

## References

[1] Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M.: Sponsored search auctions with markovian users. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 621–628. Springer, Heidelberg (2008)
[2] Athey, S., Ellison, G.: Position auctions with consumer search. Quarterly Journal of Economics 126(3), 1213–1270 (2011)
[3] Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: WSDM 2008, pp. 87–94 (2008)
[4] Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars of keywords. American Economic Review 97(1), 242–259 (2007)
[5] Fotakis, D., Krysta, P., Telelis, O.: Externalities among advertisers in sponsored search. In: Persiano, G. (ed.) SAGT 2011. LNCS, vol. 6982, pp. 105–116. Springer, Heidelberg (2011)
[6] Ghosh, A., Mahdian, M.: Externalities in online advertising. In: WWW 2008, pp. 161–168 (2008)
[7] Giotis, I., Karlin, A.R.: On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 629–638. Springer, Heidelberg (2008)
[8] Gomes, R., Immorlica, N., Markakis, E.: Externalities in keyword auctions: An empirical and theoretical assessment. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 172–183. Springer, Heidelberg (2009)
[9] Hummel, P., McAfee, R.P.: Position auctions with externalities and brand effects. arXiv:1409.4687 [cs.GT] (2014)
[10] Jerath, K., Ma, L., Park, Y.-H., Srinivasan, K.: A "position paradox" in sponsored search auctions. Marketing Science 30(4), 612–627 (2011)
[11] Jeziorski, P., Segal, I.: What makes them click: Empirical analysis of consumer demand for search advertising. Typescript, University of California Berkeley, Berkeley, CA (2014)
[12] Kempe, D., Mahdian, M.: A cascade model for externalities in sponsored search. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 585–596. Springer, Heidelberg (2008)
[13] Varian, H.: Position auctions. International Journal of Industrial Organization 25(6), 1163–1178 (2007)