

# Northern Exposure: A Field Experiment Measuring Externalities between Search Advertisements

David Reiley  
Yahoo! Labs  
4301 Great America Parkway  
Santa Clara, CA 95054  
reiley@yahoo-inc.com

Sai-Ming Li  
Yahoo! Labs  
4301 Great America Parkway  
Santa Clara, CA 95054  
eliot@yahoo-inc.com

Randall Lewis  
Yahoo! Labs and MIT Economics  
4301 Great America Parkway  
Santa Clara, CA 95054  
ralewis@yahoo-inc.com

## ABSTRACT

“North” ads, or sponsored listings appearing just above the organic search results, generate the majority of clicks and revenues for search engines. In this paper, we ask whether the competing north ads exert externalities on each other: does increasing the number of rival north ads decrease the number of clicks I receive on my own north ad? We conduct a controlled experiment to investigate this question, and find to our surprise that additional rival ads in the north tend to increase rather than decrease the click through rate (CTR) of the top sponsored listing. We propose several possible explanations for this phenomenon, and point out directions for new theoretical models of sponsored search.

## Categories and Subject Descriptors

H.5.4 [Information Interfaces and Presentation]:  
Hypertext/Hypermedia; J.4 [Social and Behavioral Sciences]:  
Economics

## General Terms

Measurement, Economics, Experimentation

## Keywords

Sponsored search, advertising, externalities, page placement

## 1. INTRODUCTION

Online search engines respond to a query with a combination of two types of search listings: sponsored and organic. Organic results are unpaid listings, while sponsored listings are paid advertisements. The placement of a sponsored result depends on the competing advertisers’ bids as well as the quality of each advertisement (where higher-quality ads are those more likely to be clicked). All else equal, ads with higher bids appear higher up

in the listings, as do ads with higher quality. Sponsored ads with sufficiently high bids and sufficiently high quality for a given query appear directly above the organic listings in what is known as the “north” position. Lower-ranked or lower-quality ads appear to the right or below the organic listings in what are known as the “east” and “south” positions, respectively. This placement strategy aims to optimize both revenue and the experience of is aimed at maximizing both user experience and revenue.

From the search engine’s point of view, there is a tradeoff in changing the number of north listings: more sponsored listings can increase revenue but may do so at the expense of user experience by displacing organic listings that might well be more relevant to the user. Search engines limit the number of north ads for this reason, using a variety of direct and indirect constraints on quality and expected revenue. For example, none of the top three search engines (Google, Yahoo!, Bing) show more than four north ads on any query, and quite frequently they show only zero or one.

Advertisers are willing to bid higher in order to get better placement, and therefore more clicks, in the generalized second price auction (GSP) used by most search engines [1,2]. Knowing that additional clicks are valuable to the advertiser, we start this paper with the question: Would an advertiser be willing to pay for exclusive placement in the north?

One might reasonably assume that north ads would exert negative externalities on each other [3]. In particular, if I occupy the top advertising position and my ad appears in the north, then the presence of other ads below me in the north would likely take clicks away from me. If this is the case, advertisers would be willing to pay a premium for exclusive north placement, with the size of the premium determined by the number of additional clicks that exclusivity would generate.

Early published models of sponsored search [1,2] ignore the subject of externalities between ads. [3] and [4] provide theoretical models for such click externalities, while [5] provides both theory and empirical evidence showing that the identity of a rival advertisement matters: the presence of a strong rival ad can cause me to receive fewer clicks than I would receive in the presence of a weak rival ad. By contrast, in this paper we look at variation in the *number* of north ads, rather than the identities of the rivals. To get clean measurements, making sure that we get sufficient observations in each number of north positions, we

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference '04, Month 1–2, 2004, City, State, Country.

perform a controlled experiment, randomly varying the number of ads on each query.

To get clean measurements of externalities, we perform a controlled field experiment, randomly varying the number of ads on each query. To our surprise, we find that rival north ads impose a positive, rather than negative, externality on existing north listings. In other words, the top north listing receives more clicks when additional sponsored listings appear below it. We suggest some possible explanations for this phenomenon, and point out directions for future research. We particularly note the importance of distinguishing north versus east ads, as well of understanding consumer substitution between organic and sponsored results.

The rest of the paper is organized as follows: section 2 describes the experimental design, section 3 describes the experiment results, section 4 provides possible explanations of the observations, and section 5 summarizes our findings and provides suggestions on future work.

## 2. EXPERIMENTAL DESIGN

This goal of the experiment is to measure how additional sponsored listings in the north impact the CTR of existing north listings. In order to gather enough data to estimate the CTR for the same position across different numbers of north ads, we randomize the number of north ads for a representative sample of search queries Yahoo! We present results aggregated across a sample of many queries, users, and sponsored listings.

By running a careful field experiment, we obtain observations with exogenously varying numbers of north ads on a representative set of queries, solving the problem of causal inference that plagues many empirical studies. If we had instead relied on naturally occurring field data, we might well have reached mistaken conclusions. For example, suppose it turns out that our search engine always showed four north ads for “car rental” and always showed one north ad for “Taurus,” and further suppose that clickthrough rates (CTRs) for north ads on “car rental” are around 20%, while CTRs for north ads on “Taurus” are 5%. If we looked at this data, we would find that the top north ad gets more clicks with three rival ads than with no rival ads, but this represents correlation rather than causality. Randomizing the number of north ads from a set of queries allows us to identify the causal effects of the number of ads.

The experiment was carried out on 2% of all Yahoo! Search users. For these users, 10% of searches were randomly selected and the number of listings placed in the north was randomly set from zero to four, disregarding all existing rules on placement. Of these modified searches, we only analyze those search results that had four or more eligible sponsored listings to show.

## 3. EXPERIMENTAL RESULTS

Before presenting the results, we provide some basic definitions and formulas used to compute estimates and standard errors for CTRs. First, the CTR for the  $i^{\text{th}}$  north position when there are  $j$  north ads present is defined as

$$CTR(i, j) = \frac{clicks(i, j)}{impressions(i, j)}$$

where the notation  $(i, j)$  is defined for  $i \leq j$  and similarly applies to clicks and impressions as to CTR. The CTR for a section of the search listings is similarly defined by restricting attention to the listings in that section (i.e. organic and sponsored north, east, or south). We make use of the well-known convenience that the binomial distribution converges to the Normal distribution and construct confidence intervals and standard errors using the usual estimator of the binomial distribution’s variance,

$$\hat{\sigma}_{CTR(i, j)}^2 = \frac{CTR(i, j) \cdot (1 - CTR(i, j))}{impressions(i, j)}$$

For the purposes of the analysis of this experiment, we assume that all impressions are independent, although certainly many searches are performed by the same Yahoo! Search user and many clicks occurred on the same page of search results. While we expect positive correlation across users to increase the variability of our CTR estimates, we would expect that same correlation to reduce the variability of the estimated differences in CTR across experimental treatments as users were permitted to be exposed to all five treatments of zero to four north ads. Regardless, we expect these differences to have little qualitative effect on our main conclusions.

### 3.1 North CTR

Figure 1 shows the CTR for each north position and number of north listings. Surprisingly, at almost every rank the CTR *increases* with the number of listings shown in the north. The effect is especially pronounced in position one. We found this positive externality so unexpected in our first experiment (October 2008) that we replicated the experiment to validate the results. Figure 2 presents the results from the second experiment (June 2009) which involved an 18% larger sample than the first experiment. We limit the remaining analysis for this paper to the second experiment for clarity and brevity of exposition. All results analyzed were qualitatively identical in the first and second experiments.






Figure 1. CTR of north sponsored listings appears to increase with the number of north ads in the first experiment.



**Figure 2. CTR of north sponsored listings also appears to increase with the number of north ads in the replication.**

As can be seen in Figure 3, the position one sponsored listing CTRs have very short 95% confidence intervals (hereafter all confidence intervals are 95% asymptotic confidence intervals). Upon comparing the CTRs we find that displaying four north ads generates a statistically significantly higher CTR for the position one listing. Strikingly, increasing the number of north ads tends to increase the CTR for the first north ad.

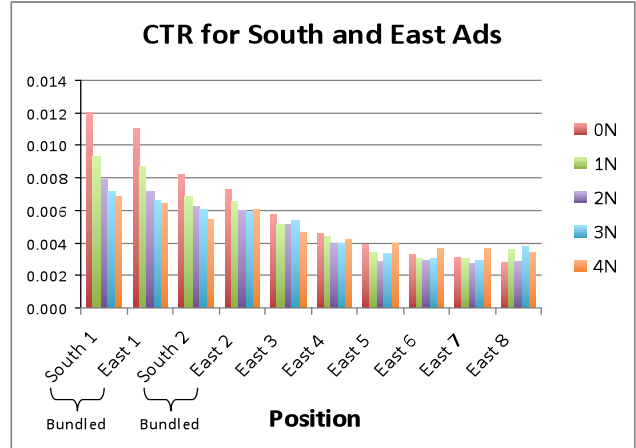


**Figure 3. Truncated histogram of position one north ad's CTR varying the total number of north ads.**

We do not observe any statistically significant difference in CTR for sponsored listings in positions two or three when the number of north ads increases from two or three. We suspect that this is primarily an issue of statistical power and hypothesize that with a larger sample a similar increasing trend would emerge for sponsored listings in positions two and three as the number of north listings increases. However, with the short confidence intervals available, we can safely conclude that the external effects are smaller than 5% of the baseline CTR without additional ads. Thus, it would seem that the largest externality among north ads is imposed when going from one north ad to more than one with the externality being 2.2%, 4.7%, or 12.0% for two, three, or four north ads, respectively.

### 3.2 South and East Ad CTR

Figure 4 shows the CTR for each south and east position and number of north listings. The CTR declines with the number of north listings for almost every position. However, the volatility of the data prevents a precise finding of decline for all positions.



**Figure 4. CTR of south and east ads declines with the number of north ads due to both ad quality and externalities.**

Further, influencing the interpretation of the decline, the ads in each position varied with the number of north ads. The ads were ranked from first to twelfth prior to their placement on the page, and then the ads were first placed into the available number of north positions, ranging from zero to four, and then into the eight east positions until there were no more available slots. Thus, while some of the clicks gained by the additional north listings may be at the expense of clicks to these sponsored listings, part of the decline in CTR when increasing the number of north ads may result from lower-quality ads occupying these positions.

Also noteworthy is that the south positions are bundled with the top two east ads on Yahoo! Search. This is important because the relevant CTR to the advertiser would then be the combined CTR. Increasing the number of north ads beyond four would be inadvisable without changing this bundling or the order of bid rankings to avoid the total CTR from the fifth north ad from exceeding the total CTR accruing to the ad located in the first east and south bundle. This is important to maintain incentive compatibility of bidding and transparency of the mechanism to the advertiser that higher bids will result in more clicks.

While Yahoo! Search bundles these listings, other major internet search engines do not. Google does not include any south ads in the results and, hence, avoids this concern. Bing bundles the two south positions with the top two north. Ask and AOL both omit any east listings and bundle several north positions with the top south positions, below which additional ad listings are sometimes included. We highlight that while concerns regarding the south ads are specific to Yahoo!, the externalities estimated for the north ads likely generalize to all major internet search engines which use similar formatting by including the primary ad units directly above the organic results in the north position.

### 3.3 Organic CTR

Figure 5 shows the CTR for each organic position and number of north listings. The CTR at each organic position decreases with the number of north listings. This suggests some of the clicks gained by the additional north listings are at the expense of organic clicks.




Figure 5. CTR of organic listings declines in the number of north ads for all organic search positions.

### 3.4 Overall CTR

We define the total CTR for a section when there are  $j$  north ads present as

$$CTR(j) = \frac{\sum_i clicks(i, j)}{searches(j)}$$

where we sum the clicks over all positions in a given section and then divide by the number of search queries with  $j$  north ads (per our random assignment). The overall or total search listing CTR is simply the sum of the section CTR over the four sections: organic, north, east, and south. Figure 6 shows this CTR for the entire search result page, broken down by section. While the CTR of north listings increase with the number of north listings at the expense of CTR in all other sections, the total CTR of all search listings combined is goes up slightly. If we assume that all clicks are equal, then more north listings are better as more clicks are generated overall with each additional north ad as can be clearly seen in Figure 7. However, it is possible that the additional clicks have lower utility to the user than existing ones. We next look at the distribution of searches with a given number of clicks to get a better understanding of the user experience impact of more north listings.

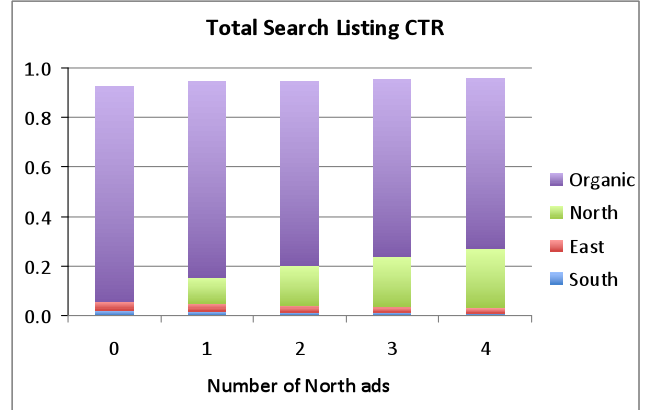


Figure 6. The distribution of CTR across search listing sections shows a large substitution of clicks from organic to north listings.




Figure 7. Average number of clicks per search increased with the number of north ads.

### 3.5 Click distribution of searches

Figure 8 shows the composition of click responses for searches under the varying number of north ads. To highlight the statistically significant differences, Figure 9 and Figure 10 show truncated histograms of the fraction of searches with zero and more than one click, respectively. The fraction of searches with no click does not change in any observable pattern with the number of north listings; the fraction of searches with multiple clicks increases steadily with the number of north listings. The range in which the fraction of searches fluctuates is smaller for those with no click relative to those with more than one click. Putting more sponsored listings in the north transforms one-click searches into multiple-click searches. Multiple clicks on the same search page is not necessarily an indication of bad user experience because users may find more than one highly relevant result and want to explore them all. However, in this experiment, since sponsored listings are generally less relevant than the organic listings they displace, increasing the fraction of searches with more than one click suggests that users have to work harder to find what they want.




Figure 8. Fraction of searches divided into zero, one, and more than one click varies with the number of north ads.




Figure 9. Fraction of searches with no click varies across different number of north ads.




Figure 10. Fraction of searches with more than one click increases with the number of north ads.

## 4. EXPLANATIONS

We offer several hypotheses that can explain the experimental result. These hypotheses are not mutually exclusive; hence, it is possible that more than one is relevant.

### 4.1 Substitution away from organic listings

Most theories of search advertising simplify the click-generation process in a way that would rule out the results obtained in this experiment. Early theories, such as [1,2] treated the consumer as a black box and make assumptions about CTRs that decline by position in the search listings. More recently, some theories have modeled consumer choice explicitly, but these simplify the choice problem by assuming away the organic search results and focusing only on the sponsored links [4,6]. Our results (see especially section 3.2) cannot be explained by such models, because we find that CTRs both for advertisements and organic search results depend on the number of north ads.

Existing models of human interaction with search results [4,5,6] make a reasonable top-down or “cascade” assumption that users, expecting the top results to be most relevant, start at the top of the page and scan through the search results sequentially until making a choice. Heuristics of this sort make sense for users, given the time cost of reading all the information on the page, and could also provide an explanation for our results if we take the organic results into account. Including an additional sponsored listing pushes the top organic results farther down the page, potentially making a sponsored result seem more relevant by comparison with what’s shown directly below it. Such results can also be enhanced by psychological framing effects, as described in [8] and [9].

One possible consequence of this behavior is that users may be more likely to find the first listing they click on does not fulfill their need, as they may have neglected the more relevant organic listings before their first click. It is more likely that users will have to go back to the search result page to find other listings. The observation that the fraction of multiple click searches increase with number of north listings support this hypothesis.

### 4.2 Comparative search behavior

Another hypothesis is that when users are presented with multiple sponsored listings with similar messages, they tend to click on more than one such listing to see what each has to offer. This has the effect of decreasing the value of each click to the advertiser as the conversion rate would likely decrease.

The search engine’s challenge is to balance the increase in clicks with the decrease in traffic quality to advertisers, recognizing that a decline in traffic quality leads to a decline in the equilibrium bid by the advertisers. This leaves the effect of increasing clicks on ad revenue ambiguous because clicks and bids are moving in opposite directions with the number of north ads.

### 4.3 Impulsive search behavior

The final hypothesis that we consider is related to impulsive shopping behavior. We hypothesize that users tend to click on ads more if they are shown more ads because they suddenly feel that they need to buy something that they did not need before. We would expect that conversions would increase due to increasing the number of ads by attracting more impulsive customers to the advertisers’ websites.

While such a boost in sales clearly represents an increase in value from the advertisers' perspective, this may represent a decline in user experience as impulsive customers may find it optimal to frequent such a search engine less to avoid the temptation of additional ads. However, if sales merely increase without any negative effect on user experience, we would conclude that the increase in the number of ads improved the matching process of users with websites.

## 5. SUMMARY AND FUTURE WORK

### 5.1 Summary

To better understand the balance of clicks and revenue between sponsored and organic search results, we have conducted an experiment to measure the externalities imposed by additional north listings on existing ones. Surprisingly, additional north listings impose a positive—not a negative—externality on existing north listings. In other words, existing north listings have a higher CTR when more sponsored listings are shown below it. The increase in the CTR does not come at the expense of the additional north ads, but partially at the expense of the organic results. However, the total number of clicks registered on the page also increases, suggesting that increasing the number of north ads provides users an incentive to click on more links on the page. We propose several explanations to these observations, and suggest more research to verify them.

### 5.2 Future research

This experiment allowed us to learn several surprising results about the effect of increasing the number of north ads and the externalities that the inclusion of additional ads has on existing ads. For further research, we propose several additional experiments and analyses to decompose online search results and investigate the heterogeneity of these effects and our hypothesized explanations.

#### 5.2.1 Examine externality heterogeneity by search keywords.

The assumption that all searches are the same is certainly far from the truth. By performing a similar experiment on perhaps 500 representative keywords, each with at least 5,000 searches, we can gain a better understanding of how these results vary at the keyword level. In particular, popular or competitive queries may exhibit different sizes of externalities than obscure queries. For the current experiment, we have at most 200 searches for a given keyword, making this analysis unfeasible with the current experiment. A power calculation yields an expected t-statistic of just greater than 1.96 for a sample of size 5,000 at the individual keyword level. This calculation assumes that the externalities are homogeneous by keyword and of magnitudes similar to those estimated in this experiment for position one with one and four north ads (for the north section,  $CTR(1,4) - CTR(1,1) = 0.0127$ ).

#### 5.2.2 Compare higher-quality organic and sponsored listings with those of lower-quality.

User click behavior on sponsored listings may depend on the quality of organic results. Search result pages with higher-quality organic listings may be harmed by increasing the number of less relevant north ads. However, if the organic search results are generally lower quality, including an additional north ad of a given quality may improve the search experience by contributing a sponsored result that is relevant to users. Better understanding how organic listings' quality interacts with sponsored listings'

relevancy may lead to insights that can improve both the user experience and the value of the sponsored listings for advertisers.

#### 5.2.3 Compare reactions of users with different number of default listings shown on screen

The total number of visible sponsored and organic listings on a search results page depends on users' screen and font sizes. Users who can see more organic listings may behave differently from those who see fewer. If the north ads completely crowd out the organic results from the page display when first loaded, users will be forced to scroll down to even find the first organic result.

By tracking which results were visible to users, we could test this source of heterogeneity in future experiments. It could be that the optimal number of north ads might change with the quality of the organic results or with the number of viewable listings on a user's screen.

#### 5.2.4 Compare conversion rate of position one clicks across different number of north ads

While users are clicking more when more ads are shown in the north, those clicks may have lower conversion rate than if only one ad is shown in the north. If the additional clicks lead to at

least proportionally more conversions, then we can safely conclude that increasing the number of north ads improves the advertiser's return.

## 6. ACKNOWLEDGMENTS

We would like to thank Kannan Achan, Erick Cantu-Paz, Rukmini Iyer, Ben Shahshahani, and Ajay Shekhawat for their advice and assistance in conducting the experiment. We thank Dan Akerberg and Preston McAfee for helpful discussions.

## 7. REFERENCES

- [1] Varian, H. R. 2007. Position auctions. *Int. J. Ind. Organ.* 25 (Dec. 2007) 1163-1178.
- [2] Edelman, B., Ostrovsky, M., and Schwarz, M. 2007. Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *Am. Econ. Rev.* 97 (Mar. 2007) 242-259.
- [3] Ghosh, A. and Mahdian, M. 2008. Externalities in online advertising. In *Proceeding of the 17th international conference on World Wide Web (New York, NY, USA). WWW '08*. ACM Press, New York, NY, USA, 161-168. DOI= <http://doi.acm.org/10.1145/1367497.1367520>
- [4] Kempe, D. and Mahdian, M. 2008. A Cascade Model for Externalities in Sponsored Search. In *Proceedings of the 4th International Workshop on Internet and Network Economics (Shanghai, China). WINE '08*. Springer-Verlag, Berlin, Heidelberg, 585-596. DOI= [http://dx.doi.org/10.1007/978-3-540-92185-1\\_65](http://dx.doi.org/10.1007/978-3-540-92185-1_65)

- [5] Gomes, R., Immorlica, N., and Markakis, E. 2009. Externalities in Keyword Auctions: An Empirical and Theoretical Assessment. In Proceedings of the 5th International Workshop on Internet and Network Economics (Rome, Italy). WINE '09. Springer-Verlag, Berlin, Heidelberg, 585-596. DOI= [http://dx.doi.org/10.1007/978-3-642-10841-9\\_17](http://dx.doi.org/10.1007/978-3-642-10841-9_17)
- [6] Athey, S. and Ellison, G. 2009. Position Auctions with Consumer Search. NBER Working Paper Series no. 15253.
- [7] Segal, I. and Jeziorski, P. 2009. What Makes Them Click: Empirical Analysis of Consumer Demand for Internet Search Advertising, Working Paper, Feb. 2009.
- [8] Tversky, A. and Kahneman, D. 1981. The Framing of Decisions and the Psychology of Choice. Science 211 (Jan. 1981) 453-458.
- [9] Schwartz, B. 2004. Paradox of Choice: Why More is Less. HarperCollins Publishers, Inc. New York, NY, USA. 126-128.