Making an Impression: Learning from Google Ads*

Gaurav Sood

October 31, 2018

Broadly, Google Ads works as follows: 1. Advertisers create an ad, choose keywords, and make a bid (on cost-per-click or CPC),[1], 2. the Google Ads account team vets whether the keywords are related to the product being advertised, and 3. people see the ad from the winning bid when they search for a term that includes the keyword or when they browse content that is related to the keyword (some Google Ads are shown on sites that use Google AdSense).

There is a further nuance to the last step. Generally, on popular keywords, Google has thousands of candidate ads to choose from. And Google doesn't simply choose the ad from the winning bid. Instead, it uses data to choose an ad (or a few ads) that yield the most profit (Click Through Rate (CTR)*bid). (Google probably has a more complex user utility function and doesn't show ads below a low predicted CTR*bid.) In all, who Google shows ads to depends on the predicted CTR and the money it will make per click.

Given this setup, we can reason about the audience for an ad. First, the higher the bid, the broader the audience. Second, it is not clear how well Google can predict CTR per ad conditional on keyword bid especially when the ad run is small. And if that is so, we expect Google to show the ad with the highest bid to a random subset of people searching for the keyword or browsing content related to the keyword. Under such conditions, you can use the total number of impressions per demographic group as an indicator of interest

*The note benefited from comments by Levi Boxell.

[1]You can bid on cost-per-view and cost-per-impression also, but we limit our discussion to CPC.

in the keyword. For instance, if you make the highest bid on the keyword 'election' and you find that total number of impressions that your ad makes among people 65+ are 10x more than people between ages 18-24, under some assumptions, e.g., similar use of ad blockers, similar rates of clicking ads conditional on relevance (which would become same as predicted relevance), similar utility functions (that is younger people are not more sensitive to irritation from irrelevant ads than older people), etc., you can infer relative interest of 18-24 versus 65+ in elections.

The other case where you can infer relative interest in a keyword (topic) from impressions is when ad markets are thin. For common keywords like 'elections,' Google generally has thousands of candidate ads for national campaigns. But if you only want to show your ad in a small geographic area or an infrequently searched term, the candidate set can be pretty small. If your ad is the only one, then your ad will be shown wherever it exceeds some minimum threshold of predicted CTR*bid. Assuming a high enough bid, you can take the total number of impressions of an ad as a proxy for total searches for the term and how often people browsed related content.

With all of this in mind, I discuss results from a Google Ads campaign. The ad campaign was for someone running for the Richmond city council. The campaign was run in the middle of October. The ad bid on the following keywords (phrases): "register to vote, Richmond ca city council candidates, local elections, Richmond city council, California democratic voter guide, east bay times election recommendations, east bay times 2018 election recommendations, Richmond CA election, council election, Richmond city council endorsements Richmond CA elections, candidates for Richmond city council, vote city elections, Richmond ca city council, Richmond city council election 2018." 86% of the impressions were generated by bids on the following keywords (phrases): elections, vote, Richmond CA elections, council election, with elections alone accounting for 50% of the impressions.

The ad was only shown to people located (and people could be just visiting the area)

in the following zip codes: 94804, 94806, 94801, 94803. The maximum bid per click was auto-capped. It cost about $1.15 per click.

Table 1 shows the impressions and clicks by age. Two things jump out. First, the CTR varies heftily and systematically by age. Under plausible assumptions, it suggests that Google doesn't know much about optimizing CTR of a random ad conditional on relevance. It makes sense given the data it observes on a new ad with a small run is limited.

Age	Clicks	Impr.	CTR
18 - 24	10	135	7.41%
25 - 34	19	324	5.86%
35 - 44	17	453	3.75%
45 - 54	30	679	4.42%
55 - 64	43	881	4.88%
65+	56	1,669	3.36%
Unknown	65	2,697	2.41%

Table 1: Ad Impressions and Clicks by Age.

To interpret the sharp trend in impressions, we need to know the rate at which young people use ad blockers vis-a-vis commensurate groups. Thankfully, PageFair allows us to back out the numbers. It seems young people use ad blockers more than older people but not by a large margin. Ad blockers are used most frequently by people between ages 25 and 34, with 22% using them. On the flipside, only 15% of people 65+ use ad blockers. We can use this information to pro-rate the trend.

Table 2 and Table 3 show commensurate tables for income and gender respectively. The relationship between impressions and household income is hefty. The ad is shown to people in the top decile nearly 5.6 as often as those in the 41-50% decile. For gender, we see a slightly unexpected pattern with women seeing the ad 1.3 times more often than men. It may in part be explained by the correlation between age and gender.

Household income	Clicks	Impr.	CTR
Top 10%	46	1,388	3.31%
11 - 20%	32	808	3.96%
21 - 30%	22	574	3.83%
31 - 40%	9	327	2.75%
41 - 50%	13	256	5.08%
Lower 50%	35	730	4.79%
Unknown	83	2,755	3.01%

Table 2: Ad Impressions and Clicks by Income.

Gender	Clicks	Impr.	CTR
Male	73	1,804	4.05%
Female	109	2,436	4.47%
Unknown	58	2,598	2.23%

Table 3: Ad Impressions and Clicks by Sex.