

Discrimination in Online Advertising

A Multidisciplinary Inquiry

[Supplementary Materials]

Amit Datta
Anupam Datta

Carnegie Mellon University

Jael Makagon
Deirdre K. Mulligan

University of California, Berkeley

Michael Carl Tschantz

International Computer Science Institute

AMITDATTA@CMU.EDU

DANUPAM@CMU.EDU

JAEL@BERKELEY.EDU

DMULLIGAN@BERKELEY.EDU

MCT@ICSI.BERKELEY.EDU

Editors: Sorelle A. Friedler and Christo Wilson

Appendix A. A primer on Google’s web advertising

The advertising ecosystem is a vast, distributed, and decentralized system with several parties. First, there are *publishers*, websites which host content. Typical examples of publishers include news websites, social networks, and blogs. As most publishers provide content for free, they earn revenue by serving advertisements on their pages. Second, there are *advertisers*, who seek to place their ads on publishers’ webpages. Advertisers include online shopping portals, firms providing web-services, and traditional brick-and-mortar shops trying to reach their customers online. Third, there are *ad networks* which connect advertisers and publishers. Fourth, there are *consumers* who consume online content. Publishers sign up with an ad network to register available spaces on their webpage. Advertisers also sign up with the ad network to register their ads. When a consumer visits a publisher website, the ad network determines which ad to place on the page for the consumer to view. To make this determination, ad networks use proprietary algorithms that among other variables take into account the expressed preferences of advertisers and publishers, and the specific user’s behavior, and, if communicated, preferences.

In this ecosystem, Google dominates as the largest ad network connecting advertisers to publishers. With nearly a third of all digital ad

revenue, Google has more than three times the market share of the next largest ad network (Facebook) (Media Buying, 2014). Google serves many different types of ads including those beside Google search results (search ads), on YouTube (video ads), and on third-party websites (display ads). As prior experiments found evidence of discrimination on display advertisements, (Datta et al., 2015) we focus on them for the remainder of this section.

Below, we provide more details about how publishers, advertisers, Google (acting as an ad network), and consumers interact for display ads. Our description is largely based upon Google’s own public-facing descriptions. We do not, and often cannot, verify that Google’s descriptions of its internal behaviors are accurate.

Publishers Publishers who serve content on the web for free often rely on advertisements to earn revenue. Most such publishers defer the task of identifying potential advertisers for their site to ad networks. To use Google’s ad network, publishers register their site on AdSense¹ and indicate *ad spaces* on their website where they would like to show ads. These could be locations on the lateral margins, at the top of the page, or inline with content. Google provides small pieces of code that publishers add to their webpage’s source code, thereby enabling Google to serve ads within the ad spaces.

1. adsense.google.com

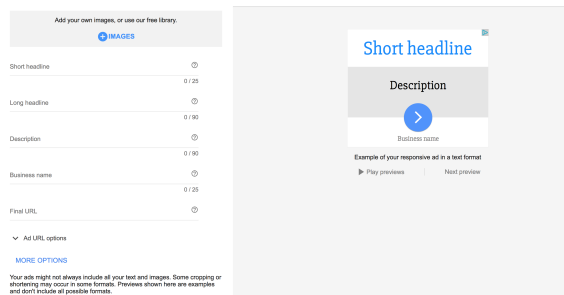


Figure 1: Creating responsive ads on Google AdWords

Advertisers Advertisers are entities who are willing to pay to have their ads shown to consumers. Advertisers wanting to serve advertisements through the Google ad network sign up on Google’s AdWords² interface and communicate to Google their ads. While there are as many as five different types of ad campaigns as of writing this paper, we focus on ‘Display Network Only’ campaigns, since it is the closest to the ad of interest from Datta et al.’s experiment. In such campaigns, Google allows serving four types of ads, namely responsive ads, image ads, gallery ads, and app/digital content ads. For simplicity, we focus on responsive ads, which modern variants of the text ads that were analyzed in the experiment. A responsive ad includes two headlines (a short and a long one), a description, a business name a URL, to which a consumer is directed upon clicking the ad, and an image. Figure 1 shows the interface for creating an ad on Google AdWords.

In addition, Google also enables advertisers to specify targeting criteria for their ads. This helps advertisers reach desired consumers more effectively. While advertisers can target their advertisement based on several criteria established by Google, we are interested in targeting based on demographics (like gender, age group, and parental status) and interests (like Cat Lovers, Mobile Enthusiasts, etc.) since they allow tailoring advertising campaigns based on data Google knows or believes about specific users. This knowledge may be based on user specified settings as discussed below, inferences drawn by Google, and data from other sources. The adver-

tiser must also select a bid, which indicates how much they are willing to pay Google when consumers view or interact with their ads. Google provides a portion of the revenue to the publisher on whose page the ad appears.

Ad Network Once the advertiser has submitted the ad and targeting criteria, Google reviews it to make sure it is safe and appropriate for consumers and complies with Google’s advertising policies, which include a requirement to “comply with all applicable laws and regulations.” (Google, 2017a) Google’s policies restrict certain advertising content, while establishing special procedures for other content. While some gender specific advertising would violate the provision requiring compliance with applicable law, we could find no independent provision limiting advertising content that discriminates based on sex even though they have restrictions that apply to other sensitive categories (like race, sexual orientation, political affiliation, etc.) (Google, 2017b). Once Google approves the ad and the advertiser has entered their billing details, the ad becomes eligible to be served.

Once the ad is eligible, it is up to Google to determine when, where, and to whom to serve the ad. When a consumer visits a publisher’s website and ads from multiple advertisers in its network are eligible for a limited number of ad spaces, Google uses a real-time auction to determine which ad appears and in what order. For each competing ad, Google calculates a metric called the *Ad Rank* which determines how competitive an ad is for a given ad space for a particular consumer. Ad Rank depends upon several factors, but the most important for our purposes, is the *expected click-through rate* (CTR), which is Google’s prediction of how likely the consumer is to click on the ad if shown. Google bases this prediction on what it knows about the current consumer and on its experience from prior click behavior of millions of other consumers.

The Ad Rank determines the position of the ad in the page and whether it shows up at all. Google computes the Ad Rank for an ad every time it competes in an auction and serves the top ranking ads in that order to the consumer. An ad with the same Ad Rank may or may not be served in subsequent auctions depending on the competition at the time. Thus, competing

2. adwords.google.com

ads may influence which consumers an ad gets served to.

Consumers Consumers are Internet users who browse the web and consume online content from publisher websites. Google provides consumers some insight and control over what types of advertisements Google serves them. Ad Settings is a Google tool that helps consumers view and control the ads they see on Google services and on websites that partner with Google. It allows a consumer to select attributes for ad-targeting and to see and modify ad-targeting inferences that Google has made about the consumer. Ad Settings attributes include both demographics and interests based on browsing behavior, which consumers can view and edit at

<http://www.google.com/settings/ads>

Google’s Ad Settings respond to regulators’ and consumers’ concerns about behavioral marketing on the web (Ur et al., 2012). They provide some transparency about how consumers are profiled for ad-targeting, and allow consumers to exercise some modicum of control over the ads they see, including the ability to limit ads targeted based on previous web activity and demographic details. In the summer of 2015, Google updated the Ad Settings page to require being logged in to their Google account to access most of its features, including viewing and editing gender, age, languages and interests.

Appendix B. Discriminatory Ads

Figure 2 shows all the ads that were approved by Google. Ad 2(a) is for a secretary job targeted towards women, while ad 2(b) is for a truck-driving job targeted towards men. These two ads are chosen to exhibit typical gender stereotypes.³ (These two are shown in the main article as well.)

The remaining ads could potentially advertise the same job, but their headlines suggest differences in the terms of employment. The ad 2(c), targeted to women, is for a lower-paying job compared to 2(d), targeted to men. The ad 2(e), targeted to women, were for lower-level job compared to the corresponding ads 2(f), targeted to

3. Google image search results for ‘secretary’ and ‘truck driver’ are heavily skewed.



Figure 2: Ads approved by Google in 2015. Ads in the left (right) column were targeted to women (men).




Figure 3: Ad disapproved by Google in 2017.

men. Ads 2(g) and 2(h) suggest higher educational requirements for women than men.

Figure 3 shows an ad disapproved by Google AdWords for the destination not working and the content being misleading.

Appendix C. History of Sex-specific Help Wanted Columns

Despite its clarity, Section 704(b) of Title VII of the 1964 Civil Rights Act did not put an end to discriminatory advertisements. This was exemplified by the fact that sex-specific help wanted columns in newspapers persisted for nearly a decade after the law took effect. The EEOC’s

initial advertising guidelines allowed for this and were one impetus for founding the National Organization for Women (NOW), and one of NOW's first actions was to challenge them (Pedriana and Abraham, 2006; National Organization for Women, 2017). In addition to challenging the EEOC's guidelines, women and women's organizations sought to bring newspaper job advertisements fully under the purview of the statute arguing that they acted as covered "employment agencies" under Title VII. This argument was rejected in *Brush v. San Francisco Publishing Co.*⁴ based largely on legislative history (Kerman, 1974).

During the following years, the EEOC revised its interpretation of the law twice until finally, under pressure from women's organizations, it adopted a flat prohibition on the use of sex-specific advertising columns by covered entities.⁵ Note this prohibition did not extend to newspapers as publishers of advertisements, but focused on the behavior of regulated employers and employment agencies.

States and localities were busy during this period as well, adopting non-discrimination laws containing aiding and abetting provisions that were aimed in part at prohibiting newspapers from segregating help wanted advertisements into sex-specific sections. These laws complimented Title VII, which did not regulate newspapers, and therefore could not limit them from creating such columns. In *Pittsburgh Press Co. v. Human Rel. Comm'n*, the Supreme Court ultimately upheld the application of the aiding and abetting provisions of these state laws to newspapers as a permissible limitation on commercial speech.

References

Amit Datta, Michael Carl Tschantz, and Anupam Datta. Automated experiments on ad pri-

vacy settings: A tale of opacity, choice, and discrimination. In *Proceedings on Privacy Enhancing Technologies (PoPETs)*, 2015.

Google. Adwords policies, 2017a. URL <https://support.google.com/adwordspolicy/answer/6008942>. Accessed Oct. 2, 2017.

Google. Personalized advertising, 2017b. URL <https://support.google.com/adwordspolicy/answer/143465>. Accessed Oct. 2, 2017.

Peter W. Kerman. Sex discrimination in help wanted advertising. *Santa Clara L. Rev.*, 15 (1):183–209, 1974. See page 194.

Media Buying. Microsoft to surpass Yahoo in global digital ad market share this year: Google, Facebook continue on as the market's leaders. *eMarketer*, July 25 2014. <http://www.emarketer.com/Article/Microsoft-Surpass-Yahoo-Global-Digital-Ad-Market-Share-This-Year/1011012>.

National Organization for Women. Highlights. <http://now.org/about/history/highlights/>, 2017.

Nicholas Pedriana and Amanda Abraham. Now you see them, now you don't: The legal field and newspaper desegregation of sex-segregated help wanted ads 1965–75. *Law & Social Inquiry*, 31(4):905–938, 2006.

Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, pages 4:1–4:15, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1532-6. doi: 10.1145/2335356.2335362. URL <http://doi.acm.org/10.1145/2335356.2335362>.

4. 31.5 F.Supp. 577 (N.D. Cal. 1970).

5. 29 C.F.R. §1604.5 "It is a violation of title VII for a help-wanted advertisement to indicate a preference, limitation, specification, or discrimination based on sex unless sex is a bona fide occupational qualification for the particular job involved. The placement of an advertisement in columns classified by publishers on the basis of sex, such as columns headed "Male" or "Female," will be considered an expression of a preference, limitation, specification, or discrimination based on sex." 33 Fed. Reg. 11539 (1968).