On the Potential for Discrimination via Composition

Giridhari Venkatadri
Northeastern University

Alan Mislove
Northeastern University

ABSTRACT

The success of platforms such as Facebook and Google has been due in no small part to features that allow advertisers to target ads in a fine-grained manner. However, these features open up the potential for discriminatory advertising when advertisers include or exclude users of protected classes—either directly or indirectly—in a discriminatory fashion. Despite the fact that advertisers are able to *compose* various targeting features together, the existing mitigations to discriminatory targeting have focused only on individual features; there are concerns that such composition could result in targeting that is more discriminatory than the features individually.

In this paper, we first demonstrate how compositions of individual targeting features can yield discriminatory ad targeting even for Facebook's restricted targeting features for ads in special categories (meant to protect against discriminatory advertising). We then conduct the first study of the potential for discrimination that spans across three major advertising platforms (Facebook, Google, and LinkedIn), showing how the potential for discriminatory advertising is pervasive across these platforms. Our work further points to the need for more careful mitigations to address the issue of discriminatory ad targeting.

CCS CONCEPTS

• Security and privacy → Social aspects of security and privacy; **• Information systems → Online advertising**; **Social networks**.

KEYWORDS

Targeted advertising, Advertising platforms, Discriminatory ad targeting

ACM Reference Format:
Giridhari Venkatadri and Alan Mislove. 2020. On the Potential for Discrimination via Composition. In *ACM Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3419394.3423641

1 INTRODUCTION

Online advertising platforms such as Facebook, Google, and LinkedIn leverage their rich user databases to allow advertisers to target ads to particular users on their platforms. While the ability to selectively target relevant users is advantageous to advertisers—potentially offering them better value for their ad budget—such targeting raises concerns that advertisers could knowingly (or unknowingly) target users in order to selectively exclude users of certain sensitive populations (such as users of particular genders, ages, races, or other historically disadvantaged groups). Such discriminatory targeting, while concerning in and of itself, could also run afoul of law for advertisements related to housing, credit, and employment, where special legal protections exist [1–3].

This concern of discriminatory targeting was first raised in the context of Facebook's platform (the largest and most mature of these platforms), where it was shown that an advertiser could explicitly exclude users with certain "ethnic affinities" (such as African American) when targeting housing ads [16]. Subsequent research demonstrated that the problem was not limited to options that explicitly mentioned a protected class, and many additional options exist that are strongly correlated with protected classes [37]. In response to the uproar—including lawsuits from the National Fair Housing Alliance [14] and the U.S. Department of Housing and Urban Development (HUD) [15]—Facebook made a number of changes to its targeting options. These changes included deploying a restricted interface for housing, credit, and employment ads that has more limited targeting options [12].

Unfortunately, there are two key omissions in terms of understanding and protecting against discrimination in ad targeting. First, the previous discussion and proposed mitigations have focused on *individual* targeting options that happen to be correlated with a particular sensitive population; indeed, Facebook's above-mentioned restrictions to mitigate discriminatory advertising primarily focused on disabling access to many individual targeting options. However, these advertising platforms typically allow advertisers to *compose* multiple such options together in various ways. Thus, two (or more) targeting options that are individually only mildly correlated with a protected class (and therefore only mildly discriminatory), may end up being more significantly correlated when used in *conjunction* with each other. For example, the population interested in electrical engineering, or the population interested in sports cars, might each be somewhat skewed towards men; however, the population interested in electrical engineering *and* sports cars might be significantly more skewed. As a result, limiting individual targeting options alone may be insufficient to prevent discriminatory advertising. Indeed, if composing targeting options in general tends to yield more skew than individual targeting options, even an honest advertiser who uses multiple targeting options may end up inadvertently running an ad in a discriminatory manner.

The second key omission is that most of the focus in the press and academia has been on Facebook, and less attention has been paid thus far to the potential for similar discrimination via targeting on other advertising platforms. It is important to study other platforms, as they may offer different targeting options (and different methods of composition), driven by varying views of user data, and varying advertiser demands.

IMC '20, October 27–29, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8138-3/20/10...$15.00
https://doi.org/10.1145/3419394.3423641

To address this situation, this paper makes four key contributions: *First*, after providing background in § 2 and detailing our methodology in § 3, we show in § 4.1 how targeting compositions enable discriminatory targeting even given Facebook's significantly curtailed individual targeting options for housing, credit, or employment ads. *Second*, in § 4.2, we perform the first examination of the targeting options present on Google, and LinkedIn's platforms. We show the existence of individual targeting options that can be used to discriminate toward or against particular ages and genders on all platforms; this is especially concerning on a platform liked LinkedIn that focuses exclusively on users' employment-related needs. *Third*, in § 4.3, we show that (i) the composition of targeting options is a vector for abuse that could potentially affect *all* three platforms studied; (ii) combining targeting options generally tends to make them more discriminatory, indicating the potential for inadvertent discriminatory targeting by even well-meaning advertisers; and (iii) skewed compositions exist even when highly skewed individual options are removed. *Fourth*, while targeting compositions typically only let an advertiser reach a small fraction of a given protected class, our results indicate that (i) this limited reach is still large enough to suffice for most advertisers, and (ii) an advertiser could increase the fraction reached by targeting across multiple targeting compositions.

The existence of skewed targeting options and compositions is, in many cases, likely a reflection of platform users' interests and preferences. For example, users of a particular protected class might be more likely to find certain products relevant, and might not be interested in ads pertaining to certain other products. However, when attempting to prevent discriminatory advertising for ads in certain categories, our results underscore the need to carefully consider compositions of targeting options when designing mitigations; we discuss specific implications in § 5.

2 BACKGROUND

We now discuss ad platforms' targeting interfaces and features, and then overview related work.

2.1 Advertising platforms

Ad platforms offer a wide variety of targeting features to advertisers; here we focus on the features offered by Facebook, Google, and LinkedIn. The set of users resulting from a given set of targeting options is referred to as an *audience*.

Attribute-based targeting allows targeting by user attributes: in addition to age, gender, and location, all platforms support a default list of attributes that an advertiser can browse and choose from [21, 23, 38]. Additionally, these platforms support (potentially open) sets of custom attributes that advertisers can either search for (offered by all platforms), or define in a custom manner (offered by Google [23]).

Ad placement targeting allows targeting of where (or in what context) their ad appears. Google has the most extensive targeting options of this kind, allowing advertisers to specify which (first- or third-party) websites, apps, and videos to show ads on, either directly [25], or by specifying particular keywords/topics [24, 28]. Facebook and LinkedIn also provide (comparatively limited) targeting options of this kind [20, 33].

Activity-based targeting allows targeting based on visits or actions on advertisers' websites and apps [17, 22, 26].[1]

PII-based targeting allows targeting specific users by uploading personally identifying information (PII), such as names, and email addresses [7, 34, 42]. The platform then internally matches the PII and creates an audience of users.

Lookalike targeting allows advertisers to target sets of users similar to those in activity-based or PII-based targeting audiences they create [18, 27, 39].

Targeting compositions allow advertisers to combine targeting options of different kinds (via logical and), and they additionally support compositions via boolean rules even for multiple targeting options of the same kind [19, 30, 35].[2]

2.2 Facebook's restricted interface

In order to settle a lawsuit, Facebook introduced a restricted ad interface for ads in the protected areas of housing, employment, and credit; this interface has limited targeting options compared to the original interface [6]. Ages and genders cannot be targeted, a smaller list of targeting attributes is supported, and excluding users with particular attributes is disallowed. Additionally, Lookalike Audiences are replaced by "Special Ad Audiences" that are, according to Facebook, "adjusted to comply with the audience selection restrictions associated with your campaign's chosen Special Ad Category and our ad policies" [6]. PII-based, activity-based, and ad placement targeting are available, however.

2.3 Related work

The potential for discrimination via individual targeting attributes (even facially neutral ones) was first demonstrated in the context of Facebook by Speicher et al. [37], subsequent to the demonstration by ProPublica [16] of the possibility of explicitly excluding users with certain "ethnic affinities" when targeting housing ads on Facebook. Our work demonstrates that this problem is not limited to Facebook, and is made worse due to platforms' support for targeting compositions.

Other work [4, 5, 11, 31, 32, 40] has demonstrated and discussed the implications of skewed outcomes (across races, genders, and political affiliations) arising from the working of the ad platforms' delivery mechanism (rather than from targeting). In addition, prior work demonstrated that particular real-world ads were delivered by Google's platform to users in a skewed manner [11, 36], without inferring specific causes for this skewed outcome, while Datta et al. [10] explored the legal implications of various potential causes for this skewed outcome.

3 METHODOLOGY

We next describe our methodology, which directly uses the ad platforms' targeting features to measure the skew of different targetings.

[1] In order to do so, the advertiser places a tracking pixel from the ad platform on their website, letting the website track visitors' actions.

[2] The form of these boolean rules varies across ad platforms and kinds of targetings; the rules only accommodate boolean-or in some cases, while they could be and of or-terms in other cases.

We focus on the sensitive attributes gender and age, as ad platforms typically have access to these and offer options to explicitly target these attributes.[3]

Metrics To account for varying underlying distributions of users across different sensitive populations and platforms, we use a metric called the *representation ratio* [37], inspired by the disparate impact metric historically used to detect discrimination in employment and housing allocation [9].

This metric focuses only on the (implicit) audience $\boldsymbol{RA}$ of users who might find a given ad relevant; within this audience, it measures whether users from a given sensitive population $\boldsymbol{RA_s}$ (represented by a value s of a corresponding sensitive attribute) are more (or less) likely to be included in a given audience $\boldsymbol{TA}$ targeted by an advertiser, compared to users with a different value of the sensitive attribute $\boldsymbol{RA_{\neg s}}$:

$$rep\_ratio_s(\boldsymbol{TA}, \boldsymbol{RA}) = \frac{|\boldsymbol{TA} \cap \boldsymbol{RA_s}|/|\boldsymbol{RA_s}|}{|\boldsymbol{TA} \cap \boldsymbol{RA_{\neg s}}|/|\boldsymbol{RA_{\neg s}}|}, \tag{1}$$

For the purposes of this paper, in line with prior work [37], we assume $\boldsymbol{RA}$ is the set of all U.S.-based users (and thus $\boldsymbol{RA_s}$ is the set of all U.S.-based users with a value s for the sensitive attribute). Thus, a representation ratio of 1 is ideal and means users from $\boldsymbol{RA_s}$ and $\boldsymbol{RA_{\neg s}}$ are equally likely to be included in the targeted audience; however, an unacceptably high (or low) value could be 1.25 or above (or 0.8 and below), as per the well-known four-fifths rule [8] for measuring disparate impact, indicating over- or under-representation (respectively) of the given sensitive population [37].

In addition, we measure the *recall* of the ad targeting, which we define as $|\boldsymbol{TA} \cap \boldsymbol{RA_s}|$ when the targeting selectively includes users from $\boldsymbol{RA_s}$, and as $|\boldsymbol{TA} \cap \boldsymbol{RA_{\neg s}}|$ when the targeting selectively excludes users from $\boldsymbol{RA_s}$ (i.e, includes users from $\boldsymbol{RA_{\neg s}}$). We next briefly describe how we target the different audiences in Equation 1, and how we measure the sizes of these audiences.

Targeting audiences To target the audiences in Equation 1, we leverage the fact that the studied ad platforms allow targeting by location, gender, and age, in addition to any other fine-grained targeting options.[4] While Facebook's restricted interface does not allow targeting by gender or age, we instead use the corresponding targeting option on Facebook's normal interface to measure the representation ratio.

To measure $|\boldsymbol{R_s}|$, we target all U.S. users, and additionally target users who have a value s for the given sensitive attribute. To measure $|\boldsymbol{TA} \cap \boldsymbol{RA_s}|$, we further add in the targeting options corresponding to $\boldsymbol{TA}$. For the given sensitive attribute (age or gender), we measure $|\boldsymbol{R_s}|$ and $|\boldsymbol{TA} \cap \boldsymbol{RA_s}|$ as above for each value of s; we then compute $|\boldsymbol{RA_{\neg s}}|$ as $\Sigma_{s' \in \neg s}|\boldsymbol{RA_{s'}}|$, and compute $|\boldsymbol{TA} \cap \boldsymbol{RA_{\neg s}}|$ as $\Sigma_{s' \in \neg s}|\boldsymbol{TA} \cap \boldsymbol{RA_{s'}}|$.

On these ad platforms, we select the campaign objective of "Reach" in order to reach the largest set of people.[5] Besides, on Google (which allows different types of ad campaigns), we focus on the "Display" campaign type as it covers Google's entire ad network and corresponds to the broadest reach.

Measuring audience sizes To measure the sizes of these audiences, we leverage the audience size estimates provided by the ad platforms' targeting interfaces.These numbers are intended to aid advertisers when they are selecting targeting options, and they give a measure of the size of the audience resulting from a given targeting. While the estimate provided by Google's ad platform, as per the UI, is the "estimated number of impressions that your settings and targeting could theoretically reach", the estimates provided by Facebook and LinkedIn measure the count of users in the audience ("the size of the audience that's eligible to see your ad", and "the number of LinkedIn members who match your targeting criteria", according to the respective interfaces). We find that the estimated number of impressions on Google's ad platform depends on a frequency capping setting which restricts how often an ad is shown to the same user [29]; we set the setting to its most restrictive value (one impression across the campaign every month per-user).

Automating size queries We use our browser's web inspector tool to identify the underlying API calls made by the targeting UIs whenever the selected set of targeting options is altered; we then automate these calls with a Python script. While the API calls made by Facebook and LinkedIn are unobfuscated, the API calls made by Google consist of obfuscated json; by manually varying the targeting options systematically, we find a mapping between the targeting options and particular keys and values in the obfuscated json.

Understanding size estimates Since audience size estimates have been shown not to be exact size estimates in the context of Facebook [41], we study the *granularity* and *consistency* of these to understand if they are obfuscated in any way. In brief, we use 100 back-to-back repeated calls for 20 random targeting options and 20 random compositions and find that across all three platforms, the returned estimates are consistent.[6]

To study the granularity of the estimates, we combine the results of over 80,000 various distinct API calls we make for each ad platform (spanning a variety of ad targetings), and find that the size estimates across all the platforms are granular: while Facebook's estimates have two significant digits (with a minimum returned value of 1,000); Google's estimates have one significant digit (until 100,000), and two significant digits thereafter; LinkedIn's estimates on the other hand have two significant digits (starting at 300). [7] Such rounding could mean the measured representation ratios (based on the rounded estimates) could be either higher or lower than the actual representation ratio (corresponding to the exact audience sizes). However, we confirm that even allowing for the representation ratios to take their least skewed values (subject to the rounding ranges), we find very similar degrees of skew in targetings across our experiments.

Obtaining targeting options As discussed in § 2, each of the three ad platforms we study provides a plethora of targeting options. However, to limit the number of possible targeting options

[3]For age, we consider the age ranges 18-24, 25-34, 35-54, and 55+, as these are most granular targeting options common to the three ad platforms we study.

[4]While LinkedIn does not have separate targeting options for targeting by gender, or age, its list of detailed attribute-based targeting options includes user genders and age ranges; these detailed targeting options can be combined by performing a logical-and of a series of logical-or terms; thus, for a given ad targeting, we additionally target a particular gender or age range by adding the corresponding targeting attribute via a logical-and.

[5]On Google and LinkedIn, we select the closest corresponding objectives "Brand awareness and reach" and "Brand awareness" respectively.

[6]It is possible that an ad platform may obfuscate the audience size corresponding to a given ad targeting by always adding the same specific noise sample.

[7]While Facebook's estimates had a minimum of 1,000; Google's and LinkedIn's estimates had a minimum of 40 and 300 respectively, with 0 returned below that minimum.

we need to study, we focus on the default list of user attributes for attribute-based targeting on each platform; for Google, in addition, we consider the default list of topics on its topic targeting feature (that lets an advertiser place ads solely on webpages corresponding to a particular topic).[8] We collect 393 and 667 attributes for Facebook's restricted and normal interface, respectively; 873 attributes and 2,424 topics for Google; and 552 attributes for LinkedIn.

Discovering the most skewed compositions To limit the query load, we avoid an exhaustive crawl and use a greedy approach to discover an approximate (lower bound) set of most skewed targeting compositions; the method simply greedily combines the most skewed individual targetings. Specifically, we approximate the 1,000 most skewed pairwise targeting compositions by combining pairwise (via a logical and) the 46 most skewed individual attributes, resulting in 1,035 pairs, and then randomly sampling.[9] To avoid very niche targetings, our method only considers individual targetings and pairs with a total reach of at least 10,000.

Limitations Our methodology has a number of limitations. *First*, our results rely on and are subject to the quality of the ad platforms' sensitive attribute data. *Second*, the size statistics provided by these ad platforms might be affected by the presence of fake accounts, or by users owning multiple accounts. *Third*, while we measure the skew in audiences arising from targeting, the operation of the ad platform's ad delivery system might introduce additional skews [4].

4 EXPERIMENTS

We now present our results, starting first with Facebook's restricted interface and then examining multiple ad platforms.

4.1 Facebook's restricted interface

We first motivate our work by using our methodology to investigate whether targeting compositions could exacerbate the potential for discrimination on Facebook's restricted interface for ads in protected categories. For each set of targetings, we use box-plots to plot the distribution of representation ratios corresponding to males, and corresponding to users of ages 18-24 respectively in Figure 1 (and corresponding to other age ranges in Figure 4 of Appendix A).[10] Throughout the paper, to avoid focusing on very niche targetings, we only show results for targetings that have a total recall 10,000 or more.

Individual targeting For gender, we focus on the Individual column in the first box plot of the first figure in Figure 1. We see that the set of 393 targeting attributes that Facebook offers on its restricted interface show some evidence of gender skew: the targetings with the 90th and 10th percentile representation ratios corresponding

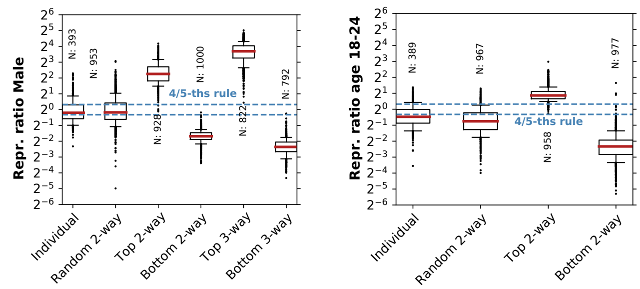

Figure 1: Distributions of representation ratios corresponding to males (left) and to ages 18–24 (right), for different sets of targetings on Facebook's restricted interface. Only targetings with a total recall above 10,000 are included.

to males are 1.84 (i.e., a male is nearly twice as likely to be picked as a female) and 0.5 (i.e., a female is twice as likely to be picked as a male), respectively. For age, again focusing on the Individual column, we observe similar results for all the age ranges studied: The 10th and 90th percentile representation ratios for the 18–24 age range is 0.39 and 1.39, respectively.

However, despite the presence of some skewed attributes, the interface is still a highly sanitized interface: the interface excludes a large number of other highly skewed individual targeting attributes observed [37] in Facebook's normal interface, including hundreds of thousands of free-form attributes (e.g., Interested in Marie Claire, which has a 0.08 representation ratio towards males).

Compositional targeting We next study whether compositions of individual targetings (via a logical and) on the sanitized interface could exacerbate the potential for discriminatory advertising. To do so, we select 1,000 random pairs of targeting attributes, referred to as "Random 2-way"; we also use the approach in § 3 to discover the top 1,000 pairs of targeting attributes most skewed towards ("Top 2-way"), and against ("Bottom 2-way") the given sensitive population.

We compare the resulting skew of these compositions to the skew exhibited by individual targeting options in Figure 1. We first observe that the random pairs of attributes often lead to more skewed distributions, such as with the 18–24 age range (where combinations tend to make the resulting audience even more skewed away the 18–24 group).

To see the extremes of this effect, we focus on the most skewed combinations. The sets of "Top 2-way" and "Bottom 2-way" targetings show additional skew, with 10th percentiles reaching as low as 0.1 and 90th percentiles as high as 8.98. Thus, while the average combination of two targeting options shows modest additional skew (compared to "Individual"), many outlying combinations show significant additional skew. For example, targeting users interested in Electrical engineering and Cars yielded a representation ratio of 12.43 towards males, while targeting each individual attribute yielded smaller representation ratios (of 3.71 and 2.18 respectively). We show other illustrative examples of "Top 2-way" targetings showing how composition can increase skew, in Tables 2 and 3 in Appendix A.

Finally, we focus how this effect scales by repeating the experiment with three targetings composed instead of two (creating "Top 3-way" and "Bottom 3-way" targetings) for gender in Figure 1. We

[8]We only consider user-attribute-based targeting for Facebook and LinkedIn as both platforms allow the composition of these targeting attributes via boolean rules. However, while Google allows boolean combinations of these attributes for specific kinds of ad campaigns (related to its search products), it does not show audience size statistics for these; when audience size statistics are available, user attributes can only be combined via a logical or. Thus, to be able to demonstrate the potential of logical and-based composition, we additionally consider the topic based targeting feature.

[9]In the case of Google, where we study compositions between targeting options belonging to two different targeting features, targeting options within the same feature cannot be composed. Thus, the number of skewed individual options from each feature necessary to obtain 1,000 skewed compositions will vary from case to case and has to be computed in each case.

[10]Each box plot shows the median representation ratio (as a thick red line), the 25-th and 75-th percentiles (as the edges of the box), the 10-th and 90-th percentiles (as the whiskers), and representation ratios in the top and bottom 10 percentiles as outlier points. The representation ratio thresholds (1.25 and 0.8) described in Section 3 corresponding to the four-fifths rule for detecting disparate impact are shown for reference.

find that the skew is indeed amplified further: the 90th percentile representation ratio for the "Top 3-way" targetings is 19.77 and the 10th percentile representation ratio for the "Bottom 3-way" targetings is 0.11, implying a further increase in the degree of skew.

Summary Our findings show that compositions of targeting options can be abused to target skewed sets of users, exhibiting a greater degree of skew towards (or away from) particular ages and genders compared to individual targeting options. That this is true in the context of a highly sanitized advertising interface—defined by the settlement of a lawsuit that was focused on discriminatory ad targeting—indicates that significant additional work is needed to ensure these systems cannot be used for discriminatory advertising. We use these findings as motivation to explore the extent to which the same is true on Facebook's full interface, as well as other platforms' interfaces.

4.2 Individual targeting

We now examine the skews arising from targeting the various default targeting attributes on each of these platforms individually. Focusing on the Individual column in the three sections of Figure 2 (and of Figure 4 of Appendix A), we make *two* observations: *First*, attributes from different platforms show varying distributions of skew. For example, LinkedIn's targeting attributes are generally more skewed towards males, with a 90th percentile representation ratio of 2.09; by contrast, Facebook's targeting attributes are more skewed toward females, with a comparatively lower 90th percentile representation ratio (toward males) of 1.45. On the other hand, Google's and LinkedIn's targeting attributes are generally more skewed away from the youngest users (ages 18–24), and skewed toward the oldest users (ages 55+). These systematic skews could potentially be due to various factors such as users' activities on these platforms, the data these platforms collect about users, advertiser needs, etc. *Second*, in all cases, we concerningly see that there exist a number of skewed targeting attributes that may violate the four-fifths rule; such skew is especially concerning in the context of LinkedIn which focuses on employment.

4.3 Compositional targeting

Next, we explore what happens with composition.

Potential for discrimination We first perform the experiments from § 4.1 on different ad platforms, plotting the distribution of representation ratios for random pairs of attributes, and for the 1,000 most skewed pairs, in Figure 2 (and in Figure 4 of Appendix A).

Across ad platforms, we find that randomly chosen pairs of targeting attributes show modest additional skew, for example, exacerbating the skew against smaller ages (18–24) on LinkedIn. This is concerning as it means that even honest advertisers using targeting compositions could be likely to be targeting users in a more skewed manner. Additionally, the most skewed pairs of targeting attributes clearly indicate the exacerbated potential for discrimination from composition, with over 90 percent of these falling outside the thresholds of the four-fifths rule.

Recall of targeting compositions We next study if an advertiser could selectively reach a large number of users (i.e., achieve a high recall) of a particular sensitive population using the previous highly

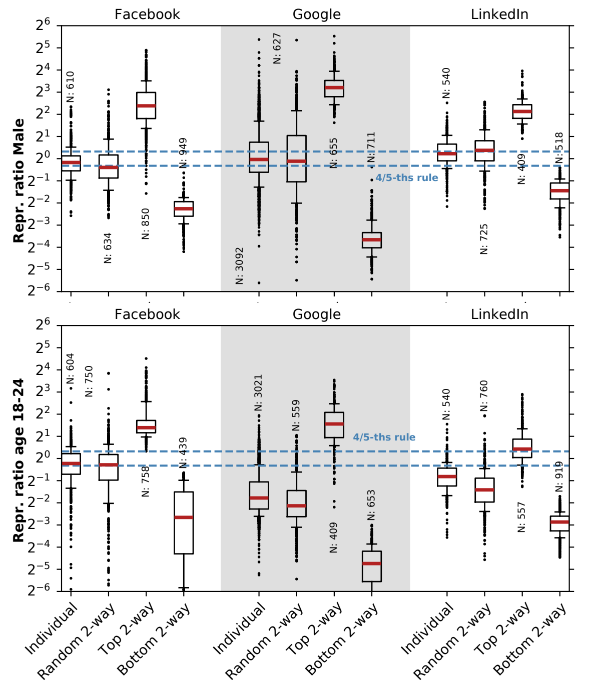

Figure 2: Distributions of representation ratios corresponding to males (top) and ages 18–24 (bottom) on different platforms.

skewed targeting pairs. For each set of targetings previously studied, we take skewed targetings in the set that fall outside the thresholds of the four-fifths rule, and plot the distribution of corresponding recalls of the given sensitive population. For reference, we also show the distribution of recalls for all individual targeting options, and for skewed individual targeting options. While we show results for gender and various age ranges in Figure 5 in Appendix A, we focus on results for females here.

We find that while skewed pairwise targeting compositions have substantial recalls, these typically correspond to small fractions of the overall target sensitive population on the platform. For example, the 90th percentile recall for the "Top 2-way" skewed compositions is 5M (4.17%), 30M (25%), 1.7M (0.14%), and 560K (0.79%) respectively for Facebook's restricted interface, Facebook's full interface, Google, and LinkedIn respectively; the respective median recalls are 570K (0.47%), 1.9M (1.58%), 170K (0.01%), and 46K (0.06%). However, since most advertisers on these platforms only spend up to a few hundreds of dollars *per ad* on average, with tens of thousands of impressions [13], these recalls may still be appealing. Besides, we find that even individual targeting options tend to only have niche recalls: the respective median recalls for all individual options across the four interfaces are 3.2M (2.67%), 5.2M (4.33%), 11M (0.92%), and 1.4M (1.97%). Finally, we unsurprisingly observe that targeting compositions tend to achieve lower recalls than individual targeting options (achieving lower median recalls, for example).

Increasing recall We then study whether an advertiser could increase their recall even further by targeting ads across *multiple* skewed compositions; this depends on the degree of overlap between the corresponding audiences. While we focus on results for females here, more complete results for this analysis are in Table 1 in Appendix A. We measure the pairwise overlaps between the sets

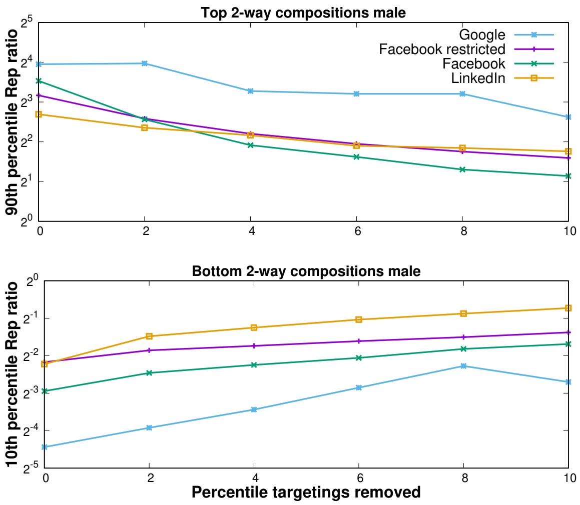

Figure 3: Effect of removal of the most skewed individual targetings on the skew of pairwise targeting compositions.

of females reached by the top 100 female-skewed targeting compositions,[11] and find a median pairwise overlap of approximately 22%, 15%, and 0% for Facebook's restricted interface, Facebook's full interface, and LinkedIn respectively;[12] this low overlap indicates the potential to increase recall further.

We further confirm that this is the case by estimating the total recall of males across the union of the top 10 male-skewed compositions.[13] While the top female-skewed composition on Facebook's restricted interface, Facebook's full interface, and LinkedIn respectively had a recall of 1.1M (0.9%), 270K (0.2%), and 28K (0.0%); the total recall across the top 10 female-skewed compositions was significantly higher, i.e., 6.1M (5.1%), 4M (3.3%), and 1.1M (1.6%) respectively.

Removing skewed individual targetings We finally study whether removing the most skewed individual targeting attributes is sufficient to mitigate against skew in targeting compositions; for each sensitive population, we successively remove the most skewed individual targeting attributes in steps of two percentile, and use the greedy method as before to obtain the "Top 2-way" and "Bottom 2-way" sets of most skewed compositions.[14] We plot the resulting variation in representation ratio for gender (males) in Figure 3; we obtain similar results for different age ranges in Figure 6 of Appendix A.

We observe that the removal of the most skewed individual targetings leads, perhaps unsurprisingly, to a drop in the skew of the corresponding compositions of targetings across platforms.[15] *However*, the compositions of the remaining targeting attributes still yields highly skewed rules; for example, even with the removal of the top 10th percentile of male-skewed individual attributes for Facebook's restricted interface, the 90th percentile of resulting "Top 2-way" representation ratios was 3.02, and the highest resulting representation ratio was 5.23.

5 CONCLUDING DISCUSSION

In this paper, we performed the first study that showed that the potential for discriminatory targeting, previously reported in the context of Facebook, exists across multiple ad platforms. Moreover, ad platforms allow advertisers to compose individual ad targeting options; we showed how such composition exacerbates the potential for intentional or unintentional discriminatory targeting, even for a highly sanitized ad interface such as Facebook's restricted interface.

Mitigations While prior work [37] showed that disabling the use of obvious stereotypically skewed targeting attributes was insufficient (as there could be facially neutral attributes that are still skewed), this paper further shows that even an approach based on removing *all* highly skewed individual targeting attributes is also likely insufficient. Thus, our work re-enforces the need to base mitigations against discriminatory advertising on the *outcome* of the targeting [37], rather than on the targeting itself. For example, especially for protected categories, ad platforms could potentially use anomaly detection based on the outcome of ad targeting to detect advertisers who consistently target skewed audiences. Any flagged advertisers could then be subject to further review about whether their use of targeting options is justifiable.

Our work specifically shows the need to base mitigations upon the outcome of the composition of targetings specified by an advertiser, rather than upon the outcomes of the individual targeting options used. At the very least, while restricting or disabling the use of skewed targetings, ad platforms should be more aggressive keeping targeting compositions in mind. In general, our work motivates the need to not be myopic and simply focus on individual components of ad platforms when designing mitigations against discriminatory advertising.

Ethics While conducting this work, we carefully considered the ethical issues and took care to ensure our work was consistent with best practices. We did not collect any individual user data; rather, we only collected high-level, obfuscated audience size statistics provided by the ad platforms to all advertisers. Additionally, our experiments did not impact users directly, as we did not run ads. We also minimized the load placed on the ad platforms by limiting both the count and rate of API queries we make.

6 ACKNOWLEDGEMENTS

We thank the anonymous reviewers and our shepherd, Marco Mellia, for their valuable comments. This research was supported in part by NSF grant CNS-1916020. Additionally, Giridhari Venkatadri acknowledges support from a Facebook Fellowship.

[11] We measure these overlaps on Facebook and LinkedIn by exploiting their support for boolean combinations (logical-and of logical-ors) of targeting attributes. We could not conduct a similar analysis on Google because, as previously mentioned, Google does not provide audience size statistics when targeting such boolean combinations.

[12] The overlaps were conservatively measured by comparing the size of the intersection to the size of the smaller set in the pair.

[13] Directly measuring the total recall of the union (logical-or) of a number of compositions (logical-ands) needs a logical-or of logical-ands. Since Facebook and LinkedIn only support a logical-and of logical-ors, we instead indirectly estimate total recall by combining multiple logical-and queries, using the inclusion-exclusion principle.

[14] As before, we only consider the subset of each set of 1K targeting compositions that has a total reach of at least 10,000.

[15] The minor deviations from monotonicity in the curve for Google could potentially be due to our greedy method giving us only an approximate set of the most skewed targeting compositions.

REFERENCES

[1] 12 CFR § 202.4 (b) – Discouragement. https://www.law.cornell.edu/cfr/text/12/202.4.
[2] 24 CFR § 100.75 – Discriminatory advertisements, statements and notices. https://www.law.cornell.edu/cfr/text/24/100.75.
[3] 29 USC § 623 – Prohibition of age discrimination. https://www.law.cornell.edu/uscode/text/29/623.
[4] M. Ali, P. Sapiezynski, M. Bogen, A. Korolova, A. Mislove, and A. Rieke. Discrimination through Optimization: How Facebook's Ad Delivery can Lead to Biased Outcomes. *CSCW*, 2019.
[5] M. Ali, P. Sapiezynski, A. Korolova, A. Mislove, and A. Rieke. Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging. *arXiv preprint arXiv:1912.04255*, 2019.
[6] About Audiences for Special Ad Categories. https://www.facebook.com/business/help/2220749868045706.
[7] About Customer Match. https://support.google.com/adwords/answer/6379332?hl=en.
[8] D. Biddle. *Adverse Impact and Test Validation: A Practitioner's Guide to Valid and Defensible Employment Testing*. Gower, 2005.
[9] S. Barocas and A. D. Selbst. Big data's disparate impact. *Cal. Law Rev.*, 104, 2016.
[10] A. Datta, A. Datta, J. Makagon, D. K. Mulligan, and M. C. Tschantz. Discrimination in Online Personalization: A Multidisciplinary Inquiry. *FAT**, 2018.
[11] A. Datta, M. C. Tschantz, and A. Datta. Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination. *PETS*, 2015.
[12] E. Dreyfuss. Facebook Changes Its Ad Tech to Stop Discrimination. *WIRED*. https://www.wired.com/story/facebook-advertising-discrimination-settlement/.
[13] L. Edelson, S. Sakhuja, R. Dey, and D. McCoy. An Analysis of United States Online Political Advertising Transparency. https://arxiv.org/abs/1902.04385.
[14] Exhibit a – Programmatic Relief. https://nationalfairhousing.org/wp-content/uploads/2019/03/FINAL-Exhibit-A-3-18.pdf.
[15] Facebook Engages in Housing Discrimination with Its Ad Practices, U.S. Says. https://www.nytimes.com/2019/03/28/us/politics/facebook-housing-discrimination.html.
[16] Facebook Lets Advertisers Exclude Users by Race. https://www.propublica.org/article/facebook-lets-advertisers-exclude-users-by-race/.
[17] Facebook: About Facebook Pixel. https://www.facebook.com/business/help/742478679120153.
[18] Facebook: About Lookalike Audiences. https://www.facebook.com/business/help/164749007013531.
[19] Facebook: Best Practices for Detailed Targeting. https://www.facebook.com/business/help/905095143159925.
[20] Facebook: Brand Safety Controls. https://www.facebook.com/business/help/1926878614264962.
[21] Facebook: Use Detailed Targeting. https://web.archive.org/web/20200429194031/https://www.facebook.com/business/help/440167386536513.
[22] Get Started with LinkedIn Website Retargeting. http://web.archive.org/web/20200601191411/https://www.linkedin.com/help/lms/answer/73934.
[23] Google: About audience targeting. https://support.google.com/google-ads/answer/2497941?hl=en.
[24] Google: About contextual targeting. https://support.google.com/google-ads/answer/2404186?hl=en.
[25] Google: About managed placements. https://support.google.com/google-ads/answer/2470108?hl=en.
[26] Google: About remarketing setup. https://support.google.com/google-ads/answer/2454000?hl=en.
[27] Google: About similar audiences on the Display Network. https://support.google.com/google-ads/answer/2676774?hl=en.
[28] Google: About topic targeting. https://support.google.com/google-ads/answer/2497832?hl=en.
[29] Google: Frequency capping: Definition. https://support.google.com/google-ads/answer/117579?hl=en.
[30] Google: How to use combined audiences. https://support.google.com/google-ads/answer/9066030?hl=en.
[31] HUD Sues Facebook Over Housing Discrimination and Says the Company's Algorithms Have Made the Problem Worse. https://www.propublica.org/article/hud-sues-facebook-housing-discrimination-advertising-algorithms.
[32] A. Lambrecht and C. E. Tucker. Algorithmic bias? An empirical study into apparent gender-based discrimination in the display of STEM career ads. 2018. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2852260.
[33] LinkedIn Audience Network – Manage Delivery Preferences. https://www.linkedin.com/help/linkedin/answer/83630.
[34] LinkedIn Contact Targeting. https://business.linkedin.com/marketing-solutions/ad-targeting/contact-targeting.
[35] LinkedIn: Use and-OR Targeting for Objective-Based Advertising Campaigns. https://www.linkedin.com/help/lms/answer/102905/use-and-or-targeting-for-objective-based-advertising-campaigns?lang=en.
[36] L. Sweeney. Discrimination in online ad delivery. *CACM*, 56(5), 2013.
[37] T. Speicher, M. Ali, G. Venkatadri, F. N. Ribeiro, G. Arvanitakis, F. Benevenuto, K. P. Gummadi, P. Loiseau, and A. Mislove. On the Potential for Discrimination in Online Targeted Advertising. *FAT**, 2018.
[38] Targeting Options for LinkedIn Advertisements. https://web.archive.org/web/20200525174210/https://www.linkedin.com/help/lms/answer/722/targeting-options-for-linkedin-advertisements?lang=en.
[39] Targeting with LinkedIn Lookalike Audiences – Overview. https://www.linkedin.com/help/lms/answer/94287.
[40] Upturn Amicus Brief in Onuoha v. Facebook. https://www.courtlistener.com/recap/gov.uscourts.cand.304918/gov.uscourts.cand.304918.76.1.pdf.
[41] G. Venkatadri, Y. Liu, A. Andreou, O. Goga, P. Loiseau, A. Mislove, and K. P. Gummadi. Privacy Risks with Facebook's PII-based Targeting: Auditing a Data Broker's Advertising Interface. *IEEE S&P*, 2018.
[42] What's a Custom Audience from a Customer List? https://www.facebook.com/business/help/341425252616329/.

A SUPPORTING RESULTS

In this Appendix, we present some additional results supporting those in the paper.

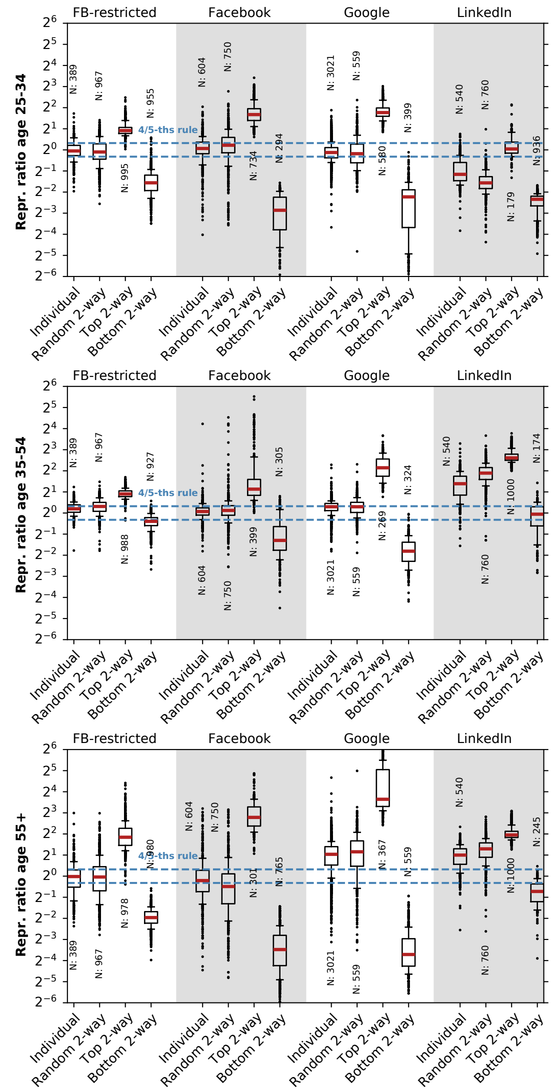

Figure 4: Distributions of representation ratios for different sets of targetings and different ad platforms, across multiple age ranges.

Skew across age ranges While we presented results for the skew relative to the youngest users considered (ages 18-24) in Figures 1 and 2 of the main body of the paper, we now present the corresponding results for older users (ages 25-34, 35-54, and 55+) in Figure 4.

We observe similar results as before across ad platforms: while even the individual targeting attributes contain highly skewed attributes, this is typically moderately exacerbated even for randomly chosen pairs of targeting attributes, and even more exacerbated when the most skewed pairs are considered. As with younger users, we see that we can effectively exclude older users (for example, users on LinkedIn aged 55+) via targeting compositions.

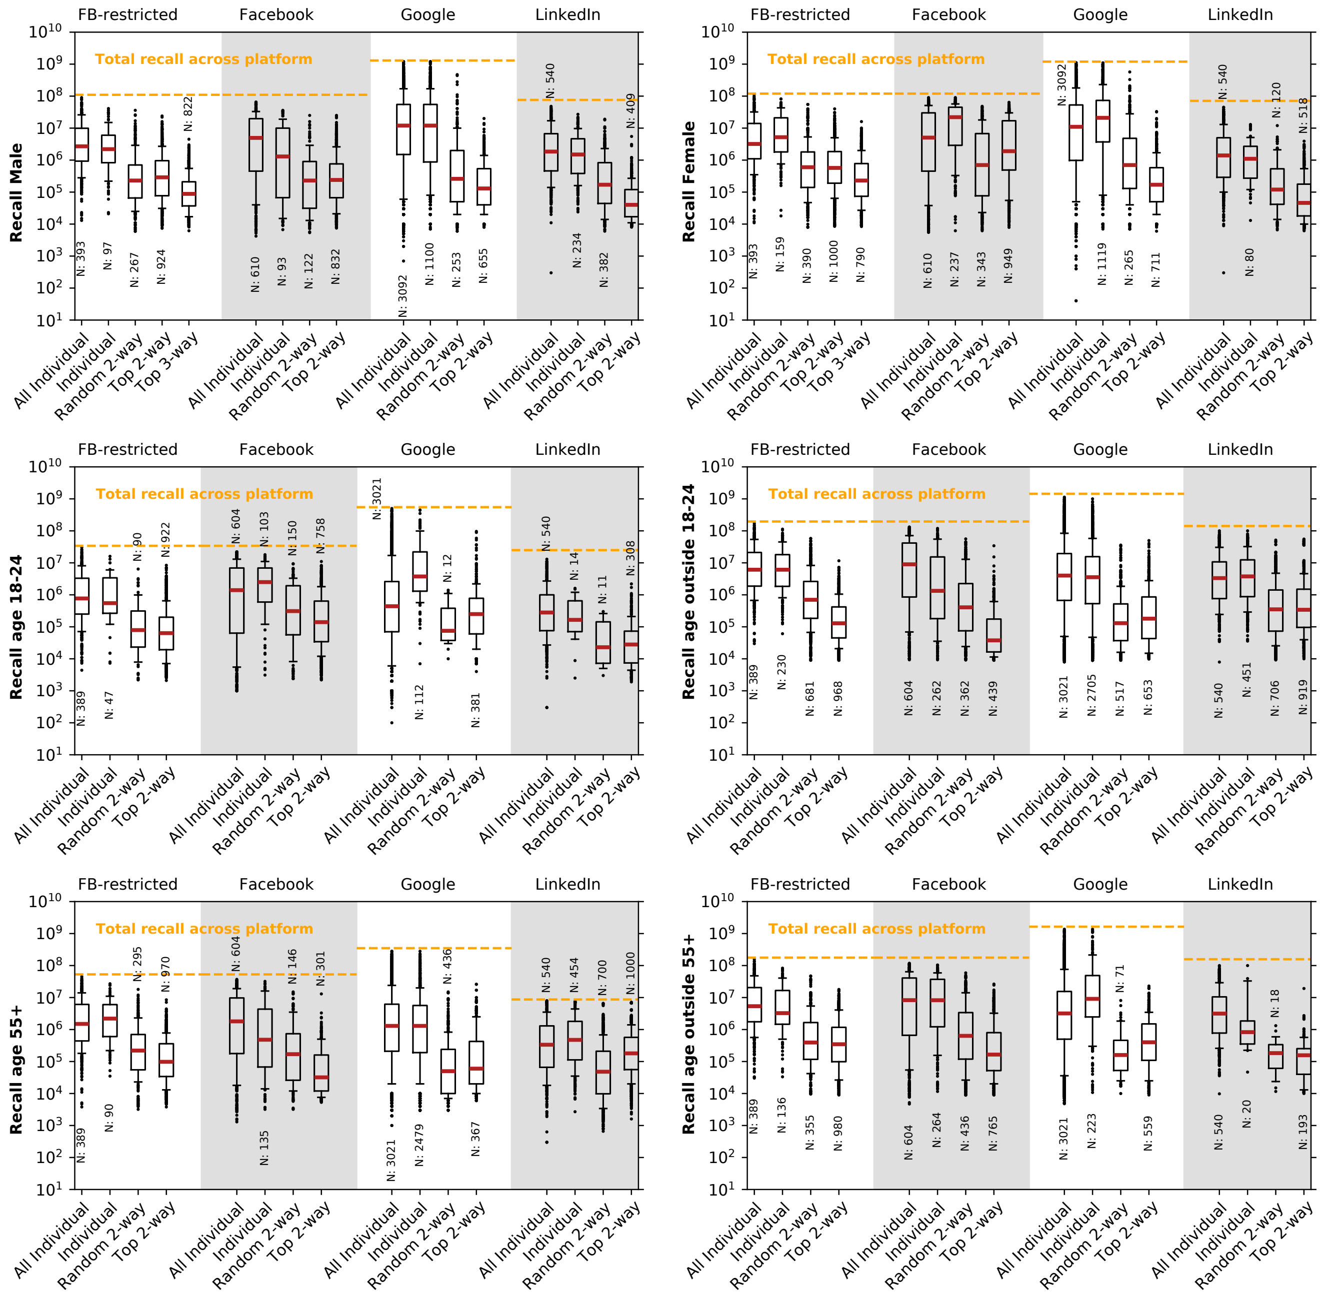

Figure 5: Distributions of recalls of a particular protected class for different sets of skewed targetings (exceeding the four-fifths thresholds) skewed toward that protected class. Results are shown for different ad platforms, and for different genders and ages, with the distribution across all individual targetings also shown for reference. Also shown is the total size of the given sensitive population for each ad platform.

Recall of sensitive populations We next complement our discussion in Section 4.3 of the recall achieved by both skewed targeting compositions; while that discussion focused on females, we now additionally present results for males, as well as for various age ranges.

As described in the paper, we first plot the distribution of the recalls of particular genders and ages achieved by various sets of skewed targetings in Figure 5. For reference, we also show the total size of the given sensitive population for each ad platform. We see results similar to those noted in Section 4.3: across genders and ages, the median recall of pairs of targeting attributes is substantial (while corresponding only to a niche percentage of the sensitive population), but considerably lower than the median recall of individual attributes. Additionally, we observe that the median recall of skewed targetings excluding particular age ranges are in most cases higher than the median recall of skewed targetings including the same particular age ranges; this is expected as exclusions of particular age ranges correspond to a wider range of ages (for the age ranges we study).

Similar to the discussion in Section 4.3, we next study whether an advertiser can achieve increased recall of a given sensitive population by running ads across multiple skewed compositions, effectively performing a logical or of the individual audiences. We first explore the potential for this increased recall by presenting the median overlap between pairs of the top 100 skewed targeting compositions (favouring various genders and ages) in the first section of Table 1; the overlap for each pair of audiences is measured as the percentage of the smaller audience that falls in the other audience. We find that the overlaps are small (often in single digits), with the largest median overlap (For Facebook's restricted interface, corresponding to female-skewed compositions) being just 22.58%; this indicates the potential, across sensitive populations and ages, for increased recall across multiple skewed targeting compositions.

Finally, as in Section 4.3, we further quantify this potential for increased recall by using the inclusion-exclusion principle to estimate the total recall obtained by combining the top 10 most skewed targeting compositions (favouring various sensitive populations). While the coarse granularity of size estimates could affect the result of adding these estimates as per the inclusion-exclusion principle, we confirmed that the estimated recalls converged as we successively added the higher-order terms specified by the inclusion-exclusion principle. We compare the recall of the most skewed composition, to the recall achieved by the top 10 compositions together, in the second section of Table 1. We see a total recall of over a million in most cases; for example, using the top 10 most skewed compositions, an advertiser on LinkedIn can achieve a total recall of over 5 million when excluding younger users (ages 18-24), and over a million when excluding older users (ages 55+).

Removing skewed individual targetings We finally extend the discussion in Section 4.3 on the impact of removing highly skewed individual targetings on the skew of the resulting targeting compositions; while that discussion focused on genders, here we present results for various age ranges. As before, for a given sensitive demographic, we remove the most skewed individual targetings in steps of 2 percentile, and discover the resulting sets of most skewed compositions using the greedy method; we plot the corresponding variation in skew for different age ranges in Figure 6.

We observe that in most cases, the removal of even the top 10 percentile most skewed individual attributes is insufficient to mitigate skew in the resulting targeting compositions. While in some cases the 90th percentile representation ratio for pairwise compositions does reduce to within the bounds of the four-fifths rule, such as for selectively including younger users of ages 18-24 on LinkedIn, even in these cases, higher degrees of targeting compositions could potentially again enable highly skewed ad targeting.

Favoured population	FB-restricted	Facebook	LinkedIn	FB-restricted		Facebook		LinkedIn	
	Median overlap			Top-1	Top-10	Top-1	Top-10	Top-1	Top-10
Male	17.33%	7.14%	0.00%	640K (0.6%)	4,232K (3.8%)	310K (0.3%)	808K (0.7%)	6K (0.0%)	169K (0.2%)
Female	22.58%	15.45%	0.00%	1,100K (0.9%)	6,136K (5.1%)	270K (0.2%)	4,014K (3.3%)	28K (0.0%)	1,122K (1.6%)
Age not 18-24	8.97%	2.22%	14.21%	1,099K (0.6%)	2,126K (1.1%)	15,305K (7.8%)	24,788K (12.6%)	15K (0.0%)	5,172K (3.7%)
Age not 55+	21.83%	3.14%	0.00%	188K (0.1%)	2,723K (1.5%)	301K (0.2%)	2,109K (1.2%)	11K (0.0%)	1,225K (0.8%)

Table 1: Exploring the potential for increased recall of a target sensitive population by targeting ads across multiple skewed audiences (corresponding to skewed targetings). For each sensitive population, the first three columns show the median pairwise intersection size between the top 100 most skewed targeting audiences toward that population. The remaining columns show the recall achieved by the top skewed targeting audience, and for the estimated (using the inclusion-exclusion rule) total recall for combining the top 10 skewed targeting audiences respectively. For each case, also shown within brackets is the percentage recall achieved of the target sensitive population on the platform.

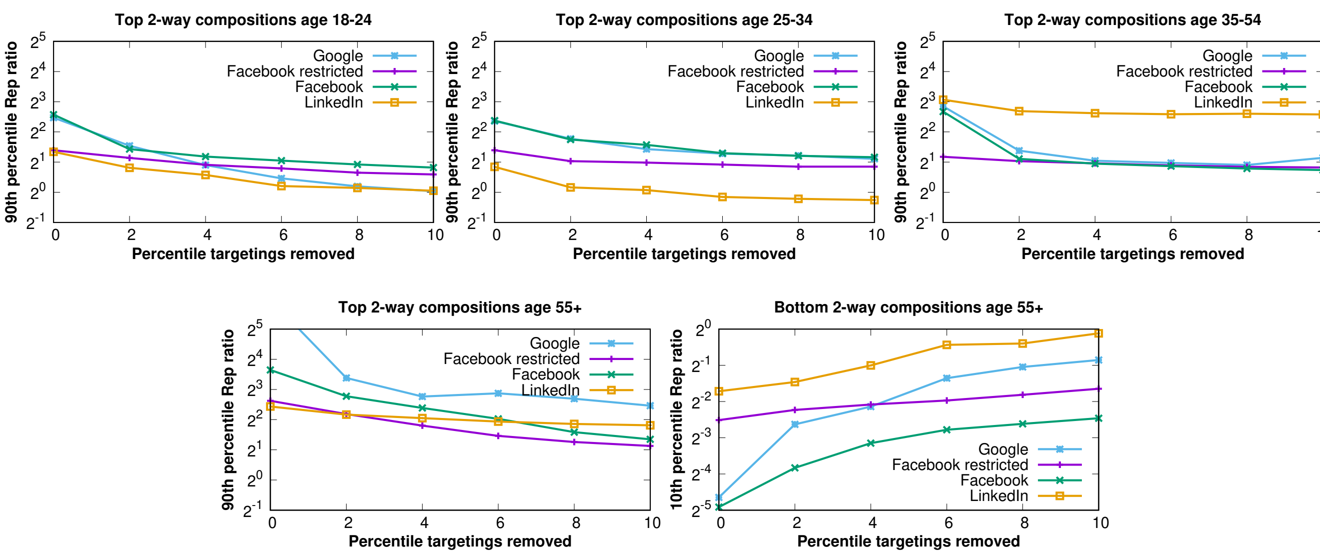

Figure 6: Effect of removal of the most skewed individual targetings on the skew of pairwise targeting compositions, for different ages.

Ad platform	Favoured gender	Targeting 1 (T1)	Targeting 2 (T2)	Rep. ratio for T1	T2	T1 and T2
FB-restricted	Male	Interests → Mechanical engineering	Interests → Automobile repair shop	4.68	4.40	18.10
		Interests → Buy to let	Interests → Sedan (automobile)	2.62	2.50	10.13
		Interests → Hatchback	Interests → Computer engineering	3.25	3.05	10.91
		Interests → Electrical engineering	Interests → Cars	3.71	2.18	12.43
	Female	Interests → Interior design magazine	Interests → Credit Sesame	2.38	2.16	6.42
		Interests → Epidemiology	Interests → Credit Sesame	2.53	2.16	6.78
		Interests → Veterinary medicine	Interests → Bungalow	2.71	2.42	5.95
		Interests → Multi-level marketing	Interests → Living room	5.00	3.03	10.48
		Interests → Product design	Interests → Grocery store	2.48	2.39	5.00
Facebook	Male	Games → Strategy games	Industries → Military (Global)	4.58	4.00	29.28
		Industries → Construction and Extraction	Games → Racing games	5.09	5.00	29.75
		Industries → Construction and Extraction	Games → Strategy games	5.09	4.58	27.91
		Games → Massively multiplayer online games	Soccer → Soccer fans (high content engagement)	2.45	2.23	12.36
		Industries → Construction and Extraction	Consumer electronics → Audio equipment	5.09	4.24	25.09
	Female	Beauty → Cosmetics	Amazon → Owns: Kindle Fire	2.59	2.51	9.17
		Facebook page admins → Health & Beauty page admins	Family and relationships → Parenting	3.38	3.25	10.80
		Beauty → Hair products	Facebook Payments users (higher than average spend)	2.75	2.29	8.66
		Shopping → Boutiques	Industries → Education and Libraries	2.92	2.43	9.17
		Clothing → Children's clothing	Industries → Community and Social Services	5.96	2.62	18.33
Google	Male	Gamers → Sports Game Fans	Martial Arts → Kickboxing	4.00	4.21	36.92
		Gamers → Shooter Game Fans	Autos & Vehicles → Custom & Performance Vehicles	4.06	5.42	24.92
		Performance & Luxury Vehicle Enthusiasts	Martial Arts → Japanese Martial Arts	4.15	5.61	23.08
		Performance & Luxury Vehicle Enthusiasts	Computer Components → Chips & Processors	4.15	5.18	21.23
		Gamers → Shooter Game Fans	Computer Hardware → Hardware Modding & Tuning	4.06	4.62	18.77
	Female	Makeup & Cosmetics → Eye Makeup	Mediterranean Cuisine → Greek Cuisine	6.16	5.27	43.33
		Holiday Items & Decorations → Christmas Items & Decor	Food → Grains & Pasta	4.84	4.55	32.50
		Infant & Toddler Feeding → Toddler Meals	Crafts → Art & Craft Supplies	4.90	6.19	34.67
		Makeup & Cosmetics → Eye Makeup	Latin American Cuisine → South American Cuisine	6.16	4.49	32.50
		Skin Care Products → Anti-Aging Skin Care Products	Crafts → Fiber & Textile Arts	4.88	5.79	28.17
LinkedIn	Male	Manufacturing → Industrial Automation	Robotics → Swarm Robotics	2.80	2.26	8.38
		Job Functions → Engineering	Transportation & Logistics → Maritime	3.74	3.11	10.78
		Desktop/Laptop Preference → Linux	Computer Software → Operating Systems	5.72	4.19	15.57
		Energy & Mining → Mining & Metals	Job Seniorities → CXO	2.94	2.55	7.27
		Manufacturing → Industrial Automation	Computer Hardware → CPUs	2.80	2.61	6.91
	Female	Health Care → Medical Practice	Job Functions → Accounting	2.41	2.17	7.48
		Corporate Services → Executive Office	Working Environments → Home-Based Business	1.90	1.87	5.15
		Consumer Goods → Cosmetics	Human Resources → Workplace Conflict Resolution	4.48	3.21	11.68
		Job Functions → Administrative	Health Care → Medical Practice	3.70	2.41	8.67
		Human Resources → Workplace Etiquette	Health Care → Medical Practice	2.73	2.41	6.22

Table 2: Illustrative examples of the "Top 2-way" skewed targeting compositions, showing how individual targeting options (skewed toward a particular gender) could sometimes be combined to get a targeting that is much more skewed than the individual targeting options.

Ad platform	Favoured ages	Targeting 1 (T1)	Targeting 2 (T2)	Rep. ratio for T1	T2	T1 and T2
FB-restricted	18-24	Interests → Vocational education	Interests → Electrical engineering	1.89	1.63	7.87
		Interests → Roommate	Interests → Moving company	1.53	1.27	4.63
		Interests → Microcredit	Interests → Mortgage calculator	1.32	1.27	3.70
		Interests → Entry-level job	Interests → Apartment Guide	1.84	1.78	4.47
		Interests → Cars	Interests → Vocational education	1.96	1.89	4.62
	55+	Interests → Income tax	Interests → Consumer Reports	2.46	2.38	7.06
		Interests → Reverse mortgage	Interests → Life insurance	7.95	3.73	21.32
		Interests → Part-time	Interests → Home equity line of credit	2.80	2.60	7.30
		Interests → Epidemiology	Interests → Government debt	2.08	2.06	5.41
		Interests → Data security	Interests → Fundraising	2.91	2.46	7.29
Facebook	18-24	Education Level → Some high school	Industries → Military (Global)	3.29	1.69	10.59
		Education Level → Some high school	Reading → Manga	3.29	2.39	8.95
		Education Level → In college	Sports → Volleyball	5.75	2.59	14.15
		Sports → Volleyball	Expats → Lived in China (Formerly Expats - China)	2.59	1.97	6.14
		Education Level → Some high school	Games → Massively multiplayer online games	3.29	2.43	7.83
	55+	Relationship Status → Widowed	Canvas Gaming → Played Canvas games (last 7 days)	8.13	7.47	29.22
		Facebook access (browser): Internet Explorer	Facebook access (OS): Windows 8	4.12	2.63	11.93
		Relationship Status → Widowed	Likely engagement with conservative political content	8.13	2.50	21.35
		Apple → Facebook access (mobile): iPhone 5	All Parents → Parents (All)	3.28	2.44	8.49
		Apple → Owns: iPhone 6 Plus	Primary email domain → AOL email users	2.96	2.49	7.58
Google	18-24	Highest education high school graduate	Business Services → Knowledge Management	1.56	1.43	11.64
		Employment → Internships	Online Communities → Virtual Worlds	1.62	1.67	6.02
		Employment → Sales & Marketing Jobs	Books & Literature → Fan Fiction	1.53	1.53	5.29
		Employment → Temporary & Seasonal Jobs	Table Games → Table Tennis	1.52	2.81	9.42
		Marital Status → In a Relationship	Software → Educational Software	1.64	1.76	5.59
	55+	Homeownership Status → Homeowners	Central Anatolia → Ankara	4.30	6.01	69.94
		Marital Status → Married	Austria → Vienna	5.00	4.93	24.03
		Retirement → Retiring Soon	Education → Alumni & Reunions	11.60	6.29	50.38
		Retirement → Retiring Soon	Movies → Classic Films	11.60	4.45	49.61
		Motor Vehicles by Brand → Lincoln	Games → Tile Games	3.83	4.70	18.93
LinkedIn	18-24	Linkedin News Editors' Top Startups (United States)	Job Functions → Operations	1.25	1.14	3.43
		Consumer Goods → Food & Beverages	Education → Higher Education	1.36	1.16	3.60
		Recreation & Travel → Recreational Facilities & Services	Member Traits → Job Seeker	1.19	1.13	2.84
		Public Administration → Political Organization	Mobile Preference → iPhone Users	1.21	1.00	2.83
		Desktop/Laptop Preference → Mac	Public Administration → Political Organization	1.23	1.21	2.81
	55+	Job Seniorities → CXO	Insurance → Life Insurance	3.71	3.13	8.02
		Job Functions → Consulting	Business Administration → Operations Management	3.01	2.90	6.26
		Job Seniorities → CXO	Corporate Finance → Corporate Financial Planning	3.71	3.42	7.63
		Agronomy and Agricultural Sciences	Job Functions → Consulting	3.02	3.01	5.83
		International Trade → Economic Sanctions	Job Functions → Consulting	3.06	3.01	5.81

Table 3: Illustrative examples of the "Top 2-way" skewed targeting compositions, showing how individual targeting options (skewed toward a particular age range) could sometimes be combined to get a targeting that is more skewed than the individual targeting options.