



Discrimination through Image Selection by Job Advertisers on Facebook

Varun Nagaraj Rao
Princeton University
Princeton, NJ, USA
varunrao@princeton.edu

Aleksandra Korolova
Princeton University
Princeton, NJ, USA
korolova@princeton.edu

ABSTRACT

Targeted advertising platforms are widely used by job advertisers to reach potential employees; thus issues of discrimination due to targeting that have surfaced have received widespread attention. Advertisers could misuse targeting tools to exclude people based on gender, race, location and other protected attributes from seeing their job ads. In response to legal actions, Facebook disabled the ability for explicit targeting based on many attributes for some ad categories, including employment. Although this is a step in the right direction, prior work has shown that discrimination can take place not just due to the explicit targeting tools of the platforms, but also due to the impact of the biased ad delivery algorithm. Thus, one must look at the potential for discrimination more broadly, and not merely through the lens of the explicit targeting tools.

In this work, we propose and investigate the prevalence of a new means for discrimination in job advertising, that combines both targeting and delivery – through the disproportionate representation or exclusion of people of certain demographics in job ad images. We use the Facebook Ad Library to demonstrate the prevalence of this practice through: (1) evidence of advertisers running many campaigns using ad images of people of only one perceived gender, (2) systematic analysis for gender representation in all current ad campaigns for truck drivers and nurses, (3) longitudinal analysis of ad campaign image use by gender and race for select advertisers. After establishing that the discrimination resulting from a selective choice of people in job ad images, combined with algorithmic amplification of skews by the ad delivery algorithm, is of immediate concern, we discuss approaches and challenges for addressing it.

CCS CONCEPTS

- Social and professional topics → Technology audits; Employment issues; Socio-technical systems; Race and ethnicity; Gender;
- Human-centered computing → Social media;
- Information systems → Online advertising.

ACM Reference Format:

Varun Nagaraj Rao and Aleksandra Korolova. 2023. Discrimination through Image Selection by Job Advertisers on Facebook. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*, June 12–15, 2023, Chicago, IL, USA. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3593013.3594115>



This work is licensed under a Creative Commons Attribution International 4.0 License.

FAccT '23, June 12–15, 2023, Chicago, IL, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0192-4/23/06.
<https://doi.org/10.1145/3593013.3594115>

1 INTRODUCTION

Targeted advertising platforms are widely used by job advertisers to reach potential employees. For example, a study commissioned by Facebook in 2018 found that one in four people in the U.S. searched for, or found a job using Facebook's platform [40]. At the same time, issues of discrimination due to targeting have surfaced and received widespread attention in recent years. Advertisers could misuse the targeting tools, including algorithmic audience creation mechanisms, to exclude people based on gender, race, location and other legally protected attributes from seeing their job ads [5, 6, 32, 70, 82, 83]. In response to lawsuits by the American Civil Liberties Union to the U.S. Equal Employment Opportunity Commission [1], and Facebook's commissioned Civil Rights Audit [25], Facebook in 2019 disabled the explicit targeting based on gender, age and zip code for some ad categories, including employment [69].

Although this is a step in the right direction, researchers [3, 42] and the U.S. Department of Justice [81] demonstrate that discrimination can take place not just due to the explicit targeting tools made available by the platforms, but also due to the impact of a bias in the ad delivery algorithm. Thus, one must look at the potential for discrimination broadly, beyond the lens of the explicit targeting.

An important and to-date, under-examined question is the potential for discrimination through the disproportionate representation of certain people in job ad images. Selective image use can lead to discrimination as follows: First, images have a unique persuasive power and social science literature [8, 80, 88, 89] demonstrates images present on recruitment websites affect the application intention of job seekers and can be used to manipulate the gender and racial composition of those who apply. Second, complementary evidence from case law [45] shows that portraying only the dominant demographic in images discourages minorities from seeking out those opportunities. And third, a skew in image selection can be algorithmically amplified by ad delivery algorithms, solely based on the demographic characteristics of the people depicted [46].

To investigate the prevalence of such discrimination, that combines both targeting and delivery in online advertising, we analyze job advertisers' selective use of people in job ad images on Facebook. We hypothesize that a job advertiser attempting to exclude or discourage individuals of a certain gender from applying to their jobs (or equivalently, to make the job more appealing to applicants of one gender, and less appealing to applicants of another gender), thereby circumventing the targeting restrictions of the platform, can do so through the selective use of people in the images chosen for their job ad campaigns. For example, a trucking company interested in hiring only men as drivers, could create job ads using images containing only men in them. Thus, we point out that in modern ad systems, a job advertiser's selective use of images

depicting people of only one demographic is almost analogous to explicitly targeting that demographic.

We collect and analyze data¹ from the Facebook Ad Library to demonstrate that job advertisers are already leveraging the selective use of people in job ad images for potential discrimination through:

- Evidence of advertisers running many campaigns using ad images of people of only one perceived gender.
- Systematic analysis for perceived gender representation in all current ad campaigns for truck drivers and nurses.
- Analysis of perceived gender and race representation for select advertisers.

After establishing that the discrimination, resulting from the selective choice of people in job ad images, combined with algorithmic amplification of skews thanks to ad delivery optimization, is of immediate concern, we discuss approaches for addressing it. Specifically, we describe the data and functionality which should be (but currently is not) provided by the ad platforms to enable public-interest researchers to detect potentially discriminatory image selection targeting. Furthermore, we underscore the necessity of transparency not only with regards to the advertiser choices, but also with regards to the ad delivery algorithm as applied to employment advertising. We then draw parallels to discrimination through selective use of language, suggesting that selective use of people in images should have similar legal protections and platform guidance. Finally, drawing on literature from jury representation and special education, we discuss approaches for measuring advertiser intent in image selection, and the challenges to their adoption.

In summary, our main contributions are:

- (1) define a new means of discriminative targeting by advertisers through selective use of people in the ad images, and demonstrate its prevalence for job ads on Facebook (Section 3).
- (2) provide desiderata for platform's actions to address our findings and discuss the challenges of establishing normative metrics on the basis of advertiser choices alone (Section 4).

2 BACKGROUND AND RELATED WORK

Before we introduce our methods and analyses, we provide background and related work on discrimination in targeted advertising systems and the corresponding legal, policy and platform responses to previously discovered issues. Subsequently, we highlight the unique role played by images in employment advertising and motivate the need to study job advertiser image selection.

2.1 Ad Targeting, Delivery and Discrimination

2.1.1 The ad system: targeting and delivery. Facebook's ad system consists of two phases - ad creation and ad delivery [26, 29]. During *ad creation* the advertiser chooses (1) their business objective, i.e. an outcome they'd like to achieve, such as increasing the number of visitors to their website, (2) contents of the ad including the images, text, and destination page link, (3) the targeting parameters, specifying the kinds of users they'd like the ad to be shown to, and (4) the budget and duration of their ad campaign. Advanced targeting

capabilities, using demographic, geographic, interest and/or behavioral characteristics of users have long been touted by Facebook as a useful tool for advertisers.

Ad delivery is the process through which a subset of the users targeted by the advertiser are chosen to see the ad. For each user, the platform runs an auction among all advertisers targeting that user, to determine which ad to show. The high-level information publicly known about the auction is that it does not necessarily select the ad with the highest bid; rather, for each candidate ad, Facebook computes its *Total Value*, defined as:

$Total Value = Advertiser Bid \times Estimated Action Rate + Ad Quality$; and then shows the user the ad with the highest total value [26, 29, 58]. All components of this equation are computed using machine learning. The *Estimated Action Rate* is Facebook's predicted likelihood of the user taking an advertiser's desired action, the *Ad Quality* is a machine learning model that combines feedback from users and predictions based on image and text of the ad creative, and even the *Advertiser Bid* is an algorithmic model, where Facebook chooses what to bid on behalf of the advertiser based on their budget and expressed preference in bid strategy [27].

2.1.2 Discrimination in ad targeting. Prior work has found significant evidence of Facebook's extensive targeting tools being used in order to exclude individuals on the basis of gender, age, or race in housing and employment advertising (which are areas governed by U.S. anti-discrimination laws). For example, Facebook tools offered the ability to exclude Black and Hispanic users living in a specific area from the targeting of housing ads [5, 82]. Furthermore, [83] found ads from ten traditionally male-dominant industries targeted just at men, including a software company, a moving company and a police department. Moreover, [6] found that advertisers excluded older workers from seeing job ads. For instance, financial analyst job ads were only targeted at users aged 25–36 years.

2.1.3 Discrimination in ad delivery. Discrimination can occur not only due to advertisers' choices during targeting, but also as a result of the ad delivery process, as the machine learning driven components of the total value equation can be biased.

Specifically, Ali et al. [3] showed that in the case of Facebook, the ad creative image influences and skews the ad delivery along gender and racial lines, in ways that cannot be explained by market effects or users' interactions with the ads. Follow-up work by Imana et al. [42] demonstrated that the skew in ad delivery by gender in the case of employment ads cannot be justified by differences in qualifications among demographics in the target audience, and is thus fully attributable to algorithmic decisions made by Facebook in its own business interests. More recently, concurrent and complementary work by Kaplan et al. [46] showed that ad delivery can be dramatically skewed solely as a result of demographic characteristics of the people depicted in the images. These works provide evidence of discrimination in ad delivery, a hypothesis that was put forward by Sweeney [78] and further investigated by [14, 17, 18, 20, 50, 51].

2.2 Legal, Policy and Platform Response

2.2.1 Removal of exclusionary targeting features. In response to the investigative reporting [5, 6, 83], lawsuits [1, 82], and a Civil Rights Audit [25], Facebook disabled the advertiser ability to explicitly

¹All data used in our analysis can be accessed at https://github.com/varunnrao/job_ad_images_facct23

target employment ads by age or gender in 2019 [69]. The ability to explicitly target by race and ethnicity had been removed earlier in 2018 [23]; however, proxies for targeting by racial categories remained available even in 2022 [28, 48]. Facebook also introduced a requirement for advertisers to self-identify when running ads on housing, employment, or credit (HEC) issues, and to acknowledge adherence to Facebook's Discriminatory Practices policy [24, 30]. Google [37] followed suit by restricting targeting criteria for HEC ads and requiring advertisers to acknowledge adherence to a personalized advertising policy. There is no advertiser-facing feature on Google to self-identify the ad category; instead, ads may be automatically labeled as belonging to an HEC category after creation and review. LinkedIn retained the ability for advertisers to target based on age or gender, but required advertisers to self-certify that for their HEC or education ads, they will not use the platform to discriminate based on age or gender [53].

2.2.2 Progress on reducing discrimination in ad delivery. Addressing discrimination in ad delivery remains more complex. In a June 2022 settlement with the Department of Justice, Facebook committed to “develop a new system to address racial and other disparities caused by its use of personalization algorithms in its ad delivery system”, called the Variance Reduction System (VRS) [81]. Specifically, the VRS system, aims to ensure that the age, gender and estimated race distribution of the audience that is shown a housing ad closely resembles the distribution of age, gender and estimated race of the audience targeted by the advertiser [58]. During an ad's delivery, Facebook periodically compares the ratio of impressions delivered to a particular age / gender / racial subgroup, with that subgroup's fraction in the advertiser's target audience. When the ratios diverge, Facebook adjusts one of the machine learning algorithms used to calculate the ad's total value in the auction, in a way that will change the likelihood that this ad will win the auction and be shown to a user of a particular subgroup. Although as of January 2023, the VRS system is implemented only for housing ads, Facebook has committed to expanding it also to employment advertising [7].

A limitation of the VRS approach for eliminating discrimination is that the demographic composition of the advertiser's targeted audience is chosen as the goal for the composition of the delivery audience. Such a choice may be problematic as long as the advertisers continue to have access to tools (e.g. custom audiences, location based audience creation) to select their targeted audience in a discriminatory manner [32, 70, 75].

As will become clear from subsequent discussion in Sections 2.3.1 and 2.3.2, the VRS approach is unlikely to address the new type of discrimination through image selection that we identify in this work. Extrapolating from the social science literature and case law [8, 45, 80, 88, 89], the selective use of people in ad images affects the ad recipients' likelihood of acting on the ad (e.g., their likelihood of *clicking* on it), whereas the VRS is entirely focused on minimizing the variance in *impressions* (i.e. showings of the ad).

2.2.3 The Facebook Ad Library and its functionality. The *Facebook Ad Library*² is a transparency tool created by Facebook that provides a keyword-based interface to a searchable collection of active ads on the platform. After choosing an ad category among (1) all ads,

(2) housing, (3) employment, (4) credit, or (4) issues, elections or politics (referred to as political ads henceforth), one can search for ads based on specific keywords or suggested advertiser names.

The search results page loads ads in the category chosen in reverse chronological order as one scrolls down the page. For all active ads, the search results page contains details about the ad content used, including one or more of the ad text, destination page link, video, and image(s), and other meta data such as the Ad ID, and the date it started running. However, it does not contain data about inactive ads. Critically, with the exception of political ads, neither the ad targeting information (such as advertiser targeting choices or campaign budget) nor the ad delivery information (such as the number of people an ad reached, location where it was delivered, or a demographic breakdown of the ad recipients) is available.

Most prior works using the Facebook Ad Library focused on political and issue advertising [22, 52, 72, 74]. We are aware of only one effort that uses the data in the ad library for inferences about employment advertising — a class action charge by Real Women in Trucking against Meta (formerly Facebook) alleging that Meta discriminates against women and older people when deciding which users receive employers' job ads on Facebook [68]. In it, the supporting evidence lists over 75 employment ads from the Ad Library that were delivered primarily to men. For example, an employer seeking truck drivers in North Carolina reached an audience that was only 5% female.

2.3 Role of Images for Job Seekers

Our work is focused on the potential ability and practice of discrimination through image selection. We now describe why images play a unique role in employment advertising from the social science and policy perspectives.

2.3.1 The social science perspective. Images play a unique role in advertising due to their distinctive characteristics, which are not present in other sources of media [56, 73]. In his 1997 book “Visual Persuasion: The Role of Images in Advertising” [56], Paul Mesaris outlines iconicity, indexicality, and syntactic indeterminacy as the fundamental traits of images that distinguish them from other modalities and uniquely affect the ad campaign. First, the iconicity of images gives advertisers access to strong emotional responses invoked in the target audience when they view an image. Second, images are indexical and can serve as documentary evidence of an advertiser's point of view. Third, images lack explicitness and syntax, and thus the audience's interpretation of the message conveyed through an image in an ad campaign is a creation of their own.

Based on social science theories, researchers argue that images present on recruitment websites and in advertising affect the application intention of job seekers and can be used to manipulate the gender and racial composition of those who apply. Recent works [88, 89] applied Spence's Signaling Theory [76], Social Identity Theory [79] and Visual Perception Theory [36] in support of this hypothesis. Based on the Signaling and Perception Theories, [89] argued that people who apply for jobs will resemble in their demographic characteristics those of the people depicted on the job websites. In an apparent attempt to leverage the effects of these theories in practice, organizations have been found to use images of diverse people in job ads to recruit diverse talent [8, 80]. In parallel,

²<https://www.facebook.com/ads/library/>

[88] leveraged the Social Identity Theory to deduce that “individuals naturally categorize themselves and others in terms of important visible characteristics (e.g., race, gender, age)”, and “develop more favorable attitudes toward in-group members and seek environments that affirm their identity”. As a result, they demonstrate that job seekers infer diversity cues from the people present in the images, which in turn influences their job application decision.

2.3.2 Policy context. Legal scholarship under both Section 704(b) of Title VII [86] and the parallel, but broader Fair Housing Act (FHA) [87] of U.S. law suggests that use of images containing only people of the dominant demographic in jobs or housing ads may be interpreted by the casual observer to convey a preference based on a protected class, irrespective of the advertiser’s intent [13, 17]. Evidence of such discrimination is detailed in the FHA case law [45]. According to [17], “the FHA case law connects the ordinary reader standard to the prohibition on sex-designated advertising columns (e.g., those with “Male Help Wanted” and “Female Help Wanted” headings) by explaining that advertisements that exclusively feature White models may discourage Black people from pursuing housing opportunities by conveying a racial message in much the same way that the sex-designated columns furthered illegal employment discrimination” (see Section 4.1.2 for additional discussion).

FHA case law can reasonably be extended to apply to the context of employment ads as evidenced by a high profile settlement involving Abercrombie & Finch in 2004 [41]. As part of that settlement, Abercrombie agreed to add more diversity to their marketing materials and reflect the makeup of the nation’s population; so as not to discourage members of minorities from applying for jobs that were otherwise dominated by “well-known collegiate, all-American - and largely White” people [39].

2.3.3 Use of images to drive a message of bias in non-advertising contexts. The influence of images in shaping biased perceptions or exaggerating diversity has been extensively researched across various contexts beyond online advertising. Moriearty [60] argued that the over-representation of Black youth as criminals in images used by media has resulted in disproportionately stricter policing and harsher sentencing for Black youth. Similarly, in an effort to attract under-represented members and create a misguided sense of belonging, universities have been found to over-represent African American students in their admissions brochures [67]. Netflix too has been accused of using misleading visual representations to entice viewers based on race; for example, some marketing posters contained only Black people even if they appeared only briefly in a show [91]. In fact, a 12 year longitudinal study of TV shows by Baruah et al. [9] found that male characters with a light skin tone occupy the majority of the screen time. Biases have also been identified in reports on vaccination and antimicrobial resistance; with White people shown in a professional capacity whereas children of color shown as vulnerable and exposed [15]. In the Computer Science community, the works of [47, 59, 66] found a gender- and race- driven representation bias in both image- and text-based Google Image Search results. Similarly, Luccioni et al. [55] found that Text-to-Image models depict care professions (e.g., dental assistant) nearly exclusively through as women, and positions of authority (e.g., director, CEO) – exclusively through men.

3 JOB ADVERTISERS’ SELECTIVE USE OF PEOPLE IN JOB AD IMAGES

A search of the Ad Library easily uncovers anecdotal examples of job advertisers across many occupations using people of only one gender in the images chosen for their ads. See Figure 1 for an illustration. Furthermore, we found examples of advertisers who invest into creating hundreds of distinct images for their campaigns, yet an overwhelming majority of those images depict people of only one gender (see Figure 6). Motivated by these examples, we perform a comprehensive study of gendered image selection by job advertisers seeking to employ truck drivers and nurses.

3.1 Perceived Gender Representation Among Truck Driver and Nurse Advertisers

3.1.1 Data Collection: To quantify how prevalent is the selective use of gendered people in job ad images, in January 2023 we scraped data from the Facebook Ad Library and saved the job ad images for two occupations - truck driver and nurses. We restricted our search to the employment ads category and to only those ads which contain images and memes. We chose these occupations because (1) they currently have a real world defacto skew towards a specific gender (91.9% of truck drivers in the US are men, 87.9% of nurses in the US are women [12]) and (2) we found thousands of ads for these occupations in the Ad Library, allowing for a large scale analysis.

Since the Ad Library does not load all search results when there are more than a few hundred, obtaining data on all advertisers running employment ads for truck drivers or nurses is not merely a task of performing one search. Instead, to stay within the number of results that load, we turned each keyword search into 26 separate searches by expanding a given keyword with an additional letter of the alphabet, i.e. we searched and downloaded results for the keywords truck driver a through truck driver z and nurses a through nurses z³. For each job advertiser whose ads appeared in the search results, we recorded their Page ID, and then performed a separate search⁴ to download all the images, excluding the videos, of all the advertiser’s active ad campaigns at the time. We explored different variations of the keywords such as singular or plural forms and whether or not to include quotes⁵, and ultimately settled on the form that produced the highest number of search results⁶. Although our data collection may have excluded some advertisers (e.g., those whose ad text contains only the words nurses or nurses followed by digits or special characters), we believe our sample is large enough to draw conclusions.

3.1.2 Data Annotation: Since an advertiser may re-use the same image in many campaigns, for each advertiser, we considered both the total number of images used and the number of distinct images used. We determined whether two ad campaign images are distinct using a hashing algorithm of the open source imagededup tool [44].

³A sample URL for the truck_driver_a keyword: https://www.facebook.com/ads/library/?active_status=all&ad_type=employment_ads&country=US&q=truck%20driver%20a&search_type=keyword_unordered&media_type=image_and_meme

⁴A sample URL for Truck Driver Recruiting America: https://www.facebook.com/ads/library/?active_status=all&ad_type=employment_ads&country=US&view_all_page_id=106206904474383&search_type=page&media_type=image_and_meme

⁵A search for the keyword enclosed in double quotes returns results with an exact match in the content of the ad.

⁶See Appendix A.1 for precise queries and select advertiser URLs.



Figure 1: True Friends Moving Company (left): 7 distinct images containing people, all exclusively depict men as movers. **Invo Healthcare (right):** 9 distinct images containing people, all exclusively depict women as healthcare staff.

Table 1: Data Summary of Image, Distinct Image, and People-Image Use by Truck Driver and Nurses Advertisers

Statistic	Truck Driver Advertisers	Nurses Advertisers
# of images	4,196	2,148
# of distinct images	2,741	1,581
# of advertisers	446	360
# of excluded advertisers	47	58
# of advertisers whose campaigns contain at least one image with people	159	259

We annotated each image for whether it contains people and if so, the perceived gender of the people in it. We did not consider people in the images who were not clearly representatives of the occupation being recruited for (e.g., children who the nurse was attending to, or family of a truck driver waiting for them at home). Further, we corroborated our results by recruiting crowdworkers through Amazon Mechanical Turk (MTurk) who were paid above the federal minimum wage. As a result, each distinct image was annotated by 2 distinct annotators. We excluded from analyses the advertisers for whom we could not reach an agreement on annotation, or were unable to decisively determine the gender of all people due to occlusion, blurriness or use of avatars. We summarize our annotated data in Table 1.

3.1.3 Findings: We find that discrimination through image selection is a prevalent practice among the advertisers we study. Specifically, among 159 truck driver advertisers whose campaigns contain people, 90 or 57% depict only men as truck drivers. Among 259 nurses advertisers whose campaigns contain people, 108 or 42% depict only women as nurses. Although we observe that the fraction of advertisers using images depicting only people of the stereotypical gender for their profession typically decreases as the number of distinct images used goes up; we find examples of extensive stereotypical image selection across all ranges of the number of images per advertiser. Specifically, among the 33 truck driver advertisers who use only 1 image, 25 or 76% depict a man; whereas among the 28 advertisers who use 6-9 distinct images, 13 or 46% depict only men. For the 29 advertisers in our dataset each using at least 15 distinct images, only 4 (but still four and not zero), or 14% depict only men. See Table 2 for the detailed statistics broken down by

the number of distinct images used and Figure 3 for examples of stereotypical image selection by two advertisers.

Detailed Analysis of Advertisers Who Put Deliberate Effort in Their Image Selection. We now restrict our analysis to the advertisers who use 15 or more distinct images across their campaigns, with at least 5 of those containing people, thus focusing on 22 truck driver and 18 nurses advertisers. We use the number of distinct images as a proxy for deliberate effort and thought the advertiser put into the visual messaging of their ad campaign, including thought towards the questions of diversity.

For each such advertiser, we present the raw statistics on the images used, and compute the fraction of images that depict men (for truck driver advertisers) or women (for nurses advertisers) among all distinct images depicting people. See Tables 6 and 7 in Appendix A.2. Analyzing the image use by these presumably more deliberate advertisers, we find only two who exclusively use images of one gender. However, the images are still dominated by the stereotypical gender choices – 415/503 (83%) of images depicting people used by truck driver advertisers use men and 270/380 (71%) of images depicting people used by nurses advertisers use women.

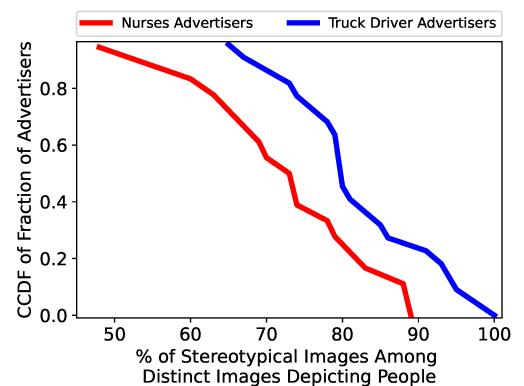


Figure 2: CCDF of the fraction of advertisers for a given percentage of stereotypical images use.

In Figure 2 we plot the fraction of advertisers whose percentage of stereotypical image use among distinct images depicting people

Table 2: A Breakdown of Advertiser Practices Based on the Number of Distinct Images Used

# of Distinct Images	% (#) of Truck Driver Advertisers			% (#) of Nurses Advertisers		
	Total	using ≥ 1 image of a person	whose images contain only men	Total	using ≥ 1 image of a person	whose images contain only women
1	187	18% (33)	76% (25)	92	70% (64)	70% (45)
2-5	133	44% (59)	71% (42)	118	91% (107)	45% (48)
6-9	35	80% (28)	46% (13)	53	92% (49)	16% (8)
10-14	13	77% (10)	60% (6)	20	100% (20)	30% (6)
≥ 15	31	94% (29)	14% (4)	19	100% (19)	5% (1)
Total	399	40% (159)	57% (90)	302	86% (259)	42% (108)

exceeds a certain value; i.e. the complementary cumulative distribution function (CCDF) of advertisers as a function of the percentage of stereotypical image use. As can be seen from it, the stereotypical image use is prevalent among these advertisers. For example, for 15/22 (68%) truck driver advertisers and 6/18 (33%) nurse advertisers, the percentage of stereotypical image use exceeds 78%. Overall, the average rate of stereotypical gender use across all advertisers is 0.73 for truck drivers and 0.82 for nurses.

As an aside, we observe that truck driver and nurses advertisers differ in their frequency of depicting people in their images. Truck driver advertisers feature more distinct images, but many of them depict trucks without people (the average fraction of images containing people across all truck driver advertisers is 0.35); whereas nurses advertisers may use fewer distinct images, but a larger fraction of them contain people (average fraction – 0.77).

In summary, we find that even among the advertisers who have invested effort into the selection of images for their ads, the selective image use resembling discriminatory targeting is prevalent.

Diverse Outreach. We now search for examples of advertisers who may be deliberately trying to diversify their workforce by exclusively depicting women in truck driver and men in nurse ads. We find limited evidence of such efforts – 5/159 truck driver and 4/259 nurses advertisers. Furthermore, most of these advertisers use a single distinct image, which limits our ability to judge whether the non-stereotypical image selection was deliberate. We highlight DHS Logistics Solution in Figure 7 of Appendix A.2 as an exception.

3.2 Perceived Gender and Race Representation Among Select Advertisers

With the goal of studying whether advertisers may be selecting images to discriminate by race in addition to gender, we now focus our analysis on a select set of advertisers and study the depiction of people in their images according to both gender and race.

Data Collection: We chose a diverse set (in terms of worker hours, domains, skill sets, demographic distribution, and income levels) of occupations and associated advertisers to study based on Bureau of Labor Statistics (BLS) data from 2021 [12], contingent worker supplement from 2018 [11] and availability of job ads in the Ad Library containing images of people, leading to the following:

Occupation Specific Advertisers (N=15): BestBuy, Doordash, Eataly, Geico Careers, Drive with HopSkipDrive, Instacart, Drive with Lyft, Nationwide Job Search for Education, Nationwide Job Search for Information Technology, Nurse Recruiter,

NYPD Recruit, Safeway, TSA, Uber, UPS Jobs.

Job Aggregator Advertisers (N=3): Monster, SimplyJobs, Talent. We map occupation specific advertisers to the closest BLS categories based on the type of ads; e.g., NYPD Recruit is mapped to the BLS category “Police officers”.

We scraped ads of the chosen advertisers from the Ad Library in October - November 2021 (and repeated the scrape in October 2022). See Table 4 and Table 5 in Appendix A.2 for a detailed summary of per-advertiser statistics in 2021.

Data Annotation: We recruited U.S. based Amazon Mechanical Turk (MTurk) workers, and compensated them above the federal minimum wage to annotate all images according to perceived gender and race. For each image, 2 workers were asked to count, if they exist, the number of people in the image according to gender (man, woman) and race (White, Black, Asian, Hispanic). We included form logic to ensure annotations were meaningful before they were submitted. We included an “Other” option for the annotators if it wasn’t possible to determine gender or race or both (~9% of total annotations). For this annotation task, we asked annotators to count all people in the images, and not just employees of the occupation being advertised. The obtained the final count of people in the image by averaging the annotations of the two annotators. As in the previous analysis, our method makes the simplifying assumption of binary gender for the sake of comparison with the binary BLS data. Furthermore, it does not annotate the images by skin tone [10] since such annotation would make a comparison with BLS data impossible. We restrict our analyses to the two most prevalent races in the data: Black and White.

Analysis: For a specific occupation, we judge whether an advertiser may be trying to skew their ads, to discourage or attract people of a specific demographic, by comparing the rate of representation of gender (woman) and race (Black, White) of people depicted in the images the advertiser uses with that of the U.S. workforce (given by BLS) for that occupation. If the rates of representation are different, we say that there is a *deviation in representation* and indicate whether it is an over- or under-representation.

For a given advertiser and across all their images, we compute the rate of representation of people as follows:

$$\% \text{ Women} = \frac{\# \text{ of Women}}{\# \text{ of Women} + \# \text{ of Men} + \# \text{ of Other}} \times 100$$

$$\% \text{ White (Black)} = \frac{\# \text{ of White (Black)}}{\# \text{ of White} + \# \text{ of Black} + \# \text{ of Asian} + \# \text{ of Hispanic} + \# \text{ of Other}} \times 100$$



Figure 3: Example advertisers with at least 10 distinct images depicting only people of stereotypical gender. Total Nurses Network (top row) depicts women in all 10 images. HireMaster (bottom row) depicts men in all 7 images.

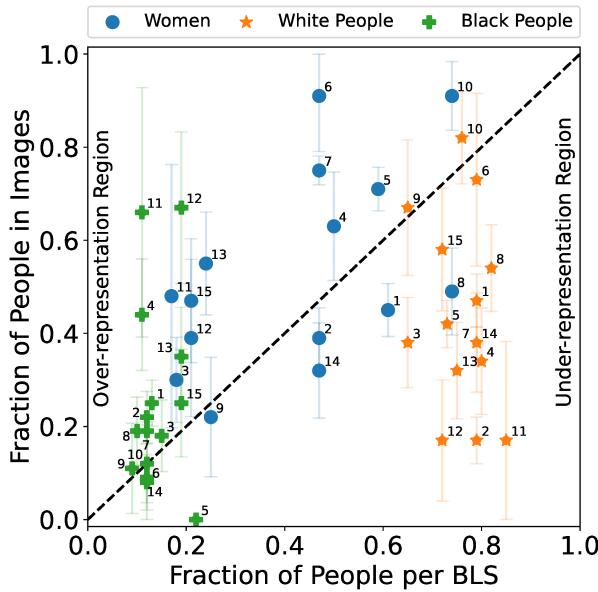


Figure 4: Rate of representation of gender and race in job ad images of select advertisers ($N=15$) compared to the rates of their corresponding BLS occupations. Column 1 of Table 4 in Appendix A.2 provides the mappings between numbers and advertiser names.

3.2.1 Findings. Figure 4 compares the rate of representation by gender (women) and race (White and Black) of people used in job ad images to the rate for that demographic characteristic per corresponding BLS occupation for each advertiser. Specifically, for a particular demographic characteristic and advertiser, the X-axis represents the fraction of people with that characteristic for that advertiser's occupation per BLS, and the Y-axis represents the fraction of people in the job ad images. We plot a 45-degree reference line to visually assess the magnitude of over- or under-representation given by the regions above and below the reference line, respectively. Full advertiser specific statistics are presented in Table 4 in Appendix A.2.

Across the 15 advertisers in Figure 4 and columns 6-8 of Table 4, we observe that most advertisers over-represent women and Black people, and under-represent White people. We draw definitive

conclusions for advertisers for whom the error bars do not cross the reference line, finding the following:

Women: 10 (over), 4 (under), 1 (inconclusive);

White People: 0 (over), 12 (under), 3 (inconclusive);

Black People: 8 (over), 1 (under), 6 (inconclusive).

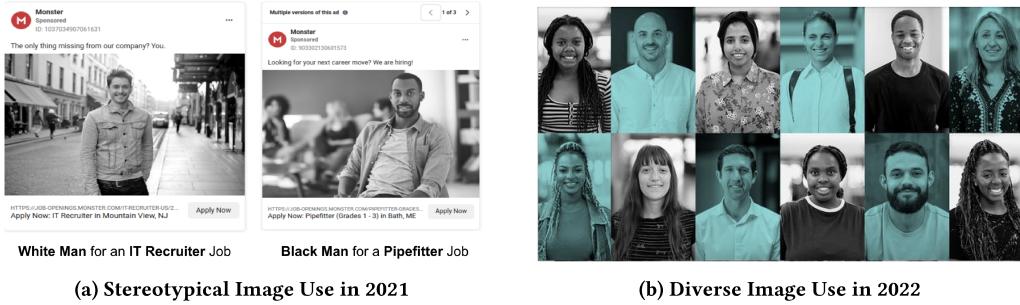
Our findings indicate that advertiser selection of people in job ad images with regards to gender and race is complex and varies across different occupations. Broadly, across a diverse set of occupations we find that some advertisers (e.g., NYPD Recruit and TSA) may be attempting to diversify their existing workforce by including more women and Black people than are employed in those professions per BLS in their images.

Evidence for Proactive Advertiser Selection by Monster: Advertisers may intentionally vary image selection based on their specific needs over time. To test this hypothesis, we studied [Monster.com](https://www.monster.com) (a job aggregator), which had the most number of campaigns, and compared image use in 2021 and 2022. We find evidence of deliberate image use and variation over time across different occupations.

In 2021, we find evidence of stereotypical image use across occupations. To enable this analysis, we extracted the occupation from the destination page URL of the ads, and used it to group our findings according to the top 3 roles - security officer, engineer, and technician. Our results are presented in Table 3. We find that the representation of women and Black people was mostly balanced for the engineer (45% women, 48% Black) and security officer (49% women and 50% Black) roles, but was skewed for technicians (28% women and 21% Black). We observed a White majority only in the case of technicians (64%).

Next, in order to delve deeper into possible advertiser intent behind image selection, we performed a reverse Google Image Search of the images. We observed that the images were generic and not present on the corresponding employer pages, implying a possible deliberate attempt by [Monster.com](https://www.monster.com) to discriminate through image selection. Examples of such stereotypical images are in Figure 5a. An expanded list of such occupations with stereotypical image use include: Black Man - pipe fitter, White Man - IT recruiter, Black Woman - security officer, White Woman - Sr. Software Engineer, Lactation Consultant.

In 2022, in contrast to the stereotypical image selection in 2021, most ads contained a single image with multiple people of diverse demographic characteristics as in Figure 5b. As a result, the percentage of women, White and Black people remains consistent across occupations as is evident in Columns 3, 5, 7 of Table 3.

Figure 5: Image selection by [Monster.com](#) across 2021 and 2022Table 3: Percentage of Women, White and Black People with 95% confidence interval for top three occupations of [Monster.com](#)

Occupation	% Women in 2021	% Women in 2022	% White in 2021	% White in 2022	% Black in 2021	% Black in 2022
engineer	45 ± 7	66 ± 7	28 ± 6	44 ± 7	48 ± 7	42 ± 7
security officer	49 ± 5	66 ± 3	33 ± 5	45 ± 3	50 ± 5	42 ± 3
technician	28 ± 9	67 ± 13	64 ± 10	44 ± 13	21 ± 8	43 ± 13

3.3 Limitations and Ethical Considerations

The limitations of our study are a result of the data and features of the Ad Library and our analyses' choices. The most significant limitation due to Ad Library's capabilities is the inability to quantify the real-world impact of the selective image use, as it does not provide information about a job ad's budget, optimization objective, or the number of people reached or engaged by the ad. We discuss desiderata for greater platform transparency in the context of discrimination through image selection by job advertisers in Section 4.1.1. The limitations due to the choices we made are:

Selection Bias: Our analyses are limited to the advertisers we chose to study and the timing of data collection; choices that may be particularly impactful for our study of racial bias in Section 3.2.1. The selections are driven by lack of functionality that would allow for a comprehensive or representative sampling by occupation.

Annotation Bias: Since we specified filtering criteria for MTurk workers who annotated the scraped images, the cohort of annotators may have been biased. Further, although we included an "Other" category when workers were not able to accurately annotate an image, they could have misreported or failed to report their perceptions. Finally, since crowdworkers operate within a social network of other crowdworkers, the mistakes could have been exacerbated by mutual influence and collaboration [38].

Limited View of Gender and Race: We restrict analyses to binary gender to allow for comparison with publicly available U.S. workforce data. Further, we restrict our analysis to the two most prevalent races in the scraped data. We view this limited categorization as only a first step towards more nuanced analyses.

Base Rate Comparison: Several possible base rates of demographic distributions by profession could be relevant for comparing the rate of representation of people, including geography-driven and employer-driven specifics. The choice of base rate influences the analysis performed and the resulting conclusions drawn from it. We chose to compare with one such base rate, given by the BLS national averages, in order to maintain consistency with prior audits related

to visual representation [47, 59] and due to infeasibility of accounting for possible regional variations in demographic distributions for most jobs and advertiser types.

Ethics: All data used was publicly accessible through the Ad Library; however, since the Ad Library only shows active ads, we expect most of the data collected for this study to not be accessible after some time. The study was approved by our institution's IRB.

3.4 Algorithmic Amplification

Building on the discussion in Section 2.1.3 of the role of the ad delivery for potential discrimination, we now illustrate how even a seemingly small difference in the rate of representation of people in job ad images can lead to significant differences across demographic groups in the delivery of those ads.

Kaplan et al. [46] demonstrated that the implied racial identity of a person depicted in an image affects the demographics of the audience that is shown the ad. For example, in their experiments, the delivery audience of a job ad in the lumber industry depicting a Black man was 55% Black, 45% White; whereas the delivery audience for the same ad depicting a White man was 72% White and only 28% Black, even when both ads were targeting a balanced audience of Black and White people, and were otherwise equivalent (see Section 6 in [46]). Consider a hypothetical advertiser running 100 identical (in terms of targeting, budget, etc.) ads for a job in the lumber industry, with only a very slight racial skew in the images used, e.g., suppose 51 of the ads use images depicting White men and 49 – depicting Black men. Assuming the ad delivery algorithm performs as measured in [46] and racially balanced targeting, the delivery audience of this advertiser will be $51 \times 72\% + 49 \times 45\% = 58.77\%$ White and 41.23% Black. In other words, a 2% difference in the racial representation in image selection would lead to a 17.5% difference in representation at delivery due to algorithmic amplification. If the advertiser ran 3 ads, 2 of them depicting White men and 1 depicting

a Black man, then the delivery audience would be 63% White and 37% Black, a 26% difference.⁷

Taken together, our empirical findings of current advertiser practices in biased image selection combined with the results on the use of implied identity by the ad delivery algorithm from prior work and the unique role of identity in images for subsequent job applications argue that selective use of people in images for job ads should receive similar scrutiny to that of targeting and delivery.

4 ADDRESSING DISCRIMINATION THROUGH IMAGE SELECTION

Having demonstrated the prevalence of selective choice of people in employment advertising, and its implications for discrimination thanks to ad delivery optimization and the unique role of images to elicit action, we now discuss approaches and challenges for addressing it. We approach the question from multiple perspectives: (1) transparency desiderata to perform audits of selective image choice, (2) transparency desiderata for the ad delivery algorithm in the context of its use of implied identity, (3) implications from parallels with practices in other domains, (4) metrics on image selection that could be used (by Facebook or public interest groups) to ascertain a particular advertiser's discriminatory intent.

4.1 Desiderata for Platform's Actions

4.1.1 Desiderata for Enhanced Platform Transparency for Job Ads and Ad Delivery Algorithm. The questions posed by our study and the challenges encountered while performing it imply that, at a minimum, the Ad Library for employment ads should provide the same data as is already being provided for social issues, elections or politics ads. Specifically, it should enable obtaining a list of all employment advertisers (rather than require scraping workarounds we had to resort to in Appendix A.1). Then for each advertiser, it should provide information on the targeting criteria and budget chosen by the advertiser for each of their ads, as well as delivery information – the number of people shown the ad, broken down by their demographic characteristics. To study changes in advertiser image selection practices over time, the Ad Library should include data for both active and past campaigns. Specific to considerations for discrimination in employment targeting, the Ad Library should provide an easy way to download the images used by the advertisers, label (or ask advertisers to self-identify) the job industry of the ad, and provide aggregate information about the people who engaged with the ad, broken down by their demographic characteristics. Furthermore, rather than aggregating the targeting data at an advertiser level or on a weekly basis as is done for political ads, employment ad data should be provided at an ad campaign level over the duration of that campaign. This desiderata echoes the one called for by [4, 21, 42], and is feasible given the existing implementation for political ads.

However, as is clear from the works related to discrimination due to ad delivery optimization, discussed in Sections 2.1.3 and 3.4, transparency of advertiser image selection practices is not sufficient

⁷If the ad delivery algorithm amplifies delivery unequally across races, then even if an advertiser depicts White and Black men in an equal number of ads, the delivery audience would be skewed, e.g. for lumber job ads – 58.5% White, 41.5% Black.

to understand the implications of such choices, neither for public-interest researchers nor for the advertisers themselves. We thus call for either disabling of the use of ad delivery optimization (and the opaque machine learning algorithms underlying the total value equation) for employment advertising or a significant increase in transparency of these algorithms when applied to job ads. A viable path for this kind of transparency has recently been proposed by Imana et al. [43]. Specifically, their proposal is a platform-supported special access API that gives public-interest researchers an ability to query the values of the estimators used in the Total Value equation, and study how they vary depending on the perceived gender and race of the person in the image and the platform-inferred demographic of the potential ad recipient. The approach of [43] addresses the concerns raised by platforms in response to requests for greater accountability – those of user privacy and protections of the platforms' proprietary code and business interests – through the use of differential privacy and the specifics of the set-up of the auditor - platform API interaction.

4.1.2 Platform's Guidance to Advertisers. The platforms should take proactive action towards ensuring that advertisers do not deliberately discriminate through their image selection. In fact, they already do so in a somewhat analogous domain of potentially discriminatory text selection by advertisers.

Similar to the case of selective image choice influencing who applies for jobs, prior social science work shows that selective use of language affects ad recipients' actions [35, 77] (e.g. removal of the word "aggressive" from the ad text increases the number of women applicants). Legal precedent establishes that the practice of printing classified ads in two separate columns: "Help wanted: Male" and "Help wanted: Female" violates anti-discrimination law [2, 49, 62, 63]. Thanks to the precedent and a relative ease for identifying discriminatory language use compared to discriminatory image use, advertising platforms such as Facebook [31] and LinkedIn [54] already have advertiser guidance about the language for creating inclusive ads, and may be enforcing them at ad review time. We suggest platforms extend the guidance to image use, e.g., by updating their policies and advertiser training materials.

Furthermore, as platforms themselves gear up to deploy Generative AI-based tools to assist advertisers in their ad campaign and creative creations [57, 64], they should be mindful not to lead the advertisers to employ image selection in a discriminatory manner.

4.2 Measuring Advertiser Intent in Image Selection

We now ask the normative question – under the assumption of full access to advertiser image selection and ad campaign budget allocation, what could be the metrics to establish whether the advertiser image selection is discriminatory or to encourage a change in their choices? This is a non-trivial question, even when all ads are allocated an equal budget. For example, since 91.9% of truck drivers in the US are men, should a U.S.-based advertiser be permitted to select 91.9% images containing men for their truck job ads? If that advertiser is using 20 distinct images, and so can only achieve representation fractions that are a multiple of 5%, should depicting only one woman be permitted? Or should the targeting always strive to achieve parity by gender and ensure that everyone

is represented? In other words, should the desiderata for visual representation mimic the employment statistics as they currently are or as we as a society may aspire them to be (e.g., according to the demographic distribution of the population or equal representation of all)? If the latter, how does one determine the aspirational ratios for jobs that are hyper-local in nature (e.g., plumbing jobs), when the demographic distribution of the population varies by geographical location? Should the advertiser or the advertising platform bear responsibility for determining the reference ratios, especially when data on qualified people of a certain demographic in a particular region for a particular job may not be easily available? Finally, to what extent should an advertiser be allowed to make trade-offs between costs of building a large set of representative visual assets for their campaigns with the potential, hard-to-quantify harms of using less representative assets?

4.2.1 Inspiration from Jury Selection: One approach to consider is through the adaptation of practices for ensuring fair and impartial jury to the advertising context. In the U.S. that means that a pool from which juries are selected should satisfy a fair cross section requirement [34, 71] via a Duren Test [19]. It measures the extent to which the jury pool demographics differ from those of the community, where the relevant community consists of individuals who are eligible for jury service. Several disparity metrics are used to compute the differences in demographics, and case law establishes permissible levels of disparity (see [61] for an overview). The challenge of adopting it in the advertising context is that the equivalents of “eligible for the jury service” in the “relevant community” is difficult to establish – both who is qualified for the job and who constitutes the relevant community for a particular job (e.g., due to commute limitations) are open questions. In the jury selection case, this data is relatively easily available through the U.S. Census. Prior work in the computer science literature [59] adopts the approach of comparing the racial and gender composition of Google Image Search results to that of the composition in the U.S. workforce across specific occupations.

Literature on ensuring compliance with the Individuals with Disabilities Education Act, and determining whether disproportionalities on the basis of race exist in the placement or discipline of students with disabilities at the school district level may also be of inspiration [33, 84]. However, it also operates in the context of full information, basing calculations on the exact number of students of particular race in a particular school district that is known – an assumption that does not hold for the advertising context.

4.2.2 Diverse and Affirmative Outreach: Additional metrics may need to be established to assess the image use of advertisers who may be trying to affirmatively reach underserved populations or attempting to diversify beyond the demographics of the existing U.S. workforce for specific occupations. For example, the primary goal of Black Career Network (Page ID 70203958046) and Black Excellence KC (Page ID 563610990802939) is to recruit Black professionals. Therefore they could be justified in selecting many (or only) Black people for their job ad images. Complementarily, in our study we found examples of advertisers attempting to diversify their workforce beyond the existing U.S. workforce of specific occupations by over-representation of women or Black people, or under-representation of White people in their job ad images.

Accounting for such practices is difficult: Is the choice of advertisers to empower an under-represented group through their ad images ethical or is it legal? Can this tension be reconciled if the advertiser self-identifies their reasons for doing so? How long should such practices be permitted? Objections have been raised to such practices in the past, e.g., White males have brought lawsuits for “reverse discrimination” against employers who have taken affirmative action to ameliorate the impacts of past discrimination [85]. The approach for accounting for such practices remains an open question, represents an ongoing policy issue [7] and is presently the subject of active judicial scrutiny.

4.2.3 Judging the Outcome. Ultimately, since the advertiser’s image choice is only one of the inputs to an extremely complex and opaque system, one can argue that the desiderata should be placed on the eventual outcome of ad delivery, rather than on the targeting, and the burden of flagging potential discrimination be shared between the advertiser, the platform, and public-interest researchers. We believe the exploration of such approaches is an important area for future work.

5 CONCLUSION

Our work highlights new representation and transparency issues at the intersection of hiring and targeted advertising. We motivate the unique role of images in job ads on social media platforms, and then study the selective use of people in job ad images according to gender, across a nearly comprehensive set of truck driver and nurses advertisers, and a select set of advertisers according to both gender and race.

We find that a large percentage of truck driver and nurse advertisers predominantly select images containing people of stereotypical gender for that occupation, thereby, effectively, targeting or excluding people by gender. On the contrary, we find that across a wider set of occupations, some advertisers may be selecting images to promote diversity among their job applicants.

Through its empirical and contextual findings, our study argues that the approach for ensuring non-discrimination in targeted advertising systems should look at advertiser actions more broadly than merely through the lens of the explicit targeting criteria made available by the platform.

Through enumerating limitations of our study and their impact on the ability to judge the real-world impact of advertiser image selection, combined with discussion of implied identity use in ad delivery optimization to further amplify demographic skew, our work contributes to the growing chorus of researchers [3, 4, 21, 42, 90] and policy makers [16] advocating for increased transparency of both the advertiser practices and the platform’s ad delivery algorithms, particularly in the employment domain.

Our findings raise important questions for future work in terms of the metrics that could be used for judging advertiser intentions in image selection, strategies for non-discriminatory image selections, and their potential interactions with the ad delivery algorithm.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive feedback and Judith Swan for feedback to improve our writing. This work was supported in part by NSF awards #1916153, #1956435, #1943584.

REFERENCES

- [1] ACLU. 2018. ACLU and Workers Take On Facebook for Gender Discrimination in Job Ads. <https://www.aclu.org/press-releases/aclu-and-workers-take-facebook-gender-discrimination-job-ads>
- [2] ACLU. 2018. How Facebook Is Giving Sex Discrimination in Employment Ads a New Life. <https://www.aclu.org/news/womens-rights/how-facebook-giving-sex-discrimination-employment-ads-new>
- [3] Muhammad Ali, Piotr Sapiezynski, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through optimization: How Facebook's Ad delivery can lead to biased outcomes. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing*.
- [4] Muhammad Ali, Piotr Sapiezynski, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2021. Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*.
- [5] Julia Angwin and Terry Parris Jr. 2016. Facebook Lets Advertisers Exclude Users by Race. <https://www.propublica.org/article/facebook-lets-advertisers-exclude-users-by-race>
- [6] Julia Angwin, Noam Scheiber, and Ariana Tobin. 2017. Dozens of companies are using Facebook to exclude older workers from job ads. *ProPublica*. <https://www.propublica.org/article/facebook-ads-age-discrimination-targeting>
- [7] Roy L. Austin Jr. 2023. An Update on Our Ads Fairness Efforts. <https://about.fb.com/news/2023/01/an-update-on-our-ads-fairness-efforts/>
- [8] Derek R Avery and Patrick F McKay. 2006. Target practice: An organizational impression management approach to attracting minority and female job applicants. *Personnel Psychology* (2006).
- [9] Sabyasachee Baruah, Digbalay Bose, Meredith Conroy, Shrikanth S. Narayanan, Susanna Ricco, Komal Singh, and Krishna Somandepalli. 2022. #SeeltBelt: What families are seeing on TV, *The Geena Davis Institute on Gender in Media*. <https://seejane.org/wp-content/uploads/GDI-What-Families-Are-Watching-TV-2022-Report.pdf>
- [10] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*.
- [11] Bureau of Labor Statistics. 2018. Electronically mediated work: New questions in the contingent worker supplement. <https://www.bls.gov/opub/mlr/2018/article/electronically-mediated-work-new-questions-in-the-contingent-worker-supplement.htm>
- [12] Bureau of Labor Statistics. 2021. Labor Force Statistics from the Current Population Survey. <https://www.bls.gov/cps/cpsaat11.htm>
- [13] Capacì v. Katz & Besthoff, Inc. 1983. Capacì v. Katz & Besthoff, Inc. , 647 pages.
- [14] Elisa Celis, Anay Mehrotra, and Nisheeth Vishnoi. 2019. Toward controlling discrimination in online ad auctions. In *International Conference on Machine Learning*.
- [15] Esmita Charani, Sameed Shariq, Alexandra M Cardoso Pinto, Raabia Farooqi, Winnie Nambatya, Oluchi Mbamalu, Seye Abimbola, and Marc Mendelson. 2023. The use of imagery in global health: An analysis of infectious disease documents and a framework to guide practice. *The Lancet Global Health* 11, 1 (2023), e155–e164.
- [16] Chris Coons, Rob Portman, and Amy Klobuchar. 2021. Coons, Portman, Klobuchar announce legislation to ensure transparency at social media platforms. <https://www.coons.senate.gov/news/press-releases/coons-portman-klobuchar-announce-legislation-to-ensure-transparency-at-social-media-platforms>
- [17] Amit Datta, Anupam Datta, Jael Makagon, Deirdre K Mulligan, and Michael Carl Tschantz. 2018. Discrimination in online advertising: A multidisciplinary inquiry. In *Conference on Fairness, Accountability and Transparency*.
- [18] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated Experiments on Ad Privacy Settings. *Proceedings on Privacy Enhancing Technologies* (2015).
- [19] Duren v. Missouri (U.S. Supreme Court) 1979. Duren v. Missouri (U.S. Supreme Court) , 357 pages. <https://www.law.cornell.edu/supremecourt/text/439/357>
- [20] Cynthia Dwork and Christina Ilvento. 2018. Fairness Under Composition. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*.
- [21] Laura Edelson, Jason Chuang, Erika Franklin Fowler, Michael Franz, and Travis N Ridout. 2021. Universal Digital Ad Transparency. In *TPRC49: The 49th Research Conference on Communication, Information and Internet Policy*.
- [22] Laura Edelson, Tobias Lauinger, and Damon McCoy. 2020. A security analysis of the Facebook Ad Library. In *IEEE Symposium on Security and Privacy (SP)*. IEEE.
- [23] Facebook. 2018. Reviewing Targeting to Ensure Advertising is Safe and Civil. <https://www.facebook.com/business/news/reviewing-targeting-to-ensure-advertising-is-safe-and-civil>
- [24] Facebook. 2019. Updates To Housing, Employment and Credit Ads in Ads Manager. Retrieved January 14, 2022 from <https://www.facebook.com/business/news/updates-to-housing-employment-and-credit-ads-in-ads-manager>
- [25] Facebook. 2020. Facebook's Civil Rights Audit – Final Report. Retrieved February 28, 2022 from <https://about.fb.com/wp-content/uploads/2020/07/Civil-Rights-Audit-Final-Report.pdf>
- [26] Facebook. 2020. Good Questions, Real Answers: How Does Facebook Use Machine Learning to Deliver Ads? Retrieved February 1, 2022 from <https://www.facebook.com/business/news/good-questions-real-answers-how-does-facebook-use-machine-learning-to-deliver-ads/#>
- [27] Facebook. 2022. About Bid and Budget Pacing. Retrieved February 24, 2022 from <https://www.facebook.com/business/help/571961726580148?id=2196356200683573>
- [28] Facebook. 2022. Removing Certain Ad Targeting Options and Expanding Our Ad Controls. <https://www.facebook.com/business/news/removing-certain-ad-targeting-options-and-expanding-our-ad-controls>
- [29] Facebook. 2023. About ad auctions. Retrieved May 10, 2023 from <https://www.facebook.com/business/help/430291176997542>
- [30] Facebook. 2023. Discriminatory Practices. <https://transparency.fb.com/policies/ad-standards/unacceptable-content/discriminatory-practices>
- [31] Facebook. 2023. Personal Attributes (Advertising Standards). Retrieved February 11, 2023 from <https://transparency.fb.com/policies/ad-standards/objectionable-content/personal-attributes>
- [32] Irfan Faizullabhoy and Aleksandra Korolova. 2018. Facebook's Advertising Platform: New Attack Vectors and the Need for Interventions. In *IEEE Workshop on Technology and Consumer Protection (ConPro '18)*.
- [33] George Farkas, Paul L Morgan, Marianne M Hillemeier, Cynthia Mitchell, and Adrienne D Woods. 2020. District-level achievement gaps explain Black and Hispanic overrepresentation in special education. *Exceptional Children* 86, 4 (2020), 374–392.
- [34] Joseph L Gastwirth, Wenjing Xu, and Qing Pan. 2015. Statistical measures for evaluating protected group under-representation: Analysis of the conflicting inferences drawn from the same data in *People v. Bryant* and *Ambrose v. Booker*. *Law, Probability and Risk* 14, 4 (2015), 279–304.
- [35] Danielle Gaucher, Justin Friesen, and Aaron C Kay. 2011. Evidence that gendered wording in job advertisements exists and sustains gender inequality. *Journal of personality and social psychology* 101, 1 (2011), 109.
- [36] James J Gibson. 2002. A theory of direct visual perception. *Vision and Mind: selected readings in the philosophy of perception* (2002), 77–90.
- [37] Google. 2020. Update to Personalized advertising policies: Housing, employment, and credit. <https://support.google.com/adspolicy/answer/9917652?hl=en>
- [38] Mary L Gray, Siddharth Suri, Syed Shoaib Ali, and Deepa Kulkarni. 2016. The crowd is a collaborative network. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*.
- [39] Steven Greenshause. 2004. Abercrombie to Alter Its Marketing to Settle Bias Suit, *The New York Times*. <https://www.nytimes.com/2004/11/16/business/abercrombie-to-alter-its-marketing-to-settle-bias-suit.html>
- [40] Alex Himel. 2018. Helping People Find Jobs and Local Businesses Hire. <https://about.fb.com/news/2018/02/jobs/>
- [41] Susan Illston. 2004. Gonzalez, et al. v. Abercrombie & Fitch Stores, Inc., et al. https://ecommons.cornell.edu/bitstream/handle/1813/80139/Gonzalez_et_al_v_Abercrombie_and_Fitch_Stores.pdf
- [42] Basileal Imana, Aleksandra Korolova, and John Heidemann. 2021. Auditing for Discrimination in Algorithms Delivering Job Ads. In *Proceedings of the Web Conference*.
- [43] Basileal Imana, Aleksandra Korolova, and John Heidemann. 2023. Having your Privacy Cake and Eating it Too: Platform-supported Auditing of Social Media Algorithms for Public Interest. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–33.
- [44] Tanuj Jain, Christopher Lennan, Zubin John, and Dat Tran. 2019. Imagededup. <https://github.com/idealo/imagededup>.
- [45] Valerie Walsh Johnson. 1991. Fair Housing Act-Ragin v. New York Times: Is it time to end the use of models in housing advertisements? *Mem. St. UL Rev* 22 (1991), 401.
- [46] Levi Kaplan, Nicole Gerzon, Alan Mislove, and Piotr Sapiezynski. 2022. Measurement and analysis of implied identity in ad delivery optimization. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 195–209.
- [47] Matthew Kay, Cynthia Matuszak, and Sean A Munson. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*.
- [48] Jon Keegan. 2021. Facebook Got Rid of Racial Ad Categories. Or Did It? <https://themarkup.org/citizen-browser/2021/07/09/facebook-got-rid-of-racial-ad-categories-or-did-it>
- [49] Peter W Kerman. 1974. Sex discrimination in help wanted advertising. *Santa Clara Lawyer* (1974).
- [50] Sara Kingsley, Clara Wang, Alex Mikhalevko, Proteeti Sinha, and Chinmay Kulkarni. 2020. Auditing digital platforms for discrimination in economic opportunity advertising. *arXiv preprint arXiv:2008.09656* (2020).
- [51] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management science* 65, 7 (2019), 2966–2981.
- [52] Victor Le Pochat, Laura Edelson, Tom Van Goethem, Wouter Joosen, Damon McCoy, and Tobias Lauinger. 2022. An Audit of Facebook's Political Ad Policy Enforcement. In *Proceedings of the 31st USENIX Security Symposium*. USENIX

Association.

[53] LinkedIn. 2021. Ad Targeting Discrimination. <https://www.linkedin.com/help/linkedin/answer/a416948>

[54] LinkedIn Corporate Communications. 2019. Language Matters - How words impact men and women in the workplace. <https://business.linkedin.com/content/dam/me/business/en-us/talent-solutions-lodestone/body/pdf/LinkedIn-Language-Matters-Report-FINAL2.pdf>

[55] Alexandra Sasha Lucioni, Christopher Akiki, Margaret Mitchell, and Yacine Jernite. 2023. Stable Bias: Analyzing Societal Representations in Diffusion Models. *arXiv preprint arXiv:2303.11408* (2023).

[56] Paul Messaris. 1997. *Visual persuasion: The role of images in advertising*. Sage.

[57] Meta. 2023. Introducing the AI Sandbox for advertisers and expanding our Meta Advantage suite. <https://www.facebook.com/business/news/introducing-ai-sandbox-and-expanding-meta-advantage-suite>

[58] Meta. 2023. Toward fairness in personalized ads. https://about.fb.com/wp-content/uploads/2023/01/Toward_fairness_in_personalized_ads.pdf

[59] Danaë Metaxa, Michelle A Gan, Su Goh, Jeff Hancock, and James A Landay. 2021. An Image of Society: Gender and Racial Representation and Impact in Image Search Results for Occupations. *Proceedings of the ACM on Human-Computer Interaction* (2021).

[60] Perry L Moriearty. 2009. Framing justice: Media, bias, and legal decisionmaking. *Md. L. Rev.* (2009).

[61] National Center for State Courts. 2010. Jury Managers' Toolbox. https://www.ncsc-jurystudies.org/_data/assets/pdf_file/0026/7478/a-primer-on-fair-cross-section.pdf

[62] National Organization For Women (NOW). 2014. Highlights. <https://now.org/wp-content/uploads/2014/02/Highlights.pdf>

[63] New York Times 1973. Law on Sex-Labeled Job Ads Is Upheld. <https://www.nytimes.com/1973/06/22/archives/law-on-sexlabeled-job-ads-is-upheld.html>

[64] Kazuyuki Okudaira. 2023. Meta to debut ad-creating generative AI this year, CTO says. <https://asia.nikkei.com/Business/Technology/Meta-to-debut-ad-creating-generative-AI-this-year-CTO-says>

[65] Barbara Orutay. 2022. Facebook parent Meta's revenue, profit decline amid ad slump. <https://apnews.com/article/technology-business-earnings-reports-8d24203456813802a4239e54d1b444d4>

[66] Orestis Papakyriakopoulos and Arwa M Mboya. 2021. Beyond algorithmic bias: a socio-computational interrogation of the google search by image algorithm. *Social Science Computer Review* (2021).

[67] Timothy D Pippert, Laura J Essenburg, and Edward J Matchett. 2013. We've got minorities, yes we do: Visual representations of racial and ethnic diversity in college recruitment materials. *Journal of Marketing for Higher Education* (2013).

[68] Peter Romer-Friedman and Mitra Ebadolahi. 2022. Real Women in Trucking. <http://guptawessler.com/real-women-in-trucking/>

[69] Sheryl Sandberg. 2019. Doing More to Protect Against Discrimination in Housing, Employment and Credit Advertising. Retrieved January 13, 2022 from <https://about.fb.com/news/2019/03/protecting-against-discrimination-in-ads/>

[70] Piotr Sapiezynski, Avijit Ghosh, Levi Kaplan, Alan Mislove, and Aaron Rieke. 2022. Algorithms that "Don't See Color": Measuring Biases in Lookalike and Special Ad Audiences. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*.

[71] Richard Seltzer, John M Copacino, and Diana Roberto Donahoe. 1995. Fair Cross-Section Challenges in Maryland: An Analysis and Proposal. *U. Balt. L. Rev.* 25 (1995), 127.

[72] Márcio Silva, Lucas Santos de Oliveira, Athanasios Andreou, Pedro Olmo Vaz de Melo, Oana Goga, and Fabrício Benevenuto. 2020. Facebook ads monitor: An independent auditing system for political ads on Facebook. In *Proceedings of The Web Conference*.

[73] Veronica Smith. 2008. Visual Persuasion: Issues in the translation of the visual in advertising. *Meta: Journal des Traducteurs/Meta: Translators' Journal* (2008).

[74] Vera Sosnovik and Oana Goga. 2021. Understanding the Complexity of Detecting Political Ads. In *Proceedings of the Web Conference*.

[75] Till Speicher, Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Ribeiro, George Arvanitakis, Fabricio Benevenuto, Krishna P Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Potential for discrimination in online targeted advertising. In *Conference on Fairness, Accountability and Transparency*.

[76] Michael Spence. 1973. Job Market Signaling. *The Quarterly Journal of Economics* (1973).

[77] Anne Sladens. 2019. Goldman Sachs Removed This One Word From Some Recruiting Materials—and Saw Female Hires Soar. <https://fortune.com/2019/12/10/goldman-sachs-removed-this-one-word-from-some-recruiting-materials-and-saw-female-hires-soar/>

[78] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Commun. ACM* (2013).

[79] Henri Tajfel and John C. Turner. 1974. Social identity and intergroup behaviour. *Social science information* (1974).

[80] Ruth E. Thaler-Carter. 2001. Diversify Your Recruitment Advertising. More companies are adding a diversity element to their recruitment ads. *HR MAGAZINE* 46, 6 (2001), 92–101.

[81] The United States Department of Justice. 2022. Justice Department Secures Groundbreaking Settlement Agreement with Meta Platforms, Formerly Known as Facebook, to Resolve Allegations of Discriminatory Advertising. <https://www.justice.gov/opa/pr/justice-department-secures-groundbreaking-settlement-agreement-meta-platforms-formerly-known>

[82] Ariana Tobin. 2019. HUD sues Facebook over housing discrimination and says the company's algorithms have made the problem worse. <https://propublica.org/article/hud-sues-facebook-housing-discrimination-advertising-algorithms>

[83] Ariana Tobin and Jeremy B. Merrill. 2018. Facebook Is Letting Job Advertisers Target Only Men, *ProPublica*. <https://www.propublica.org/article/facebook-is-letting-job-advertisers-target-only-men>

[84] U.S. Department of Education. 2016. Racial and Ethnic Disparities in Special Education. <https://www2.ed.gov/programs/osepidea/618-data/LEA-racial-ethnic-disparities-tables/disproportionality-analysis-by-state-analysis-category.pdf>

[85] U.S. Equal Employment Opportunity Commission. 1981. CM-607 Affirmative Action. <https://www.eeoc.gov/laws/guidance/cm-607-affirmative-action>

[86] U.S.C. 1964. Title VII of the Civil Rights Act of 1964. <https://www.eeoc.gov/statutes/title-vii-civil-rights-act-1964>

[87] U.S.C. 1968. Fair Housing Act. <https://www.justice.gov/crt/fair-housing-act-1>

[88] H. Jack Walker, Hubert S. Feild, Jeremy B. Bernerth, and J. Bret Becton. 2012. Diversity cues on recruitment websites: Investigating the effects on job seekers' information processing. *Journal of Applied Psychology* 97, 1 (2012), 214.

[89] Kenneth White, Megan McCoy, Kim Love, Eun Jin Kwak, Erin Bruce, and John Grable. 2019. The role of signaling when promoting diversity and inclusion at the firm level: A financial advisory professional case study. *Advances in Business Research* 9, 1 (2019), 1–16.

[90] Kyra Yee, Uthaipon Tanitpongipat, and Shubhanshu Mishra. 2021. Image cropping on Twitter: Fairness metrics, their limitations, and the importance of representation, design, and agency. *Proceedings of the ACM on Human-Computer Interaction* (2021).

[91] Lara Zarum. 2018. Some Viewers Think Netflix Is Targeting Them by Race. Here's What to Know. <https://www.nytimes.com/2018/10/23/arts/television/netflix-race-targeting-personalization.html>

A APPENDIX

A.1 Download of Truck Driver and Nurses Advertiser Data

A.1.1 Search Query Keywords for Obtaining the List of Advertisers.

- Truck Driver: truck driver b, ... truck driver j, truck driver l, ... truck driver z
Further, we narrowed our queries and searched through keywords containing all two letter combinations for the letters a and k, since the number of search results for the single letters were too many for the page to load completely. The specific keywords include: truck driver aa, truck driver ab, ... truck driver az and truck driver ka, truck driver kb, ..., truck driver kz
- Nurses: nurses a, ..., nurses z

A.1.2 Creation of Advertiser URL for Scraping.

For each unique advertiser in the search results, we obtained the Page ID and created an advertiser specific URL which was used to scrape the respective pages.

The Base URL: <https://www.facebook.com/ads/library/?>

URL Parameters:

```
{
  'view_all_page_id' : page_id,
  'search_type': 'page',
  'media_type': 'image_and_meme',
  'country': 'US',
  'active_status': 'all',
  'ad_type': 'employment_ads'
}
```

Sample URLs:

- CDLLife Jobs: https://www.facebook.com/ads/library/?view_all_page_id=201384326737496&search_type=page&media_type=image_and_meme&country=US&active_status=all&ad_type=employment_ads
- Total Nurses Network: https://www.facebook.com/ads/library/?view_all_page_id=85406439698&search_type=page&

media_type=image_and_meme&country=US&active_status=all&ad_type=employment_ads

A.2 Additional Data and Analysis of Job Advertisers Image Selection

A.2.1 Job Advertisers Analyzed According to Gender and Race.

Data Collection Period: The identified ads were scraped from the Ad Library during October - November 2021. We also scraped data one year later, in October 2022 to perform a longitudinal analysis. However, compared to 2021, we observed many fewer advertisers (13/18) and each advertiser had significantly fewer ad campaigns and in-turn very few images to facilitate any meaningful analysis. This may be a result of the prevailing economic conditions in 2022 - a general fall in ad sales on Facebook and stiff competition from TikTok [65]. Therefore, other than for Monster.com, we restrict our analysis in this paper to ads which were run in 2021.

Standard Errors: We calculate the 95% confidence interval of the percentages in Table 4 according to $p \pm 1.96\sqrt{\frac{p(1-p)}{n}}$, where p is the fraction of women, White and Black people, and n is the total number of annotations across all the images of an advertiser. Here, we annotated all images (not just distinct images) by 2 distinct annotators on MTurk.

Interannotator Agreement: We used Cohen Kappa to calculate the inter-annotator agreement averaged out across all annotators for each label. The kappa value for gender is 0.86 (2021), 0.72 (2022) and race is 0.64 (2021), 0.52 (2022). The moderate agreement is not surprising due to the subjective nature of perceived gender and race classification.

The decrease in the annotator agreement in 2022 compared to that in 2021, may in part be a result of advertisers using more avatars in their job ad images, rather than human subjects. The classification of perceived gender and race of avatars can be more challenging than that for human subjects.

Additional data and analysis is presented in Tables 4, 5.

A.2.2 Job Advertisers Analyzed According to Gender. Additional data and analysis is presented in Tables 6, 7, and Figures 6, 7.

Table 4: Columns 1-5: Occupation-specific Job Advertisers (Rows 2-16) and Job Aggregator (Rows 17-19) Data Summary from 2021. Columns 6-8: Percentage of Women, White and Black People Corresponding to the Respective BLS Data, Across All Job Ad Images for data obtained in 2021. The data from our analysis with 95% confidence interval is present first followed by the BLS Data within parenthesis

Serial Number	Advertiser	Occupation	# of Images	# of Distinct Images	# of Distinct Images Containing People	% Women (BLS)	% White (BLS)	% Black (BLS)
1	BestBuy	Sales and Office Occupations	147	3	3	45 ± 6 (61)	47 ± 6 (79)	25 ± 5 (13)
2	Doordash	GIG	108	16	12	39 ± 7 (47)	17 ± 5 (79)	22 ± 6 (12)
3	Eataly	Chefs and Head Cooks	48	14	13	30 ± 9 (18)	38 ± 10 (65)	18 ± 8 (15)
4	Geico Careers	Insurance Sales Agents	33	12	11	63 ± 12 (50)	34 ± 11 (80)	44 ± 12 (11)
5	Drive with HopSkipDrive	Bus Drivers, School	181	8	8	71 ± 5 (59)	42 ± 5 (73)	0 (22)
6	Instacart	GIG	11	4	4	91 ± 12 (47)	73 ± 19 (79)	9 ± 12 (12)
7	Drive with Lyft	GIG	383	35	23	75 ± 3 (47)	38 ± 3 (79)	19 ± 3 (12)
8	Nationwide Job Search for Education	Education and Training Occupations	55	45	45	49 ± 9 (74)	54 ± 9 (82)	19 ± 7 (10)
9	Nationwide Job Search for Information Technology	Computer and Mathematical Occupations	20	13	12	22 ± 13 (25)	67 ± 15 (65)	11 ± 10 (9)
10	Nurse Recruiter	Healthcare Practitioners	29	8	8	91 ± 7 (74)	82 ± 10 (76)	12 ± 8 (12)
11	NYPD Recruit	Police Officers	6	6	6	48 ± 28 (17)	17 ± 21 (85)	66 ± 27 (11)
12	Safeway	Laborers and Freight, Stock, and Material Movers, Hand	16	6	6	39 ± 17 (21)	17 ± 13 (72)	67 ± 16 (19)
13	TSA	Protective Service Occupations	39	11	11	55 ± 11 (24)	32 ± 10 (75)	35 ± 11 (19)
14	Uber	GIG	40	21	19	32 ± 10 (47)	38 ± 11 (79)	8 ± 6 (12)
15	UPS Jobs	Laborers and Freight, Stock, and Material Movers, Hand	27	21	17	47 ± 13 (21)	58 ± 13 (72)	25 ± 12 (19)
1	Monster	Multiple	499	104	102	44 ± 3	36 ± 3	37 ± 3
2	SimplyJobs	Multiple	104	89	45	45 ± 7	52 ± 7	14 ± 5
3	Talent	Multiple	109	77	38	46 ± 7	75 ± 6	2 ± 2

Table 5: Advertiser List Analyzed According to Gender and Race Corresponding to Page ID on Facebook

Advertiser	Page ID
BestBuy	12699262021
Doordash	534754226586678
Eataly	443671242427553
Geico Careers	52380474954
Drive with HopSkipDrive	103480144441214
Instacart	369288959794283
Drive with Lyft	1023523827667350
Nationwide Job Search for Education	102379295254304
Nationwide Job Search for Information Technology	110279211123776
Nurse Recruiter	112414138834087
NYPD Recruit	143154955728131
Safeway	78143372410
TSA	782005221949498
Uber	120945717945722
UPSJobs	93397977942
Monster	87877000648
SimplyJobs	697292896998339
Talent	108919760729002

Table 6: Image selections of truck driver advertisers who use ≥ 15 distinct images and ≥ 5 distinct images containing people.
 We calculate the 95% confidence interval of the fraction in Column 6 according to $p \pm 1.96\sqrt{\frac{p(1-p)}{n}}$, where p is the fraction of people of stereotypical gender and n is the total number of annotations across all the distinct images of an advertiser.

Advertiser Name	# of Images	# of Distinct Images	# of Distinct Images Depicting People	# of Distinct Images Depicting Men	Fraction of Men Images Among Distinct Images Depicting People
Best Driving Job	169	83	30	21	0.70 ± 0.07
Best Truck Driver Job	315	199	47	40	0.85 ± 0.03
Better Driver Jobs	50	37	21	17	0.81 ± 0.09
Big Truck Driving Jobs	83	57	20	16	0.80 ± 0.07
C R England	72	71	45	42	0.93 ± 0.04
CDL Job Now	143	80	35	26	0.74 ± 0.07
CDLLife Team Drivers	24	24	5	4	0.80 ± 0.11
CDLLife.com	463	241	40	31	0.78 ± 0.04
CDLLife Jobs	278	160	27	23	0.85 ± 0.04
Drivers Job Choice	26	25	14	11	0.79 ± 0.11
Findatruckerjob.com	231	149	58	53	0.91 ± 0.03
Go Your Way Maine	19	19	9	6	0.67 ± 0.15
Hammond Lumber Company Careers	31	19	10	8	0.80 ± 0.13
Higher Paying Driver Jobs	58	53	18	17	0.94 ± 0.04
Hiremaster	150	87	8	8	1.00
HMD Trucking Inc	27	26	20	13	0.65 ± 0.13
Knight Transportation	36	34	32	25	0.78 ± 0.10
Nationwide Job Search For Transportation	27	22	21	20	0.95 ± 0.06
Ryder System Jobs	47	33	26	19	0.73 ± 0.11
The Jobs Driver	56	32	7	6	0.86 ± 0.09
Top Pay For Drivers	27	19	5	4	0.80 ± 0.13
Ultimate Trucking Jobs	75	50	5	5	1.00
Total	2,407	1,520	503	415	Average: 0.82

Table 7: Image selections of nurses advertisers who use ≥ 15 distinct images and ≥ 5 distinct images containing people.
 We calculate the 95% confidence interval of the fraction in Column 6 according to $p \pm 1.96\sqrt{\frac{p(1-p)}{n}}$, where p is the fraction of people of stereotypical gender and n is the total number of annotations across all the distinct images of an advertiser.

Advertiser Name	# of Images	# of Distinct Images	# of Distinct Images Depicting People	# of Distinct Images Depicting Women	Fraction of Women Images Among Distinct Images Depicting People
Applicat Healthcare	112	68	48	33	0.69 ± 0.08
Ascension Careers	41	34	31	23	0.74 ± 0.10
Camden Clark Medical Center	21	20	18	14	0.78 ± 0.13
Children's Hospital Of The Kings Daughter's (CHKD)	39	22	14	11	0.79 ± 0.12
Consumer Direct Care Network	20	19	19	17	0.89 ± 0.10
Dartmouth Health Careers	24	23	20	12	0.60 ± 0.14
Encompass Health Careers	62	25	23	17	0.74 ± 0.12
Gifted Healthcare	27	25	20	14	0.70 ± 0.13
Health Carousel Travel Nursing	25	17	8	7	0.88 ± 0.11
John Knox Village Of Central Florida	68	44	23	11	0.48 ± 0.10
Kettering Health	50	31	31	20	0.65 ± 0.12
Memorial Health Careers	25	24	23	19	0.83 ± 0.11
MyMichigan Health	24	20	19	12	0.63 ± 0.15
NuWest Travel Nursing	67	37	24	13	0.54 ± 0.11
Providence Careers	19	19	11	8	0.73 ± 0.14
Providence Health Services Careers	24	23	13	9	0.69 ± 0.13
The Guthrie Clinic	17	16	16	13	0.81 ± 0.14
VNS Health	19	19	19	17	0.89 ± 0.10
Total	684	486	380	270	Average: 0.73



Figure 6: Best Auto Service & Repair is an example of an advertiser using people of demographic characteristics similar to their existing non-diverse workforce in job ad images. Specifically, they used 100 distinct images containing people where nearly all images (93/100) exclusively depict men. This is reflective of the 2021 BLS data for “Installation, maintenance, and repair occupations” where women make up only 4.2% of the current workforce.



Figure 7: Diverse Outreach by DHS Logistics Solution: 3 out of 6 distinct images contain people and all 3 depict women.