



February 9, 2023

U.S. Equal Employment Opportunity Commission
Headquarters
131 M Street, NE
Washington, DC 20507

Submitted via Federal eRulemaking Portal only at <https://www.regulations.gov>

RE: EEOC Draft Strategic Enforcement Plan, 2023–2027

We write to provide comments in response to the Equal Employment Opportunity Commission (“EEOC”)’s Draft Strategic Enforcement Plan for 2023–2027, published on January 10, 2023.¹

Upturn is a non-profit organization that advances equity and justice in the design, governance, and use of technology. Through research and advocacy, we drive policy change by investigating specific ways that technology and automation shape people’s opportunities, particularly in historically disadvantaged communities.

Our comments primarily address the **Subject Matter Priorities** set forth under Principle 1 of the Draft Plan. We identify the ways in which our past research supports the Commission’s intended priorities and explain where the Commission can clarify or refine the Draft Plan to further support civil rights in the workplace.

Priority #1: Eliminating Barriers to Recruitment and Hiring. The Draft Plan identifies a number of recruitment and hiring practices and policies that discriminate against historically disadvantaged individuals, including racial and ethnic minorities; women; older workers; and people with disabilities.

First, the Commission indicates an intent to focus on the use of automated systems to target job advertisements, recruit applicants, or make or assist in hiring decisions,

¹ Equal Employment Opportunity Commission, Draft Strategic Enforcement Plan, 88 Fed. Reg. 1379–1385 (Jan. 10, 2023), available at <https://www.govinfo.gov/content/pkg/FR-2023-01-10/pdf/2023-00283.pdf> [hereinafter, “Draft Strategic Enforcement Plan”].

where such systems intentionally exclude or adversely impact protected groups.² The Draft Plan thus correctly recognizes that the use of automated systems can lead to both intentional discrimination and facially neutral practices that have an unjustified disparate impact.

It is especially vital that the Commission direct its resources, including its investigative tools, to illuminate how individuals are forced to interact with automated systems when seeking employment, and how these systems, even when ostensibly “neutral” and “validated,” have disparate impacts on historically marginalized communities. For applicants, automated employment systems are opaque and difficult to understand. As EEOC Commissioner Keith Sonderling has noted, “enforcement has been difficult in this area because the employees do not know that they’re being subjected to this technology.”³ Employers, too, may struggle to understand how recruitment and hiring tools created by third party vendors operate in practice. The EEOC can and must play a central role in rectifying existing information asymmetries in modern recruitment and hiring processes, since individuals seeking employment cannot identify, much less attempt to redress, discrimination without a clear understanding of how various automated systems have been applied to them.

Moreover, as we explained in Help Wanted, our 2018 report on predictive tools used in hiring processes, recruitment and hiring usually consist of a series of small decisions, rather than a single determination.⁴ For its enforcement to be most effective, the EEOC must use its investigative and other authorities to study the systems employers use to screen job applicants, focusing on processes where candidates are scored, ranked, or rejected, and on how each of these processes might impact applicants of color, women, applicants with disabilities, older people, and other groups that have been systematically denied equal access to employment opportunities. In prioritizing understanding the use of automated systems in recruitment and hiring, the Commission can expose each decision point and deepen the public’s understanding of how various systems may interact to obstruct employment access.

² *Id.* at 1381.

³ Paige Smith, *Artificial Intelligence Bias Needs EEOC Oversight, Official Says*, Bloomberg Law, Sept. 1, 2021, <https://news.bloomberglaw.com/privacy-and-data-security/artificial-intelligence-bias-needs-eeoc-oversight-official-says>.

⁴ See generally Miranda Bogen & Aaron Rieke, Upturn, Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias (Dec. 2018), <https://www.upturn.org/reports/2018/hiring-algorithms/> [hereinafter, “Help Wanted”].

Second, the Commission indicates an intent to focus on job advertisements that exclude or discourage certain demographic groups from applying for jobs.⁵ On this issue, we urge the EEOC to proactively investigate online ad delivery systems and algorithms, not just facially discriminatory advertising copy.

It has long been unlawful to steer job ads away from people based on their race, gender, age, and other protected statuses.⁶ Over the last decade, however, an increasing number of Americans have relied on the internet and online platforms to search for and apply to jobs.⁷ Online platforms use algorithms to decide which people will see which job ads. In many cases, such algorithms deliver ads based on protected characteristics—like race, sex, age, or proxies for them.⁸

Recently, Gupta Wessler LLC and Upturn filed a class action with the EEOC against Meta Platforms, Inc. (“Meta” or “Facebook”) on behalf of Real Women in Trucking, demonstrating that Facebook, a social media application owned by Meta, disproportionately steers certain types of job ads away from users based on their gender and age.⁹ For example, an employer published a job advertisement on Facebook seeking to hire truck drivers in the Durham/Raleigh, North Carolina area. The eligible audience for this advertisement was people of all genders who were 18 or older. But, when Facebook’s own ad delivery algorithm decided which people would see this advertisement, Facebook

⁵ Draft Strategic Enforcement Plan, *supra* note 1, at 1381.

⁶ *See, e.g.*, Charge of Discrimination, Real Women in Trucking v. Meta Platforms at 7–12, Dec. 1, 2022, *available at* <https://www.upturn.org/static/files/20221201-real-women-in-trucking-eeoc-charge.pdf> [hereinafter, “RWIT v. Meta”].

⁷ *Id.* at 2 (noting that, in 2015, the Pew Research Center found that 90% of the people who had searched for work in the previous two years relied on the internet to do so, and 84% had submitted a job application online).

⁸ Upturn has researched Facebook advertising for several years, producing various reports identifying the ways in which the world’s largest social media platform has perpetuated discrimination. *See, e.g.*, Piotr Sapiezynski, et al., Algorithms that Don’t See Color: Comparing Biases in Lookalike and Special Ad Audiences, Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (Dec. 2019; rev’d May 2022), <https://arxiv.org/abs/1912.07579>; Muhammad Ali, et al., Discrimination Through Optimization: How Facebook’s Ad Delivery Can Lead to Skewed Outcomes, Proceedings of the ACM on Human-Computer Interaction 2019 (Apr. 2019), <https://arxiv.org/abs/1904.02095>; Aaron Rieke & Miranda Bogen, Upturn, Leveling the Platform: Real Transparency for Paid Messages on Facebook (May 2018), <https://www.upturn.org/work/leveling-the-platform/>; Aaron Rieke, Facebook, Race, and Ads: The Story So Far and What Should Happen Next (Dec. 2016), <https://www.upturn.org/work/facebook-race-and-ads/>.

⁹ RWIT v. Meta, *supra* note 6, at 22–39 & Exhibit A (collecting specific examples of gender and age discrimination in Facebook’s delivery of job advertisements).

showed the ad to 94% men and only 5% women, and to only 11% people over the age of 55, even though people 55 and older make up more than 28% of Facebook users who are looking for a job.¹⁰

Shortly after we filed Real Women in Trucking’s charge with the EEOC, Meta announced, as part of a housing discrimination lawsuit settlement with the Department of Justice, a new process ostensibly designed to minimize discrimination in ad delivery across Meta platforms.¹¹ Meta has stated it plans to voluntarily extend this process, called the Variance Reduction System (“VRS”), to employment ads (as well as credit ads) later this year. According to Meta, VRS “uses new machine learning technology in ad delivery so that the actual audience that sees an ad more closely reflects the eligible target audience for that ad.”¹²

It remains to be seen whether VRS will sufficiently reduce discrimination against protected groups in the delivery of Facebook ads regarding major life opportunities, including employment.¹³ We have two preliminary concerns.

First, VRS applies only to certain types of ads: housing, employment, and credit (“HEC”).¹⁴ Advertisers are responsible for classifying their ads as HEC-related before they

¹⁰ *Id.* at 4.

¹¹ See Roy L. Austen, Jr., Meta, An Update on Our Ads Fairness Efforts, Jan. 9, 2023, <https://about.fb.com/news/2023/01/an-update-on-our-ads-fairness-efforts/> [hereinafter, “Meta Fairness Update”]; Meta AI, A New System to Help Ensure Ads Are Delivered Fairly to Different Demographic Groups, Jan. 9, 2023, <https://ai.facebook.com/blog/advertising-fairness-variance-reduction-system-vrs/>; U.S. Department of Justice, Press Release, Justice Department and Meta Platforms Inc. Reach Key Agreement as They Implement Groundbreaking Resolution to Address Discriminatory Delivery of Housing Advertisements, Jan. 9, 2023 <https://www.justice.gov/opa/pr/justice-department-and-meta-platforms-inc-reach-key-agreement-they-implement-groundbreaking> [hereinafter, “DOJ Press Release”].

For a technical analysis of VRS, see Miranda Bogen, et al., Meta, Toward Fairness in Personalized Ads (Jan. 2023), <https://bit.ly/3x8n17f> [hereinafter, “Toward Fairness”].

¹² Meta Fairness Update, *supra* note 11.

¹³ As part of its settlement with the Department of Justice, Meta has agreed to certain VRS compliance metrics. For example, according to the Department of Justice, by December 31, 2023, “for the vast majority of housing advertisements on Meta platforms, Meta will reduce variances to less than or equal to 10% for 91.7% of those advertisements for sex and less than or equal to 10% for 81.0% of those advertisements for estimated race/ethnicity.” DOJ Press Release, *supra* note 11.

Although Meta has stated its intention to “voluntarily” apply VRS to employment and credit ads, no such firm compliance metrics exist for these areas (at least, not publicly). Without such metrics, the public cannot assess how effective VRS is in reducing discrimination in ad delivery for employment or credit advertisements.

¹⁴ Meta Fairness Update, *supra* note 11.

determine the eligible audience for the advertisement. Only if the advertiser self-identifies an advertisement as HEC-related will the advertiser immediately move into a special portal that limits the advertiser’s options for targeting the advertisement.¹⁵ In this portal, an advertiser is prohibited from ad targeting “based on gender, age, or interests that appear to describe people of a certain race, religion, ethnicity, sexual orientation, disability status, or other protected class.”¹⁶ Similarly, if an HEC advertiser wishes to target an advertisement based on location, the “location targeting must have a minimum 15-mile radius.”¹⁷

Yet the drop-down menu that an advertiser must use to classify an ad as HEC-related is not a “required” field. Although Facebook claims to apply a “classifiers” algorithm to block HEC-related ads that advertisers have failed to properly classify, our investigation in the Real Women in Trucking case indicates that this “classifiers” algorithm does not always work.¹⁸ Given that Facebook does not require all advertisers to say whether an advertisement is HEC-related, and the apparent weakness of its “classifiers” algorithm, it remains possible for advertisements related to key life opportunities (including employment) to be unclassified as such, thereby evading VRS (or other corrective systems). For VRS and Facebook’s limitations on discriminatory advertiser audience selections to apply comprehensively, Facebook should require advertisers to affirmatively select an ad category before they can proceed.¹⁹

Second, VRS aims to narrow the gap between an eligible ad audience and the users actually shown an ad. Therefore, whether VRS is effective in reducing bias in ad delivery remains a function of the eligible ad audience. If advertisers can persist in targeting users in discriminatory ways (thereby skewing the eligible ad audience), the corrective impact of VRS will be limited.

¹⁵ RWIT v. Meta, *supra* note 6, at 15.

¹⁶ Toward Fairness, *supra* note 11, at 9–10.

¹⁷ *Id.*

¹⁸ RWIT v. Meta, *supra* note 6, at 16 (describing the “classifiers” algorithm) & 37–39 (identifying employment-related advertisements which Facebook’s “classifiers” algorithm failed to properly classify).

¹⁹ As of January 25, 2023, the “Special Ad Categories” drop-down menu on the Facebook Ad Manager portal included four options: (1) credit, (2) employment, (3) housing, and (4) social issues, elections or politics. An additional category, “none of the above,” should be added, and selection of one category should be required before the advertiser can proceed. This will require classification of all advertisements on Facebook, enabling more accurate and robust analysis of Facebook’s ad delivery system.

Facebook should provide additional public information about its ad delivery system to help regulators and members of the public understand bias in that system. For example, Facebook could compile and provide summary statistics on advertiser accounts that have the highest “disparity” rates even after VRS has been applied. In its Ad Library, Facebook could publish demographic comparisons for each particular advertisement run—including for housing, employment, and credit ads—that show the eligible target audience distributions as compared to the actual delivery distributions across each protected class. These kinds of metrics would enable regulators and advocates to assess the effectiveness of VRS and the impacts of Facebook’s ad delivery system overall.²⁰

Time and again, Meta has moved to make changes in its advertising systems only after advocates and government regulators have identified discriminatory trends and taken legal action against Facebook. Robust oversight is essential, and the EEOC can play an important role (including by working with the Department of Justice) in assessing the actual impact of VRS on Meta’s ad delivery.

Of course, Meta is not the only online platform that uses algorithms to connect job seekers with relevant employment opportunities.²¹ Despite the proliferation of online platforms and their increasing centrality in an individual’s ability to obtain employment, very little is understood about the platforms’ algorithmic ad ranking and recommendation systems. The EEOC can and should prioritize investigating how these systems affect access to employment opportunities. Doing so will not only help ferret out discrimination against protected groups, but will also diminish the persistent information asymmetries that impede individuals from asserting their civil rights under equal employment laws. For example, the Commission could study gender and race disparities on job platforms like LinkedIn, ZipRecruiter, Indeed, and Monster. These platforms allow recruiters to search for potential job candidates based on inputs such as job titles, location, and skills. However, after these simple inputs, relatively little is understood about how these

²⁰ Such data could also help advertisers understand the demographics of their audience selections in the first place, which would make it possible for advertisers to know whether demographic disparities in the delivery of their ads are due to Facebook’s algorithmic decision making or, alternatively, due to their own audience selections.

²¹ See, e.g., Testimony of ReNika Moore (ACLU), Navigating Employment Discrimination in AI and Automated Systems: A New Civil Rights Frontier, Meeting of the EEOC, Jan. 31, 2023, <https://www.eeoc.gov/meetings/meeting-january-31-2023-navigating-employment-discrimination-ai-and-automated-systems-new/moore>, at nn. 64–69 and accompanying text (discussing such companies as LinkedIn, ZipRecruiter, Indeed, CareerBuilder, and Monster).

companies' own algorithmic ranking processes affect which candidates are shown to recruiters.²²

Third, the EEOC indicates an intent to focus on restrictive application processes or systems, including online systems that are difficult for individuals with disabilities or other protected groups to access.²³ We agree that this is an important priority.

The EEOC can and should play a proactive role in obtaining additional information from employers about their selection procedures. In July 2021, Upturn published *Essential Work*, a report identifying various technologies that applicants for low-wage hourly jobs encounter.²⁴ We concluded that many employers apply a blend of traditional selection procedures and new hiring technologies when evaluating applicant pools.²⁵ Here, too, civil right advocates and individual workers encounter a significant information asymmetry. It is nearly impossible to fully understand an employer's digital hiring practices from the outside. More sustained EEOC attention, and updated regulations where appropriate, may incentivize employers to think more critically about the ways in which their hiring practices perpetuate or exacerbate discrimination.

Take, for example, personality tests. Personality tests are often "normed" around a largely white, middle-class population, which could contribute to discrimination against

²² Job platforms such as LinkedIn and Monster incorporate a variety of data, including users' behavioral data, into their algorithms, which then curate a list of recommendations for both job seekers and employers. These algorithms are generally optimized towards generating applications from job seekers and the likelihood of a successful hire for an employer. LinkedIn's algorithm eventually detected "behavioral patterns exhibited by groups with particular gender identities" causing the algorithm to adjust recommendations in a way that disadvantaged women. *See, e.g.,* Sheridan Wall & Hilke Schellmann, *LinkedIn's job-matching AI was biased. The company's solution? More AI*, MIT Technology Review, June 23, 2021, <https://www.technologyreview.com/2021/06/23/1026825/linkedin-ai-bias-ziprecruiter-monster-artificial-intelligence/>. To mitigate biases caused by its algorithms, LinkedIn has created features like "Representative Results" and "Diversity Nudges," which proactively infer job seekers' gender to create a "gender representative ranking approach" on LinkedIn's Recruiter platform. *See* Sahin Cem Geyik & Krishnaram Kenthapadi, *Building Representative Talent Search at LinkedIn*, LinkedIn Engineering, Oct. 10, 2018, <https://engineering.linkedin.com/blog/2018/10/building-representative-talent-search-at-linkedin> (discussing "Representative Results"); LinkedIn Help, *Diversity Nudges in Recruiter - Overview*, <https://www.linkedin.com/help/recruiter/answer/a794260?trk=hc-articlePage-sidebar> (discussing "Diversity Nudges").

²³ Draft Strategic Enforcement Plan, *supra* note 1, at 1381.

²⁴ *See generally*, Aaron Rieke, Urmila Janardan, Mingwei Hsu & Natasha Duarte, Upturn, *Essential Work: Analyzing the Hiring Technologies of Large Hourly Employers* (July 2021), <https://www.upturn.org/reports/2021/essential-work/> [hereinafter, "Essential Work"].

²⁵ *Id.*; *see also* Help Wanted, *supra* note 4, at 1.

other groups.²⁶ “Norming” is the process by which numerical test scores are given qualitative meaning: the population around which a test is normed will be set to a distribution, with the most common scores constituting what is most “normal.”²⁷ If the population is homogenous and not representative of those to whom the personality test will ultimately be administered, the test will likely be inaccurate for individuals who do not fit the original population’s characteristics. Specifically, such tests may be biased against those who do not fit a certain cultural, class, religious, able-bodied, neurotypical or racial “norm.”²⁸

Personality tests can screen out individuals with depression or anxiety, while a game-based assessment may screen out an applicant with ADHD.²⁹ These results may violate the Americans with Disabilities Act (“ADA”), which requires that traits considered in an employment assessment be job-related and necessary for the job in question.³⁰ As the Center for Democracy and Technology has noted, employers “have an obligation to reject pseudo-science, focus on what specific skills are required for the position, and think critically about how they can fairly and accurately assess those skills.”³¹

In our research, however, we found that most personality tests have little or no relationship to the actual work involved in a particular job.³² In addition to potential violations of the ADA, we have identified hiring questionnaires that reflect union-avoidance preferences.³³ For example, we found instances where employers asked applicants if they “question authority” or prioritize their well-being over job performance.³⁴ We also came across questions that appeared to gauge a worker’s

²⁶ Essential Work, *supra* note 24, at 25 & n.54 (internal quotation marks omitted).

²⁷ *Id.* at 25 & n.55.

²⁸ *Id.* at 25–26.

²⁹ Center for Democracy and Technology, Algorithm-Driven Hiring Tools: Innovative Recruitment or Expedited Disability Discrimination? 3, 12 (2020), <https://cdt.org/wp-content/uploads/2020/12/Full-Text-Algorithm-driven-Hiring-Tools-Innovative-Recruitment-or-Expedited-Disability-Discrimination.pdf>.

³⁰ *Id.* at 12.

³¹ *Id.*

³² Essential Work, *supra* note 24, at 25 & nn. 51–53.

³³ *Id.* at 26–27.

³⁴ *Id.* at 27.

willingness to work for low pay; for example, one question posed a false dichotomy, asking applicants if they preferred a job where “there are high performance expectations” or one where employees are “highly compensated for [their] work.”³⁵ Such questions may screen out job applicants who would agitate for better working conditions (including higher pay) or who would unionize.³⁶

We urge the EEOC to increase its use of Commissioner charges and directed investigations to tackle systemic discrimination resulting from restrictive application processes or systems. We also urge the Commission to use its statutory research authority creatively and aggressively to help develop a more detailed and accurate picture of how employers—especially large hourly employers—are using hiring technologies.³⁷

Fourth, the Draft Plan identifies as a priority “screening tools or requirements that disproportionately impact workers based on their protected status, including those facilitated by artificial intelligence or other automated systems, pre-employment tests, and background checks.”³⁸ As explained, the use of such tools has proliferated rapidly, yet most Americans are unable to identify the precise ways such tools affect their working lives.

Background checks, in particular, can be a significant barrier to employment for those who need it most.³⁹ They often include credit reports and criminal records, both of which reflect racial discrimination in the financial and criminal legal systems.⁴⁰ A mere request for an applicant’s “consent” to conduct a background check can dissuade an individual with a criminal record or poor credit history from submitting a job application.⁴¹ Moreover, background check reports are plagued with errors, such as records “matched” to the wrong person and inaccurate case dispositions.⁴²

³⁵ *Id.*

³⁶ *Id.* at 27.

³⁷ *See, e.g.*, 42 U.S.C. § 2000e–4(g)(5).

³⁸ Draft Strategic Enforcement Plan, *supra* note 1, at 1381.

³⁹ Essential Work, *supra* note 24, at 21 & n.38.

⁴⁰ *Id.* at 21 & nn. 39–40.

⁴¹ *Id.* at 21–22.

⁴² *Id.* at 22.

The EEOC has acknowledged the potential harms of background checks. In its 2012 guidance on criminal records, the Commission noted that “national data supports a finding that criminal record exclusions have a disparate impact based on race and national origin.”⁴³ Accordingly, the EEOC concluded, “national data provides a basis for the Commission to investigate Title VII disparate impact charges challenging criminal record exclusions.”⁴⁴ Yet courts continue to reject disparate impact claims predicated on national statistics.⁴⁵ The EEOC should clarify its guidance on background checks and pursue employers whose background checks yield disparate impacts.

The Commission can also do more to induce employers to identify and adopt less discriminatory alternatives in hiring practices. As a general matter, most employers face too little pressure from regulators to meaningfully evaluate their selection methods and consider less discriminatory alternatives.⁴⁶ Employers also are not incentivized to compare alternative selection procedures to find the least discriminatory means of accomplishing their hiring goals. Under Title VII, hiring practices that have a disparate impact—even those justified as “job related”—are unlawful if the employer could have used a less discriminatory alternative.⁴⁷ However, in practice, this standard manifests as a burden on the plaintiff in litigation.⁴⁸ Sustained attention from the EEOC on hiring practices would motivate many employers to adopt better, less discriminatory approaches.

In sum, we are pleased to read the EEOC’s Draft Strategic Enforcement Plan for 2023–2027 and look forward to assisting the Commission in its efforts to eradicate

⁴³ Equal Employment Opportunity Commission, Enforcement Guidance on the Consideration of Arrest & Conviction Records in Employment Decisions Under Title VII of the Civil Rights Act, Apr. 25, 2012, <https://www.eeoc.gov/laws/guidance/enforcement-guidance-consideration-arrest-and-conviction-records-employment-decisions>.

⁴⁴ *Id.*

⁴⁵ See, e.g., *Mandala v. NTT Data, Inc.*, 975 F.3d 202 (2d Cir. 2020) (affirming dismissal of lawsuit where plaintiffs relied on national statistics, rather than on the particular applicant pool in question, to show that excluding applicants with criminal records had a disparate impact).

⁴⁶ Essential Work, *supra* note 24, at 39 & nn. 150–151.

⁴⁷ *Id.* at 39 & n.152.

⁴⁸ *Id.* at 39–40.



discrimination in employment. We welcome further conversations on these important issues. If you have any questions, please contact Mitra Ebadolahi (Senior Project Director, mitra@upturn.org), Natasha Duarte (Project Director, natasha@upturn.org), and Urmila Janardan (Policy Analyst, urmila@upturn.org).