Guidance on Application of the Fair Housing Act to the Advertising of Housing, Credit, and Other Real Estate-Related Transactions through Digital Platforms

I.   Introduction

This guidance from HUD's Office of Fair Housing and Equal Opportunity explains how the Fair Housing Act ("Act") applies to the advertising of housing, credit, and other real estate-related transactions through digital platforms. In particular, it addresses the increasingly common use of automated systems, such as algorithmic processes and Artificial Intelligence ("AI"),[1] to facilitate advertisement targeting and delivery.

New technologies can be used to target advertising toward some consumers and away from others.[2] This can be done deliberately—for example, when advertisers choose to have their ads directed in a particular way—but it can also can occur through the operation of complex automated systems designed to make ad delivery more efficient in accomplishing an advertiser or ad platform's purposes.[3] These systems may conclude, for example, that women are more likely than men to click on advertising for certain products, and so direct such ads only to women (or, more precisely, to people they have estimated to be women). Or they may conclude that Black people respond more frequently to certain ad variants than others, and so direct only those ads to Black people.[4] Importantly, this can happen without the advertiser's direction or knowledge, and can even frustrate an advertiser's intention that an ad be distributed more broadly.[5]

Such targeting and delivery, which may be permissible in other contexts, risks violating the Act when used for housing-related ads. As described further below, the Act prohibits discrimination in a variety of housing-related transactions based on seven protected classes— race, color, religion, sex (including sexual orientation and gender identity), national origin, familial status (children under eighteen being present, seeking of legal custody, or pregnancy),

---

[1] *See* Exec. Order No. 14,110, Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, 88 Fed. Reg. 75,191 (Oct. 30, 2023). The EO sets forth the definition of "artificial intelligence" or "AI," pursuant to 15 U.S.C. § 9401(3), as "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to (A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action."

[2] *See* Thomas Beauvisage, Jean-Samuel Beuscart, Samuel Coavoux & Kevin Mellet, *How Online Advertising Targets Consumers: The Uses of Categories and Algorithmic Tools by Audience Planners*, New Media & Soc'y 1, 1–3 (2023) [hereinafter, Beauvisage et al., *How Online Advertising Targets Consumers*].

[3] *See* Levi Kaplan, Nicole Gerzon, Alan Mislove & Piotr Sapienzynski, *Measurement and Analysis of Implied Identity in Ad Delivery Optimization*, Proc. of 22nd ACM Int. Measurement Conf. 195, 195–96 (2022) [hereinafter, Kaplan et al., *Measurement and Analysis*].

[4] *Id.* at 205.

[5] *Id.* at 207.

or disability.[6] Ad targeting risks discriminating on the basis of protected characteristics in violation of the Act in multiple ways, including: denying consumers information about housing opportunities; targeting vulnerable consumers for predatory products or services; discouraging or detering potential consumers; advertising different prices or conditions to consumers; steering home-seekers to particular neighborhoods; or charging advertisers higher amounts to show ads to some consumers.[7] Discriminatory advertising can contribute to, reinforce, and perpetuate residential segregation and other harms addressed by the Fair Housing Act.

This guidance describes the responsibilities and potential liability of both advertisers and ad platforms, and how ad targeting and delivery functions may risk violating the Act when deployed for housing-related ads. This guidance concludes with recommendations to avoid violations of the Act.

The term "advertiser," as used in this guidance, refers to entities or individuals placing advertisements for any one of the full range of products and services covered by the Act, which includes rental housing and property management services, residential real estate and related services, mortgages and mortgage broker services, home insurance, and more.[8] The term "ad platform," as used in this guidance, refers to products or systems used to direct and deliver advertisements to consumers in digital spaces—sometimes on a single website or mobile application, such as a social media or real estate website, and other times across many websites, mobile applications, or other channels. Ad platforms can be highly complex systems that may employ a wide range of tools to determine which ads are delivered to which users. These tools often include audience categorization and selection tools, custom and mirror audience tools, and ad delivery models, ad auctions, and other algorithmic systems.[9]

II.   Legal Background

The Fair Housing Act prohibits discrimination in the sale, rental, and financing of dwellings and in other housing-related services because of race, color, religion, sex, national origin, familial status, or disability.[10] The Act prohibits intentionally discriminatory practices, as well as those with an unjustified discriminatory effect.[11] The Act also prohibits discriminatory statements.[12] Any entity that "play[s] a substantial role" in a discriminatory housing decision or

---

[6] 42 U.S.C. §§ 3601–19.

[7] *See, e.g.*, Havens Realty Corp. v. Coleman, 455 U.S. 363, 368 (1982) (steering); Gladstone Realtors v. Village of Bellwood, 441 U.S. 91, 93–96 (1979) (steering); Clark v. Universal Builders, Inc., 501 F.2d 324, 333–334 (7th Cir. 1974) (pricing); Horne v. Harbour Portfolio VI, LP, 304 F. Supp. 3d 1332, 1341–42 (N.D. Ga. 2018) (predatory targeting). *See generally* Varun Nagaraj Rao & Aleksandra Korolova, *Discrimination Through Image Selection by Job Advertisers on Facebook*, Proc. of 23rd ACM Conf. on Fairness, Accountability, & Transparency 1772, 1772 (2023); *see also, e.g.*, *Factsheet: Surveillance Advertising: What About Discrimination,* Consumer Fed. of Am. (Aug. 26, 2021), https://consumerfed.org/consumer_info/factsheet-surveillance-advertising-discrimination/.

[8] *See, e.g.,* Swanson v. Citibank, N.A., 614 F.3d 400, 406 (7th Cir. 2010) (applying the Act to lenders); Fair Hous. Council of San Fernando Valley v. Roommates.com, LLC, 521 F.3d 1157, 1166–1175 (9th Cir. 2008) (en banc) (housing search website); NAACP v. Am. Fam. Mut. Ins. Co., 978 F.2d 287, 290 (7th Cir. 1992) (insurers); United States v. Am. Inst. of Real Estate Appraisers, 442 F. Supp. 1072, 1079 (N.D. Ill. 1977) (appraisers).

[9] *See generally* Beauvisage et al., *How Online Advertising Targets Consumers*.

[10] 42 U.S.C. §§ 3601–19.

[11] See 24 C.F.R. § 100.500; Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Cmtys. Project, Inc., 576 U.S. 519, 540 (2015).

[12] 42 U.S.C. §§ 3604(c).

outcome can be liable under the Act, even if that entity or person is not the provider of housing or housing-related services.[13]

Intentional discrimination can include using a protected characteristic—or a proxy for a protected characteristic[14] —as the basis for the decision to offer, not offer, or provide different terms for housing or a housing-related transaction. This is true even if that decision is made in whole or in part by an automated system, including a system using machine learning or another form of AI.

Even when there is no intent to discriminate, a policy or practice violates the Act if it has a discriminatory effect and (i) the policy or practice is not necessary to achieve a substantial, legitimate, non-discriminatory interest or (ii) the interest can be served by a less discriminatory alternative.[15] Discriminatory effects liability is assessed using a three-step burden-shifting framework and requires a fact-specific analysis. In the first step of the analysis, a plaintiff (or HUD in an administrative enforcement action) has the burden to prove that a policy or practice has a discriminatory effect, meaning that it has a disparate impact based on one or more of the characteristics protected by the Act or it perpetuates segregation.[16] If such a discriminatory effect is proven, the burden shifts to the defendant (or respondent in an administrative action) to prove that the policy or practice is necessary to achieve a substantial, legitimate, nondiscriminatory interest.[17] If a defendant or respondent successfully meets this burden, the burden shifts back to the plaintiff or HUD to prove that such interest could be served by another practice that has a less discriminatory effect.[18]

The Act also explicitly forbids discriminatory statements and advertising practices. The Act specifically prohibits making, printing, or publishing any housing-related advertisement which indicates any preference, limitation or discrimination because of a protected characteristic.[19] HUD's regulations implementing the Act prohibit "[u]sing words, phrases, photographs, illustrations, symbols or forms which convey that dwellings are available or not available to a particular group of persons because of [protected characteristics]"[20]; "[s]electing media or locations for [housing-related advertising] which deny particular segments of the

---

[13] Commonwealth v. Priority Home Sols., LLC, CL20000688-00, at 6 (Va. Cir. Ct. Dec. 6, 2021); Sabal Palm Condos. of Pine Ridge Ass'n, Inc. v. Fischer, 6 F. Supp. 3d 1272, 1293 (S.D. Fla. 2014).

[14] A proxy is a variable that is highly correlated with a protected class, as for example ZIP Code of residence can be highly correlated with race.

[15] 24 C.F.R. § 100.500; *see also Inclusive Cmtys. Project, Inc.*, 576 U.S. at 541 (2015) (citing HUD's 2013 rule). HUD's 2023 Discriminatory Effects Rule formally restores, with minor unrelated modifications, HUD's 2013 Discriminatory Effects Rule. An intervening version of the Rule published in 2020 was enjoined before it took effect, so the framework established by the 2013 Rule has been in effect without interruption since it was issued. *See* Mass. Fair Hous. Ctr. v. U.S. Dep't of Hous. & Urb. Dev., 496 F. Supp. 3d 600 (D. Mass. 2020).

[16] 24 C.F.R. § 100.500(a), (c)(1).

[17] *See* 24 C.F.R. § 100.500(c)(2); *see also Inclusive Cmtys. Project*, 576 U.S. at 541.

[18] 24 C.F.R. § 100.500(c)(3); *accord Inclusive Cmtys. Project*, 576 U.S. at 527.

[19] 42 U.S.C. § 3604(c); 24 C.F.R. § 100.75(a). Courts consistently interpret this to mean that a defendant can violate 42 U.S.C. § 3604(c) if the notice, statement, or advertisement indicates discrimination to an "ordinary reader" or "ordinary listener," regardless of whether the defendant intended to discriminate. *See, e.g.,* Rodriguez v. Vill. Green Realty, Inc., 788 F.3d 31, 41, 52–54 (2d Cir. 2015); Corey v. Sec'y, Dep't of Hous. & Urb. Dev. *ex rel.* Walker, 719 F.3d 322, 326 (4th Cir. 2013); White v. Dep't of Hous. & Urb. Dev., 475 F.3d 898, 905–06 (7th Cir. 2007).

[20] 24 C.F.R. § 100.75(c)(1).

housing market information about housing opportunities because of [protected characteristics]"[21]; and "[r]efusing to publish [housing-related advertising] or requiring different charges or terms for such advertising because of [protected characteristics]."[22]

III. Audience Targeting Tools

Ad platforms have provided a range of tools for advertisers to select their intended target audience for ads, including audience categorization, custom audience, and mirror audience tools.

A. Audience Categorization Tools

"Audience categorization tools" refers to tools that ad platforms offer advertisers to segment and select potential audiences by category, such as gender, age, income, location, interests, activities, or connections. These audience categorization tools may take different forms, such as drop-down menus, toggle buttons, search boxes, or maps, and may allow both inclusion and exclusion functions. An example of how categorization tools may work is the display of a toggle button in an ad placement interface that prompts an advertiser to select "men" or "women" as the potential audience for ad delivery. Other examples of categorization tools would be the display of a dropdown menu that prompts advertisers to select potential audience members by field of employment, or a map that prompts advertisers to select residents of certain neighborhoods as potential audience members.

In some instances, consumers may self-identify and disclose their gender, location, or other characteristics when they sign up for a product, make a purchase, or even sign into their browser. For example, consumers may disclose their gender in response to question prompts when creating and filling out a profile on a social media site, or consumers may disclose their address when signing up for a company's mailing list. In other instances, ad platforms infer consumers' characteristics from available data such as their purchase or browsing history, activities, and movements. Often such inferences are drawn not just from information about the consumer, but also from information about people with whom the consumer interacts. Examples of this would be a platform inferring a consumer is female based in part on past purchases of women's clothing; inferring a consumer resides on a particular block based in part on their phone's regular nighttime presence in that location; or inferring a consumer's sexual orientation based in part on friend groups, associations, or the content of social media posts.

Advertisers and ad platforms may violate the Act by segmenting and selecting audiences for housing-related ads based in part on protected characteristics or proxies, whether self-disclosed or inferred. For example, some categorization tools have directly

---

[21] 24 C.F.R. § 100.75(c)(3).
[22] 24 C.F.R. § 100.75(c)(4).

segmented potential audiences by protected characteristics, such as gender, parental status, country of origin, religion, or affiliation with disability-rights groups.[23] Other categorization tools have segmented audiences based on proxies for protected characteristics.[24] Examples of such proxies are precise geographic location (*e.g.*, census blocks), language spoken, purchases of diapers, cribs, or other child-needs, or engagement with culturally-specific media.[25]

One way that ad platforms' and advertisers' use of these tools for housing-related ads may violate the Act is by limiting protected class groups from accessing information about housing opportunities. For example, they may do so when used to exclude families with children or people with service animals from the eligible audience for a rental ad, or when used to exclude residents of predominately Black and Hispanic neighborhoods from the eligible audience for a home insurance ad. Such exclusion risks violating the Act even if done for purportedly 'benign' purposes—such as an advertiser's belief that a particular property is inappropriate for children because it has potentially hazardous features like stairs or pools.[26]

The use of categorization tools to target housing-related ads may also violate the Act by subjecting vulnerable protected class groups to targeting of predatory products or harmful housing practices, also known as 'reverse redlining.' For example, a lender may violate the Act by targeting ads for high-cost loans only to consumers who have limited English-proficiency, or a company running a housing-related scam may violate the act by targeting ads to immigrants from a particular country (and using cultural markers to build trust to perpetrate the scam).

Finally, the use of categorization tools to target housing-related ads may violate the Act when used to show different content to different groups on the basis of protected characteristics—which may result in steering, pricing discrimination, or other discriminatory outcomes.[27] Advertisers may frame an ad in one way to one group, and a different way to another group, based on their own views on the efficacy of doing so or relying on tools intended to predict which ads perform better with which groups (such as A/B testing or machine learning). A well-publicized example of this from a non-housing context has been movie trailers tailored differently for White or Black audiences.[28] In the housing context, this customization may violate the Fair Housing Act if, for example, it results in meaningfully different information about housing opportunities being made available to consumers on the

---

[23] *See e.g.*, Complaint ¶¶ 2–5, United States v. Meta Platforms, Inc., No. 1:22-cv-05187 (S.D.N.Y. filed June 21, 2022), https://www.justice.gov/crt/case-document/file/1514111/dl; *see also* Beauvisage et al., *How Online Advertising Targets Consumers* at 7–9.

[24] *See e.g.*, Complaint ¶¶ 32, 34, 49–58, *Meta Platforms, Inc.*; Charles Duhigg, *How Companies Learn Your Secrets*, N.Y. TIMES MAG. (Feb. 16, 2012), https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html.

[25] In complex models, combinations of a number of factors, each of which may be weighted differently, can also function as proxies for protected characteristics.

[26] *See e.g.*, Community Svcs., Inc. v. Wind Gap Mun. Auth., 421 F.3d 170, 177 (3d Cir. 2005); United States v. Reece, 457 F. Supp. 43, 48 (D. Mont. 1978).

[27] *See, e.g.*, *Clark*, 501 F.2d 324, 333-334 (7th Cir. 1974) (pricing); *Havens Realty Corp.*, 455 U.S. 363 (1982) (steering); *Gladstone Realtors*, 441 U.S. 91 (1979) (steering).

[28] *See e.g.*, Chi Chi Inzundu, *Straight Outta Compton had 'different trailers for different races'*, BBC (Mar. 18, 2016), https://www.bbc.com/news/newsbeat-35841138.

basis of a protected characteristic, serves to deter consumers from seeking housing opportunities on the basis of a protected characteristic, [29] or serves to steer consumers to housing opportunities in specific neighborhoods on the basis of a protected characteristic.[30]

For example, it may violate the Act to for a property management company to target a rental ad to predominately Black neighborhoods that includes prominent disclaimers stating "no criminal records!" and "good credit only!," while targeting an otherwise similar ad for the same property to White neighborhoods that does not contain those disclaimers. Such an ad could deter or discourage applicants on the basis of race. It may also violate the Act to target a mortgage ad to men that includes details about available interest rates, while an otherwise similar ad targeted to women instead contains the phrase "getting a mortgage is easy!." Recipients of the second ad would be given less information to make a home purchase decision because of their gender.

Advertisers and ad platforms should be alert to the risk of using categorization tools for housing-related ads and take steps to avoid discriminatory delivery through these features. Advertisers should not utilize categorization tools for housing-related ads that segment audiences on the basis of protected characteristics or close proxies, and ad platforms should consider disabling some or all categorization functions for housing-related ads. Notably, some platforms have already taken steps to limit audience categorization tools for housing-related ads to avoid enabling ad placements that may violate the Act.[31] These efforts often rely on a combination of machine learning and advertiser self-certification to flag ads as pertaining to content covered by the Act before audience selection options are offered to an advertiser placing an ad. Advertisers should be alert to this flagging function and accurately certify as to whether an ad being posted is a housing-related ad and so covered by the Act. Ad platforms should regularly audit the accuracy of advertiser certification and machine learning-based flagging tools, and take corrective action when necessary to ensure housing-related ads are being identified.

Ad platforms that offer categorization tools for housing-related ads should audit the categorization tools offered to advertisers and identify any resulting differences in ad delivery based on protected class to determine which categorization tools should not be offered for housing ads. Performing such audits can ensure that the ad platform provides the least discriminatory option among those that satisfy a legitimate need. It should be kept in mind that ad platforms may regularly generate new targeting categories, in part based on machine learning techniques. Newly derived targeting options may operate effectively as proxies for

---

[29] Hous. Rts. Ctr. v. Sterling, 404 F. Supp. 2d 1179, 1190–92 (C.D. Cal. 2004) (discussing some of the mechanisms by which housing providers may unlawfully deter consumers).

[30] *Havens Realty Corp.*, 455 U.S. 363 (1982) (steering); *Gladstone Realtors*, 441 U.S. 91 (1979) (steering).

[31] *See e.g.*, *An Update on Our Ads Fairness Efforts*, FACEBOOK.COM, https://about.fb.com/news/2023/01/an-update-on-our-ads-fairness-efforts/ (last visited April 26, 2024); *Update to Personalized advertising policies: Housing, employment, and credit (October 2020)*, GOOGLE, https://support.google.com/adspolicy/answer/9917652?hl=en (last visited April 26, 2024).

April 29, 2024

protected classes, [32] so ad platforms should regularly test that any targeting options provided to advertisers for housing-related ads do not result in discriminatory delivery. [33]

While these practices can help minimize one form of risk, both advertisers and ad platforms should be aware that limiting explicit categorization options does not protect against potential discrimination in downstream delivery functions like those described in section IV below.

B. Custom and Mirror Audience Tools

A number of platforms have also offered features that deliver ads only to a specified "custom" audience or to an audience estimated to be similar to a custom audience in interests, behaviors, or likelihood to interact with an ad.

Custom audience tools deliver ads to an audience specified by an advertiser. These tools may function by prompting an advertiser to upload a list of identifying information such as phone numbers, emails, or names, and then delivering ads only to members of that audience. Advertisers may acquire these lists from existing customer databases, data brokers, or other sources. Custom audience tools may also function by identifying consumers who have taken a specific action tracked by an advertiser or ad platform, such as visiting a particular website, making a particular purchase, attending a particular event, or interacting with a particular person or organization.

Mirror audience tools are designed to find consumers who are similar to or mirror consumers on a customized list—also called a 'source audience.'[34] After prompting advertisers to upload data that identifies specific consumers, the mirroring tools use algorithmic techniques to find an expanded audience of potential consumers who are similar to the source audience, which may include consideration of demographics, income, location, behaviors, interests, habits, associations, or other traits—potentially leading to discriminatory outcomes. [35]

The use of custom and mirror audience tools for housing-related ads may violate the Act when the source audience is limited by protected characteristics, and when mirroring functions to introduce, replicate, or enhance such limitations. Custom and mirror audiences

---

[32] *See* Till Speicher, Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Ribeiro, George Arvanitakis, Fabricio Benevenuto, Krishna P. Gummadi, Patrick Loiseau, Alan Mislove, *Potential for Discrimination in Online Targeted Advertising,* 81 Proc. of Machine Learning Res. 1, 9–10 (2018) (highlighting how a malicious advertiser could use facially neutral online engagement activity data to exclude gay users from an advertisement); Valerie Schneider, *Locked Out by Big Data: How Big Data, Algorithms and Machine Learning May Undermine Housing Justice*, 52 Colum. Hum. Rts. L. Rev. 251, 267 (2020) (explaining how algorithms easily find proxies when they are not allowed to consider protected class status); *see also* Alice Xiang, *Reconciling Legal and Technical Approaches to Algorithmic Bias*, 88 Tenn. L. Rev. 649, 667 (explaining that, even if a system is trained to ignore protected classes as a variable, the system can recreate those variables).

[33] Valerie Schneider, *Locked Out by Big Data: How Big Data, Algorithms and Machine Learning May Undermine Housing Justice*, 52 Colum. Hum. Rts. L. Rev. 251, 299 (2020).

[34] *See* Piotr Spaiezynski, Avijit Ghosh, Levi Kaplan, Aaron Rieke & Alan Mislove, *Algorithms that "Don't See Color": Measuring Biases in Lookalike and Special Ad Audiences*, Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society 609 (2022), https://doi.org/10.1145/3514094.3534135.

[35] *Id.* at 610.

may be used to effectuate discriminatory intent, such as if a perpetrator of a housing scam uploads a list of consumers known to have limited English proficiency in order to target an ad for the scam, and the ad platform mirrors that list. But custom and mirror audience tools may also drive discriminatory delivery of housing-related ads even in the absence of discriminatory intent. For example, when placing an ad for a home, a real-estate agent may upload a custom audience from attendees at a recent open house held for a nearby home. If all the attendees at the open house were White, the new home ad will be eligible to be delivered only to White consumers. Using mirroring functions with this audience may replicate and expand biases in the source audience, resulting in delivery of the ad to people who did not attend the open house but who may also all be White (because of similarities identified by the algorithm based on the millions of data points described above covering interest, activities, and associations).

As another example, in order to attract new tenants, a property management company may post an ad and provide a source list of existing tenants, who are mostly White and childless, to an ad platform. The ad platform's mirroring function may generate an expanded audience of other White and childless people based on their shared online behavior. This would deny information about this housing opportunity based on the protected class characteristics of the source list, even though there would be high demand from other members of the community if they knew about the housing opportunity.

Advertisers and ad platforms should be alert to the risk posed by utilization of custom and mirror audience tools for housing-related ads and take steps to avoid discriminatory delivery through these features. Ad platforms should consider disabling custom and mirror audience functions altogether for housing-related ads.[36] When offering those features, ad platforms should present advertisers with clear guidance and prominent disclaimers about the appropriate use of custom and mirror audience tools for housing-related ads. Advertisers and ad platforms should audit source lists to ensure they are not unjustifiably limited on the basis of protected characteristics. Advertisers and ad platforms should also audit ad delivery outcomes to ensure they are non-discriminatory.

Again, following these best practices can be helpful but, as with categorization tools, limiting discrimination in source audiences and mirroring functions does not protect against potential discrimination in the downstream ad delivery functions described below.

IV.   Algorithmic Delivery Functions

In addition to eligible audience selection tools—such as the categorization, custom, and mirror audience tools described above—ad platforms may use machine learning and other forms of AI to decide which ads are actually delivered to which consumers, and at what location, time, and price. As there is limited space in which to place ads for any given

---

[36] *See e.g.*, Settlement Agreement at 6, United States v. Meta Platforms, Inc., No. 1:22-cv-05187 (S.D.N.Y. filed June 27, 2022).

consumer, ad platforms employ sophisticated technologies to conduct complex calculations to determine which ads, from among the ads the consumer is eligible to receive, to actually deliver to the consumer.

Ad platforms may use machine learning and other forms of AI to determine which consumers within the targeted audience are most likely to achieve the advertiser's objective for a specific ad, such as clicking on it. This estimate may be combined with other information—such as how much particular advertisers are willing to pay to place an ad and an estimate of the quality of the ad and other factors—to determine which ad to show a particular consumer at a given moment.[37] Algorithmic delivery functions may violate the Act when they direct housing-related ads to or away from consumers based on protected characteristics—potentially resulting in steering, pricing discrimination, or other discriminatory outcomes.[38] As with audience selection functions, algorithmic delivery functions may operate to exclude protected groups from an ad's audience or to concentrate delivery to a protected group—an outcome particularly problematic for predatory products.

Consider the following illustration of how a delivery system might steer ads away from protected class groups in a manner that may run afoul of the Act, even where an advertiser has no intent to create that result. A mortgage lender, conscious of the need to avoid redlining, may specifically want to advertise evenly across a metropolitan area. The mortgage lender may therefore select consumers who live in an entire metropolitan area as the target audience for its ad for a residential mortgage product. However, an ad delivery system may not deliver the ad to every consumer in the selected area. Instead, relying on the mechanisms described above, the system may deliver the ad to a subset of consumers it determines most likely to engage with the ad. Thus, notwithstanding the mortgage lender's expressed intent to display the ad to consumers throughout the metropolitan area, the ad delivery system may actually deliver the ad only to consumers who live in wealthier and Whiter neighborhoods of the metropolitan area, because the system has predicted those consumers are most likely to engage with the ad based on factors such as their income or profession, past interactions with lending ads, or home purchase history of friends, family, and neighbors.

Similarly, if the same mortgage lender seeks to run several ads throughout the metropolitan area with images of various and diverse human models—*e.g.*, an ad with a White family, an ad with a Hispanic family, an ad with a single woman, and ad with a person in a wheelchair—the delivery system may steer particular versions of the ad to particular consumers, in part, based on protected characteristics. For example, research has shown ad

---

[37] These estimates are often utilized in an instant "auction," where the ability to present an ad to the consumer is up for auction. As a result of each auction, the ad with the highest estimated value is delivered to the consumer, with a price determined by the outcome of the auction and the advertiser's specified objective. On large platforms, there may be millions of ad auctions running at any moment.

[38] *See, e.g.*, *Clark*, 501 F.2d at 333-334 (7th Cir. 1974) (pricing); *Havens Realty Corp.*, 455 U.S. 363 (1982) (steering); *Gladstone Realtors*, 441 U.S. 91 (1979) (steering).

delivery systems may deliver an ad with an image of a White family more often to White consumers and an ad with an image of a Black family more often to Black consumers in part because of people's tendency to click on them.[39] Research has also indicated ad delivery systems function likewise along lines of other protected classes such as gender.[40] Additional research has shown that the text content, the content of the destination link of the ad, and image content (beyond images of people) may similarly influence the outcomes. [41] Advertisers and ad platforms should be alert to the potential operation and consequences of this steering for housing-related ads.

Relatedly, systems that incorporate pricing may discriminate by charging advertisers higher amounts to show ads on the basis of protected characteristics. Research shows, for example, it often costs more to advertise to women than men (because of prior purchasing or ad engagement histories).[42] Ad platforms may react to such higher prices by showing the ad only to consumers the platform is charging less to deliver ads to. For example, advertisers who want to pay less per ad view may end up inadvertently advertising only to men. Differential pricing may have additional implications, like effectively charging advertisers more to show ads to consumers on the basis of a protected characteristic.[43]

Finally, discriminatory delivery on the basis of protected characteristics may occur because of differences in the level of confidence a delivery system has about its predictions for consumers—which in turn can be based on disparities in the number and type of past interactions, as well as disparities in data the system was trained on.[44] If an algorithm has greater confidence in its predictions about White consumers relative to housing ads, that may cause it to deliver certain ads much more frequently to White consumers.[45]

Advertisers and ad platforms should be alert to the risk of utilization of ad delivery technologies and consider the recommendations below to reduce the risk of violating the Act.

---

[39] Kaplan et al., *Measurement and Analysis* at 196.

[40] *See generally id.* at 195.

[41] *See e.g.,* Muhammad Ali, Piotr Sapienzynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, Aaron Rieke, *Discrimination through Optimization: How Facebook's Ad Delivery Can Lead to Bias*, 3 PROC. OF ACM ON HUMAN-COMPUTER INTERACTION 199 (2019), https://www.khoury.northeastern.edu/~amislove/publications/FacebookDelivery-CSCW.pdf.

[42] *Men are Cheap: Efficient Gender Targeting with Facebook Ads*, RESOLUTION MEDIA & KENSHOO SOC., (2012), https://skai.io/wp-content/uploads/2013/02/Social-Media-Insights-from-Resolution-Media-and-Kenshoo-Social-Part-3-Men-Are-Cheap.pdf; Dina Fine Maron, *Science Career Ads Are Disproportionately Seen by Men*, SCI. AM. (July 25, 2018), https://www.scientificamerican.com/article/science-career-ads-are-disproportionately-seen-by-men/.

[43] *See e.g.*, Anja Lambrecht & Catherine Tucker, *Algorithm-Based Advertising: Unintended Effects and the Tricky Business of Mitigating Adverse Outcomes*¸13 NIM MKT. INTELLIGENCE REV. 24, 27 (2021), https://www.nim.org/fileadmin/PUBLIC/12_NIM_MIR_Issues/MIR_Die_dunklen_Seiten_des_digitalen_Marketings/MIR_Die_dunklen_Seiten_des_digitalen_Marketings_EN/lambrecht_tucker_vol_13_no_1_eng.pdf.

[44] Anja Lambrecht, Catherine Tucker, Apparent Algorithmic Discrimination and Real-Time Algorithmic Learning in Digital Search Advertising 1, 26-28 (Mar. 29, 2024) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3570076.

[45] *See generally id.*; *see also* Catherine Tucker, *Algorithmic Exclusion: the Fragility of Algorithms to Sparse and Missing Data*, (Brookings Ctr. for Regs. & Mkts., Working Paper), https://www.brookings.edu/wp-content/uploads/2023/02/Algorithmic-exclusion-FINAL.pdf.

April 29, 2024

V. Conclusion and Recommendations

The ad targeting and delivery functions described above may violate the Act when they unlawfully deny consumers information about housing opportunities based on the consumers' protected characteristics. They may also violate the Act when used to target vulnerable consumers for predatory products or services on the basis of a protected characeristic, display content that could discourage or deter potential consumers on the basis of a proteced characteristic, steer home-seekers to particular neighborhoods, or offer different terms and conditions on the basis of a protected characteristic—among other discriminatory outcomes Advertisers and platforms should be alert to the risks of deploying targeted advertising tools for ads covered by the Act.

Advertisers and platforms should consider the following potential recommendations to reduce the risk of violating the Act.

Advertisers should:

- Utilize ad platforms that are taking steps to manage the risk of discriminatory delivery of housing-related ads through audience selection tools and algorithmic functions. Before using an ad platform, advertisers should ensure that they obtain necessary information and disclosures from the ad platform regarding how the platform mitigates these risks, such as through the steps below.

- Follow ad platform instructions to ensure that advertisements related to housing are identified as such to the ad platform, enabling the appropriate treatment.

- Carefully consider the source, and analyze the composition, of audience datasets used for custom and mirror audience tools for housing-related ads to mitigate risk of generating discriminatory target audiences, and make considered use of any tools provided by the ad platform for evaluating the projected demographics of a targeted audience.

- Monitor outcomes of advertising campaigns for housing-related ads, to the extent possible, to identify and mitigate discriminatory outcomes.

Ad platforms should:

- Ensure that housing-related ads are run in a separate process and specialized interface designed to avoid discrimination in audience selection and ad delivery.

- Avoid providing targeting options for housing-related advertisements that directly describe or relate to FHA-protected characteristics, or that are effectively proxies for FHA-protected characteristics, either alone or in combination.

April 29, 2024

- Conduct regular end-to-end testing of advertising systems to ensure that any discriminatory outcomes are detected, such as by running pairs of ads for equivalent housing opportunities at the same time and comparing the demographics of the delivery audience.

- Proactively identify and adopt less discriminatory alternatives for AI models and algorithmic systems, including by assessing data used to train AI models and verifying that the technologies measure lawful attributes that predict valid outcomes.

- Ensure that algorithms are similarly predictive across protected class groups and make adjustments to correct for any disparities in predictiveness or direct the algorithm to develop additional information that will enhance predictiveness for certain groups.

- Ensure that ad delivery systems are not resulting in differential charges on the basis of protected characteristic, or charging more to advertisers to deliver ads to a non-discriminatory audience.

- Document, retain, or publicly release in-depth information about ad targeting functions and internal auditing.