October 18, 2019

Via Electronic Submission to: www.regulations.gov

Dr. Ben Carson, Secretary
Department of Housing and Urban Development
451 7th Street SW, Room 10276
Washington, DC 20410

Re: HUD's Implementation of the Fair Housing Act's Disparate Impact Standard [Docket No. FR-6111-P-02] [RIN 2529-AA98]

Comments of the Consumer Advocacy & Protection Society (CAPS)
University of California, Berkeley, School of Law

Dear Secretary Carson:

The Consumer Advocacy and Protection Society (CAPS),[1] a student-run organization dedicated to the promotion of consumer law and consumer protection at the University of California, Berkeley, School of Law, appreciates the opportunity to comment in response to the Department of Housing and Urban Development's (HUD) proposal to amend its interpretation of the Fair Housing Act's (FHA) disparate impact standard. Many of us have worked as clinical

---

[1] Consumer Advoc. & Protection Soc'y (CAPS), https://consumer.berkeley.edu/ (last visited Oct. 11, 2019).

students or volunteers with the East Bay Community Law Center's Tenants' Rights Workshop and/or Consumer Justice Clinic, clinics of Berkeley Law that provide legal services to low-income individuals. Some of us have also worked on fair housing enforcement, including for the Department of Justice and at other levels of government. Based on these experiences and our own research, we urge HUD to rescind its proposed amendments to its 2013 "Implementation of the Fair Housing Act's Discriminatory Effects Standard." Specifically, we urge HUD to reconsider the three proposed affirmative defenses available to defendants when the allegedly discriminatory policy or practice relies on an algorithmic model. We must also stress the critical importance that housing providers collect data related to protected classes so that discrimination can be monitored and addressed, and plaintiffs can access the data necessary to establish prima facie cases of discriminatory impact where they are warranted.

Introduction

We have grave concerns about the Proposed Rule in its entirety. The Proposed Rule is antithetical to *Texas Department of Housing & Community Affairs v. Inclusive Communities Project, Inc*., which reaffirmed the longstanding validity of disparate impact liability.[2] This comment focuses on one of the most egregious aspects of the Proposed Rule: the three affirmative defenses that would effectively allow defendants to hide behind algorithmic models in continuing to perpetrate discriminatory housing practices.

---

[2] 135 S. Ct. 2507 (2015); *see also Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971) (establishing the disparate impact standard).

Algorithmic models increasingly determine Americans' housing options and thus warrant much more federal attention.[3] While algorithms can reduce disparate impact, they also can perpetuate and obfuscate discrimination behind a veneer of objectivity. According to one recent study of mortgage lenders between 2012 and 2018, technology-enabled "FinTech" companies utilizing algorithms still charge Latinx and African-American borrowers 5.3 basis points more for purchase mortgages than similarly situated white borrowers.[4] Instead of targeting insidious new methods of discrimination, however, the Proposed Rule would decimate existing protections against disparate impact wherever housing providers, lenders, insurers and a host of other actors ("Providers") rely on algorithms.

The first affirmative defense would allow a defendant to defeat a claim of discriminatory effect by showing that the algorithmic model predicts risk and does not use individual inputs that are close proxies for protected classes. Its narrow emphasis on individual inputs ignores the potential for the interaction between multiple inputs to produce discriminatory effects. This defense also ignores the potential for inputs that are not close proxies for protected characteristics to still impact protected classes differently, independent of predicted risk. The second affirmative defense, which would allow a defendant to defeat a claim of discriminatory impact by showing that the challenged model is produced by a recognized third party, falsely assumes that an industry standard for algorithmic nondiscrimination exists. In the absence of such a standard, this defense disincentivizes Providers from using algorithms that do not

---

[3] Kriston Capps, *How HUD Could Dismantle A Pillar of Civil Rights Law*, CityLab, Aug. 16, 2019, https://www.citylab.com/equity/2019/08/fair-housing-act-hud-disparate-impact-discrimination-lenders/595972/.

[4] Robert P. Bartlett et al., *Consumer Lending Discrimination in the FinTech Era*, UC Berkeley Public Law Research Paper 4 (Dec. 7, 2017), https://ssrn.com/abstract=3063448. (Finding that FinTech lenders discriminate 40% less than traditional lenders, 45% of which also use algorithms).

discriminate. The third affirmative defense, which would allow a defendant to defeat a claim of discriminatory impact by showing that a neutral third party has analyzed the challenged model, poses similar problems as the first two. Moreover, in the absence of a clear standard for what makes an algorithm nondiscriminatory, parties on both sides would be forced to expend extraordinary resources over dueling experts at the pleading stage.

Finally, Providers' growing reliance on algorithms necessitates more collection of data on protected classes in the housing market. But instead of advancing the ability of parties to combat discrimination in the housing market, the Proposed Rule needlessly discourages Providers from collecting this important data.

I. The Proposed Rule's affirmative defenses will allow defendants to defeat legitimate disparate impact claims arising from algorithmic models.

Even where plaintiffs can overcome the Proposed Rule's exceptionally high hurdles to establishing a prima facie case of disparate impact, the Rule allows defendants to easily defeat the claim by wielding three affirmative defenses to algorithm-based discrimination. Under the Proposed Rule, the defendant may overcome the plaintiff's prima facie claim if she:

> (i) Provides the material factors that make up the inputs used in the challenged model and shows that these factors do not rely in any material part on factors that are substitutes or close proxies for protected classes under the Fair Housing Act and that the model is predictive of credit risk or other similar valid objective;
>
> (ii) Shows that the challenged model is produced, maintained, or distributed by a recognized third party that determines industry standards, the inputs and methods within the model are not determined by the defendant, and the defendant is using the model as intended by the third party; or
>
> (iii) Shows that the model has been subjected to critical review and has been validated by an objective and unbiased neutral third party that has analyzed the challenged model and found that the model was empirically derived and is a demonstrably and statistically sound algorithm that accurately predicts risk or other valid objectives, and that none of the factors used in the algorithm rely in any

material part on factors that are substitutes or close proxies for protected classes under the Fair Housing Act.[5]

The Proposed Rule's affirmative defenses undercut the FHA's prohibition on discrimination in housing. The first affirmative defense does not adequately protect against discrimination within algorithmic inputs, nor does it protect more generally against discriminatory algorithmic outputs. The second affirmative defense provides a safe harbor for defendants that use discriminatory third-party algorithms and puts up new roadblocks for plaintiffs seeking to vindicate their rights under the FHA. Similarly, the third affirmative defense increases the burden on plaintiffs at the pleading stage and creates loopholes for defendants through its vague language.

A. Algorithms have the potential to decrease housing discrimination, but HUD must protect against the new vulnerabilities and challenges algorithms create.

The first affirmative defense proposed by HUD would allow defendants to defeat disparate impact claims by showing that individual algorithmic inputs are not "substitutes for a protected characteristic" and that the model is "predictive of risk or other valid objective."[6] The Proposed Rule does not adequately safeguard against discrimination within individual algorithmic inputs, nor does it consider how even nondiscriminatory inputs can interact to result in discriminatory outputs. A 2014 White House report warned that algorithms and big data technology "could be used to 'digitally redline' unwanted groups" through new, and often

---

[5] Implementation of the Fair Housing Act's Disparate Impact Standard, 84 Fed. Reg. 42854 (proposed Aug. 19, 2019) (to be codified at 24 C.F.R. pt. 100.500 (c)(2)).

[6] *Id.*

complex, forms of discrimination.[7] The Proposed Rule would not protect against those vulnerabilities.[8]

Testing each input for fairness does not necessarily guarantee that an algorithmic output is not discriminatory.[9] Algorithms that utilize machine learning models do not rely on explicitly programmed instructions, but rather learn from data and identify patterns without human interaction.[10] Thus, "forbidding [nondiscriminatory] inputs alone does not assure equal pricing" in credit pricing models.[11] By stringing together several variables correlated with race, these algorithms can actually increase pricing disparities between protected groups.[12] The potential for interaction between multiple inputs demonstrates a flaw in the proposed approach of requiring defendants to justify only individual inputs. As a factor-by-factor analysis of a model is not necessarily sufficient to determine the end results, the "mechanics behind" factors in algorithms must also be considered in order to guarantee there is no disparate impact.[13]

Aside from concerns that individual input testing is insufficient, the Proposed Rule does not adequately safeguard against bias within each individual factor. Under the Proposed Rule, algorithmic factors that are predictive of risk or another valid objective are permissible, as long as they are "not substitutes or close proxies for protected classes."[14] But even factors that are not "substitutes or close proxies" for a protected class may still have an impermissible disparate

---

[7] *See Exec. Office Of The President, Big Data: Seizing Opportunity, Preserving Values* (May 2014), available at https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf

[8] *Id*.

[9] *See* Talia B. Gillis & Jann L. Spiess, *Big Data and Discrimination*, 86 U. Chi. L. Rev. 459, 460 (2019).

[10] *See id.*; *Exec. Office Of The President*, *supra* note 7.

[11] Gillis, *supra* note 9.

[12] *Id.* at 469.

[13] *See Exec. Office Of The President, supra* note 7.

[14] Implementation of the Fair Housing Act's Disparate Impact Standard, 84 Fed. Reg. 42854 (proposed Aug. 19, 2019) (to be codified at 24 C.F.R. pt. 100.500 (c)(2)).

impact. According to some scholars, in order to adequately safeguard against discrimination in an algorithm, each factor must (a) be predictive of risk and (b) have no correlation with a protected characteristic independent of predicted risk.[15] This second part of the test is crucial because it eliminates factors that are correlated to a protected characteristic in a manner beyond that which can be attributed to actual risk. As the Proposed Rule is written, defendants would be free to use all but the most obvious proxies for protected characteristics as algorithmic inputs. The complexity of algorithmic decisions requires a more robust rule in order to ensure defendants cannot get away with discrimination based on data technicalities.

B. The third-party affirmative defense creates a safe harbor from disparate impact liability for the entire industry.

Under the Proposed Rule's second affirmative defense, Providers may defeat a disparate impact claim simply by showing that the challenged model is "produced, maintained, or distributed by a recognized third party that determines industry standards" as long as the Provider uses the model as intended.[16] This affirmative defense would apply in most cases because Providers most often use models created by third parties.[17] Thus, the Proposed Rule would shield Providers from disparate impact liability arising from third-party algorithmic models.

The third-party exemption is deeply flawed for three reasons. First, the Proposed Rule would provide a get-out-of-liability free card for "industry standard" algorithms even though the industry standard is no guarantee of nondiscrimination. Second, the Proposed Rule would

---

[15] Bartlett et. al., *supra* note 4, at 22.

[16] Implementation of the Fair Housing Act's Disparate Impact Standard, 84 Fed. Reg. at 42862.

[17] Jamie Williams et. al, *EFF to HUD: Algorithms Are No Excuse for Discrimination*, Elec. Frontier Found. (Sept. 26, 2019), https://www.eff.org/deeplinks/2019/09/dangerous-hud-proposal-would-effectively-insulate-parties-who-use-algorithms.

disincentivize Providers from using algorithms that do not discriminate. Finally, shifting liability to third parties would make it more difficult for protected classes to vindicate their rights under the FHA. Providers must be held accountable when their practices have a disparate impact on a protected class, regardless of whether they choose to rely on third-party algorithms.

1. The Proposed Rule would allow Providers to escape liability for following the industry standard practice, but there are no industry standards for algorithmic fairness.

The Proposed Rule carves out a safe harbor for Providers using algorithms created by a "recognized third party that determines industry standards."[18] There is, however, no industry standard practice for avoiding or correcting discriminatory bias in algorithms.[19] Without such industry best practices in place, using an "industry standard" is no guarantee against discrimination and should not be treated as a get-out-of-liability-free card.

Standard industry practices have historically perpetuated widespread systemic discrimination. Indeed, like redlining, algorithms are powerful tools for masking discrimination under the guise of neutral decision making.[20] Currently, there is no industry standard for correcting algorithmic bias, although researchers are working toward that goal.[21] Some researchers caution that there is no cookie-cutter method for removing systemic bias in

---

[18] Implementation of the Fair Housing Act's Disparate Impact Standard, 84 Fed. Reg. at 42862.

[19] *See* Capps, *supra* note 3; *see also* Andrew D. Selbst, *A New HUD Rule Would Effectively Encourage Discrimination by Algorithm*, Slate (Aug. 19, 2019), https://slate.com/technology/2019/08/hud-disparate-impact-discrimination-algorithm.html.

[20] Daniel Greene et al., *Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial Intelligence and Machine Learning*, Hawaii International Conference on System Sciences (2019), *available at https://pdfs.semanticscholar.org/625d/a63503d70be79cf5a6454686baae4a9256fa.pdf?_ga=2.131984847.1904461500.1571103987-523142922.1571103987*.

[21] *See* Selbst, *supra* note 19.

algorithms.[22] Even the most advanced systems cannot guarantee nondiscrimination when a model trained using data from one city is applied somewhere else, refuting the idea that such an industry standard even could exist.[23] According to one data ethicist working on such technologies, most industry-focused "ethical A.I." efforts remain ineffective because their models are too simplistic.[24]

2. The Proposed Rule would disincentivize Providers from using nondiscriminatory third-party algorithms.

The Proposed Rule would disincentive Providers from seeking third-party algorithms that guarantee nondiscrimination because Providers face no FHA liability arising from those algorithms. Nor do Providers have an incentive to monitor whether the outputs that third-party algorithms produce are discriminatory. Moreover, even if a Provider does become aware that its third-party algorithms are discriminatory, it still has no incentive to change its practices or inform the responsible third party of its findings. In fact, lenders have an incentive to continue utilizing discriminatory algorithms because they can be more profitable.[25]

Under this regime, Providers will exercise no pressure on third parties to invest and offer alternative, nondiscriminatory algorithms. Third-party models that discriminate and are widely adopted will likely become the "industry standard," thus allowing their users to escape liability--regardless of the Provider's knowing complicity. As algorithmic modeling becomes the industry

---

[22] *See id.*

[23] *See id.*

[24] *See id.*

[25] Bartlett et. al., *supra* note 4, at 4 (charging minority borrowers more because they have fewer choices or shop around less is discriminatory).

norm for mortgage lenders[26] and other Providers, it is more important than ever to incentivize Providers to use models that do not have a disparate impact on protected classes.

3. The Proposed Rule would make it harder for plaintiffs to seek a remedy under the Fair Housing Act.

The Proposed Rule's attempt to shift liability from Providers to third parties may make it harder for plaintiffs to challenge algorithmic models.[27] In filing disparate impact claims, plaintiffs must point to a specific policy and how that policy results in a discriminatory effect.[28] However, in a case with a third-party algorithm, the plaintiff may not know that a third party is responsible for the disparate impact, which third party is responsible, or whether the disparate impact is the result of interaction between multiple algorithms. Model makers likely will attempt to rely on trade secret law to resist disclosing any information about how their algorithms are designed or function.[29]

Case law surrounding third-party liability is sparse, and courts will first have to decide whether the FHA covers algorithm makers before reaching the merits of any case.[30] This year, a Connecticut district court held that a third-party tenant screening company could be held liable for a criminal history screening tool if it has a disparate impact on prospective tenants.[31] However, it remains to be seen whether other courts will reach similar conclusions in regards to other algorithmic tools.

---

[26] *Id.* at 2 (noting that 45 percent of the largest mortgage lenders utilize algorithms, and the percentage of total mortgages that are priced using algorithms is even higher).

[27] Implementation of the Fair Housing Act's Disparate Impact Standard, 84 Fed. Reg. 42854, 42859 (proposed Aug. 19, 2019) (to be codified at 24 C.F.R. pt. 100.500 (c)(2)).

[28] *Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Cmtys. Project, Inc*., 135 S. Ct. 2523 (2015).

[29] Williams et al., *supra* note 17.

[30] *Id.*

[31] *Connecticut Fair Hous. Ctr. v. Corelogic Rental Prop. Sols., LLC*, 369 F. Supp. 3d 362, 373 (D. Conn. 2019).

Finally, third parties may counter that any disparate impact was the result of provider misuse.[32] In situations where the Provider and the third party point fingers at one another, the plaintiff may be left in the middle without redress. Plaintiffs must be able to hold Providers and third parties jointly liable so neither can escape responsibility for the harm that they inflict.

C. The Proposed Rule would drastically increase litigation costs at the pleading stage, inappropriately burdening plaintiffs, defendants, and the court.

The Proposed Rule's third affirmative defense, which would allow defendants to defeat a disparate impact claim simply by showing that a "neutral third party" has analyzed and approved of the challenged model,[33] has the same shortcomings as the other two proposed affirmative defenses discussed above. But this third affirmative defense also is misguided because (1) it forces plaintiffs to engage experts just to file their complaint and forces defendants to engage experts just to evaluate a complaint, ratcheting up the cost of bringing a disparate impact suit far beyond what is reasonable and necessary; and (2) in the absence of clear industry standards, it provides courts with inadequate guidance to determine either "neutrality" or "qualified expert."

First, allowing use of a "neutral third party" to immunize a defendant from liability for discrimination burdens the pleading stage with unnecessary costs. The Supreme Court already requires plaintiffs to plead cases with specificity.[34] Many plaintiffs must engage experts to plead prima facie cases of disparate impact, at significant expense.[35] Proving discrimination already can be cost-prohibitive, effectively closing the courthouse door to individuals who have been

---

[32] Implementation of the Fair Housing Act's Disparate Impact Standard, 84 Fed. Reg. at 42859.

[33] *Id.* at 42862.

[34] *Ashcroft v. Iqbal*, 556 U.S. 662 (2009); *Bell Atl. Corp. v. Twombly*, 550 U.S. 544 (2007).

[35] Jon Greenbaum et al., *Toward a More Just Justice System: How Open are the Courts to Social Justice Litigation?*, Lawyers' Committee for Civil Rights Under Law 32–33 (2016), *available at* https://lawyerscommittee.org/wp-content/uploads/2016/08/Toward-A-More-Just-Justice-System.pdf.

harmed by discrimination. Under the Proposed Rule, a defendant could defeat a claim merely by pointing to analysis done by an alleged "neutral third party." To keep her claim alive, a plaintiff likely would need to hire her original expert, or another expert, to undertake additional analysis to show how the third-party analysis is incomplete. This added burden could multiply the cost of advancing a disparate impact claim and turn the pleading stage into a mini-trial, devolving preliminary hearings into resource-intensive battles of the experts. Moreover, any expert trying to challenge a third-party analysis would be stymied prior to discovery. The Proposed Rule would force courts to wade into thorny issues of expert certification and trade secrets at the most preliminary stage of litigation. Requiring more expert witnesses at the pleading stage would balloon costs and difficulties for all involved.

Second, the current definitions and qualifications of a "neutral third party" are inadequate and unworkable, further exacerbating the difficulties of conducting a mini-trial at the pleading stage. The Proposed Rule states only that the model must be "validated by an objective and unbiased neutral third party" for this third affirmative defense to apply.[36] The Rule would give the court no standards for what would qualify a third party as neutral. "Neutrality" can be amorphous and in the eye of the beholder. But widely accepted third-party certification, at this juncture, does not exist. In the absence of clearly established industry standards, even a truly neutral third party will have no baseline to which to compare the challenged model, and neither will the court. Plaintiffs, in response, will be left with little guidance on how to challenge a third party's neutrality.

---

[36] Implementation of the Fair Housing Act's Disparate Impact Standard, 84 Fed. Reg. at 42862.

Algorithm-based housing discrimination, like all housing discrimination, violates the FHA and demands redress. But each of the three affirmative defenses included in the Proposed Rule, if adopted, would allow Providers to evade liability for legitimate claims.

II. HUD should not actively discourage Providers from collecting data regarding protected classes.

The Proposed Rule also clarifies that nothing in the disparate impact rule requires, or even encourages, Providers to collect new data on protected classes. This instruction conflicts with the goals and spirit of the FHA to combat racial disparities in housing. Data is at the heart of disparate impact analysis. Federal and state governments, fair housing entities, and individuals who believe that they are experiencing discrimination all benefit from the availability of data detailing the various protected classes which these Providers serve. Given that there is no clear standard or consensus on what makes a nondiscriminatory algorithm, as discussed above, the Proposed Rule would only further discourage data collection which could otherwise inform policy decisions, enable effective enforcement, and help clarify what a nondiscriminatory algorithm might look like.

First, protected class data is necessary to monitoring and understanding evolving trends in racial discrimination among different groups in the U.S.[37] Congress enacted the FHA in order to diminish and ultimately eradicate discrimination in the sale and rental of homes.[38] This goal directly answers the Constitution's never-ending mission to "establish justice, insure domestic tranquility…promote the general welfare, and secure the blessings of liberty to ourselves and our

---

[37] Nat'l Academies of Sciences, Eng'g, & Medicine, Measuring Racial Discrimination, 221 (2004).
[38] Fair Housing Act, 42 U.S.C. § 3603.

posterity."[39] Instead, the diminished collection and dissemination of data on protected classes will "promote impunity," "hinder adequate policymaking," and dispose of the very evidence that could bring discrimination to light.[40] Without this data, discrimination will be free to metastasize with no recourse. By discouraging the gathering of protected class data, HUD is essentially tying its own hands—avoiding the fight to end discrimination by simply not collecting data of where there might be discrimination in the first place.

Second, protected class data is necessary in the very lawsuits that plaintiffs may bring against Providers. In *Inclusive Communities*, the Supreme Court stated that while "a disparate-impact claim that relies on a statistical disparity must fail if the plaintiff cannot point to a…policy…causing that disparity…[a] plaintiff who fails to allege facts at the pleading stage or *produce statistical evidence* demonstrating a causal connection cannot make out a prima facie case of disparate impact."[41] Essentially, the Supreme Court explained that statistical evidence of disparate impact is insufficient on its own, yet it acknowledged that the data is a significant portion of a plaintiff's claim.[42] By reassuring Providers that they do not need to collect protected class data, HUD reduces the availability of one of the critical tools that plaintiffs have when deciding if and how to bring a lawsuit.

---

[39] U.S. Const. pmbl.

[40] Press Release, U.N. Gen. Assembly: Combating racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action (Nov. 3, 2015), https://www.ohchr.org/EN/NewsEvents/Pages/DisplayNews.aspx?NewsID=16704&LangID=E (last visited Oct. 18, 2019).

[41] *Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Cmtys. Project, Inc*., 135 S. Ct. 2507, 2523 (2015) (emphasis added).

[42] *Id.*

Because the Proposed Rule clarifies and reiterates that there is no requirement for Providers to collect protected class data, it acts counter to the FHA's goal of decreasing discrimination, and because it will further limit plaintiffs' ability to find justice, HUD should not implement this Proposed Rule. HUD instead should pursue the opposite goal and at least encourage, if not incentivize or even require, Providers to collect this protected class data, so that everyone working to make our country *less* divided can have the tools they need to do so.

Conclusion

The Supreme Court in *Griggs v. Duke Power Company* set forth the proposition that "[c]ourts should avoid interpreting disparate-impact liability to be so expansive as to inject racial considerations into every housing decision."[43] But neither should courts interpret disparate-impact liability so narrowly as to blindly ignore discrimination where it does infect housing decisions. To fulfill its mission of building "inclusive and sustainable communities free from discrimination,"[44] HUD must not allow Providers to evade liability for the discriminatory effects of their policies by hiding behind algorithmic models and declining to collect data. For the reasons stated in this comment, we urge HUD to reconsider its Proposed Rule.

---

[43] *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971).

[44] Mission, U.S. Dep't of Housing & Urban Development, https://www.hud.gov/about/mission (last visited Oct. 11, 2019).

Snehee Khandeshi
snehee@berkeley.edu

Brenna Kearns
brenna_kearns@berkeley.edu

Tal Solovey
tsolovey@berkeley.edu

Consumer Advocacy and Protection Society
University of California, Berkeley School of Law
https://consumer.berkeley.edu/
caps@law.berkeley.edu