Loyola Consumer Law Review

Volume 33 | Issue 2 | Article 6

2021

Detecting Disparate-Impact Discrimination in the Big-Data Era

John Yinger
the Maxwell School, Syracuse University, joyinger@syr.edu

DETECTING DISPARATE-IMPACT DISCRIMINATION IN THE BIG-DATA ERA

*John Yinger**

I. INTRODUCTION

The terms at the core of this paper are well known: A business practice has a "disparate impact" if it is applied to every customer but leads to poorer outcomes for one legally protected class than for another. This impact becomes "disparate-impact discrimination," which is illegal, if this disparate impact cannot be justified on the basis of "business necessity." In this context, a business practice can be labeled a "necessity" if it furthers a legitimate business interest. Disparate-impact discrimination contrasts with "disparate-treatment discrimination," which arises when businesses use less-favorable rules for people in legally protected classes. An effective civil rights enforcement system must be able to detect and combat both types of discrimination.

This paper provides an economist's view of disparate-impact discrimination, explains why outlawing disparate-impact discrimination is such an important feature of civil rights law enforcement, reviews the recent debate about the regulations for enforcing the prohibition against disparate-impact discrimination, and makes some recommendations for federal anti-discrimination policy. The focus is on disparate-impact discrimination that may arise in the actions of large lenders, large landlords, or other businesses that have measurable objectives and deal with many customers.[1]

---

* This paper was prepared for the symposium on "Racial Justice in Consumer Law" sponsored by the Loyola University Chicago School of Law and its Consumer Law Review, March 4-5, 2021. John Yinger is a professor of economics and public administration at the Maxwell School, Syracuse University.

[1] Although disparate-impact discrimination may also arise in labor markets, home insurance markets, and automobile loans, these markets have special features that are beyond the scope of this paper. *See, e.g.*, Solon Baroca & Andrew

Many court decisions have supported the application of civil rights laws to disparate-impact discrimination. An important recent example is the U.S. Supreme Court's decision in *Texas Department of Housing and Community Affairs v. Inclusive Communities Project* (*Inclusive Communities*), which affirmed that the Fair Housing Act (FHA) covers disparate-impact discrimination.[2] As in other types of civil rights enforcement, the courts also play a crucial role in guiding the details of enforcement actions involving disparate-impact discrimination. This paper focuses on the broad issues involved in these actions, but does not, for the most part, review the legal details.

II. WHAT IS DISPARATE DISCRIMINATION AND WHY DOES IT ARISE?[3]

An analysis of disparate-impact discrimination begins with the observation that many types of business decisions are guided by a prediction about the outcome of that decision based on the traits of the customer involved at the time the decision is made. Lenders predict the probability that a mortgagor will default, for example, and landlords predict whether a particular tenant will

---

D. Selbst, *Big Data's Disparate Impact*, 104 CAL L. REV. 671 (2016) (on labor markets); Latonia Williams, *African American Homeownership and the Dream Deferred: A Disparate Impact Argument against the Use of Credit Scores in Homeownership Insurance Underwriting*, 15 CONN. INS. L.J. 295 (2008-2009) (on home insurance markets); Mark A. Cohen, *Imperfect Competition in Auto Lending: Subjective Markup, Racial Disparity, and Class Action Litigation*, 8 REV. OF L. & ECON. 21 (2012) (on automobile loans).

[2] *Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Cmtys. Project, Inc.*, 135 S. Ct. 2507 (2015). *See also* Robert G. Schwemm. *Fair Housing Litigation After Inclusive Communities: What's New and What's Not*, 115 COL. L. REV. SIDEBAR 206 (2015) (for the impact of *Inclusive Communities* on fair housing cases); Winnie F. Taylor, *The ECOA and Disparate Impact Theory: A Historical Perspective,* 26 J. OF L. & POL'Y 575 (2018) (for an analysis of the applicability of the Equal Credit and Opportunity Act to disparate-impact discrimination).

[3] The analysis in Sections 2 and 3 draws heavily on Stephen L. Ross & John Yinger, *Uncovering Discrimination: A Comparison of the Methods Used by Scholars and Civil Rights Enforcement Officials*, 8 AMER. L. & ECON. REV. 562 (2006) and STEPHEN L. ROSS & JOHN YINGER, THE COLOR OF CREDIT: MORTGAGE LENDING DISCRIMINATION, RESEARCH METHODOLOGY, AND FAIR-LENDING ENFORCEMENT (MIT Press, 2002). These references provide a more rigorous and complete version of this analysis.

pay his rent on time. The transaction is completed or its terms are favorable if and only if the customer's predicted "performance" is deemed satisfactory. A business predicts a customer's future behavior based on the observed performance of comparable customers in the past. If a business's customer base is sufficiently large, the predicted performance of current customers is based on a statistical analysis of the performance of past customers. Some of the U.S. Department of Housing and Urban Development (HUD) regulations described in Section 4 call this an "algorithmic" model. This type of model is often matched with business practices based on technological innovations, called FinTech, such as mortgage underwriting over the internet and the use of non-traditional data.

This type of statistical analysis identifies a "weight" to be placed on each customer trait and leads to a "score" for each customer defined as a weighted average of all their relevant traits. When a business has many customers and the data it uses to obtain customer scores are available, an enforcement agency can determine, using methods discussed in Section 3, whether the business engages in either disparate-treatment or disparate-impact discrimination. Of course, these conditions often are not met, and a civil rights enforcement agency must develop methods to detect discrimination for small businesses and with limited data. Although this paper focuses on detecting discrimination when many traits are observed for many past customers, some of the key lessons apply to other circumstances, as well.

Consider the following stylized analysis of a business decision, D, such as the decision to grant someone a mortgage loan or rent someone an apartment. The framework developed here also can be applied to a decision about the term of a mortgage loan or of an apartment lease. A business makes this decision based on the observable traits of the customer, such as her income, and of the transaction, such as the loan-to-value ratio. Let X indicate the set of traits a business uses to make decision D. The business places a weight on each of the X variables to obtain a score for each customer.[4] The higher the score, the more favorable the decision from the customer's point of view. The business determines these weights by observing which traits in X best predict the business's preferred outcome as observed in previous transactions. In the case of a mortgage lender, for example, the preferred outcome is likely to be a low probability of loan default within the first four or

---

[4] A credit score is an example of a "score" as the term is used here, but the definition in this paper is more general.

five years after the loan was granted.[5] Let P measure the "performance" of a transaction, defined as the extent to which the transaction meets the business's goals. Then the weight for each X in determining a current customer's score is the coefficient of that X in a statistical analysis of P as a function of the Xs. In other words, the weight placed on a given X in a current transaction equals the statistically determined impact of that X on the firm's performance goals as seen in past transactions. The required statistical analysis could be conducted by the business itself or by an appropriate third party.

This framework leads to precise definitions of disparate-treatment and disparate-impact discrimination. Disparate-treatment discrimination exists when membership in a legally protected class, say M, is one of the variables included in X and the weighted value of M (based on a statistical procedure) is included in a customer's score. For example, a landlord cannot use a scheme in which a white customer is offered an apartment but a black customer with identical traits is not.

A necessary condition for disparate-impact discrimination to arise is that the scoring scheme on which a business makes its decisions diverges from the best-possible scoring scheme based on available information. This divergence is undoubtedly favorable to some customers and unfavorable to others. This condition is not sufficient, however, because the customers who are treated less favorably may not belong to any legally protected class. The obvious sufficient condition is that the business's deviations from the best-possible scoring scheme place a burden on a protected class—even though those deviations apply to all customers.

Suppose, for example, that the best-possible scoring scheme to guide a landlord's rental decisions includes all an applicant's income, with the same weight on income from full-time and part-time jobs. Now suppose that Hispanics are more likely to have part-time work than Anglos, and that a particular landlord uses a scoring scheme that includes the share of income from part-time work. Even if this scheme is used for all customers, it places Hispanic applicants at a disadvantage and constitutes disparate-impact discrimination.

---

[5] As discussed by Robert Bartlett, Adair Morse, Richard Stanton & Nancy Wallace, *Consumer-Lending Discrimination in the FinTech Era* (Nat'l Bureau of Econ. Rsch., Working Paper No. 25943, 2019), the courts have consistently ruled out profit as an appropriate performance measure for lenders. After all, as discussed below, discrimination may be profitable under some circumstances.

One key point to preview is that disparate-impact discrimination cannot be detected unless the decisions made by a business can be compared to the decisions that would be made with the "best possible" scoring scheme. For the purposes of this paper, the best possible scoring scheme is the one developed by an enforcement agency using data supplied by the business (and perhaps other data if they are available) and the most appropriate statistical methods. The enforcement agency needs to observe the business's screening and pricing decisions, but it does not need to know the business's scoring scheme. In the standard framework in federal regulations, a business accused of disparate-impact discrimination against a given protected class has the right to show that its own scheme leads to more favorable outcome for the protected class than does the scheme developed by the enforcement agency.

The logic of disparate-impact discrimination can be applied not just to the inclusion in a scoring scheme of inappropriate traits, such as the share of income from part-time work in the above hypothetical example, but also to traits that are given the "wrong" weight, that is, a weight not associated with the best available scoring scheme. Suppose, for example, that the enforcement agency's scoring scheme based on its own statistical analysis finds that income from part-time work and income from full-time work have the same regression coefficient in the performance regression and hence the same weight in the enforcement agency's scoring scheme. With more part-time work among Hispanics, therefore, a scoring scheme that weights income from a part-time job at half the weight of income from a full-time job would constitute disparate-impact discrimination against Hispanics.

Disparate-impact discrimination also can arise when the statistical analysis linking customer performance, P, to customer and transaction traits, X, omits variables for group membership, M. At first glance, this appears to be a contradiction. How can including group membership in a statistical analysis help to prevent discrimination? The answer to this question is a based on a statistical concept known as "omitted variable bias." When a variable with explanatory power is omitted from a regression then the coefficient of each included variable is biased, and the bias is proportional to the correlation between the included and omitted variables.[6] When M is omitted, therefore, the coefficients of all

[6] This type of bias is not confined to the case of linear regression. It also arises with logit or probit regressions and to models developed through data mining. *See, e.g.*, Scott Menard, *An Introduction to Logistic Regression*

included variables that are correlated with M are biased—and so are the scoring weights based on those coefficients. As shown in the previous paragraph, inaccurate weights such as these, can be a source of disparate-impact discrimination. To avoid disparate-impact discrimination, therefore, a statistical analysis of the relationship between customer performance, P, and observed traits, X, should include minority status in the regression as an element of X, but should (to avoid disparate-treatment discrimination) ignore M in calculating the associated customers' scores.[7]

One final point about this economic framework is that it does not depend on the distinction between intended and unintended discrimination. This distinction is important in the context of a trial because direct evidence of intent, which always strengthens a case against an alleged discriminator, is often missing in a disparate-impact case. This distinction is not appropriate in a conceptual framework for discriminatory behavior, however. As I use the terms, disparate-treatment discrimination can be unintended and disparate-impact discrimination can be intentional.[8]

Moreover, a business may have an economic incentive to practice disparate-impact discrimination. Suppose, for example, that (1) a lender cannot observe whether a loan applicant has family members who can bail her out if she gets into financial trouble and (2) in the pool of applicants available to this lender blacks are less likely than whites to have this financial backstop. Under these

---

Diagnostics, APPLIED LOGISTIC REGRESSION . ANALYSIS (2011). (DOI: https://dx.doi.org/10.4135/9781412983433 , p. 3.)

[7] In a review of Ross, *supra* note 3, Cynthea E. Geerdes, *Now You See It, Now You Don't: When Color-Conscious Means Color-Blind*, 13 J. OF AFFORDABLE HOUS. & CMTY DEV. LAW 176, 177 (2004) restated this point in a particularly clear way: "The framework they construct ... leads them to devise a new, startling, counterintuitive method for preventing most mortgage discrimination in loan approval. Basically, the method requires that minority status first be included as a variable in the underwriting model and then subsequently be discounted on average. In the context of automated underwriting, their goal is to achieve a self-correcting system in which all hidden bias towards minorities can be removed."

[8] A striking example of intentional disparate-impact discrimination is provided by Valerie Schneider, *In Defense of Disparate Impact: Urban Redevelopment and the Supreme Court's Recent Interest in the Fair Housing Act*, 79 MO. L. REV. 539 (2014). To keep African-Americans who had been displaced by Hurricane Katrina from moving back to St. Bernard Parish, which is adjacent to New Orleans, the Parish passed an ordinance that restricted property owners in the Parish (who were mostly white) to renting only to their own blood relatives.

circumstances, the lender can make more accurate predictions of a customer's probability of loan default if it includes in its credit scoring scheme variables that are highly correlated with race. In other words, this business can increase its profits without practicing disparate-treatment discrimination by designing a scoring scheme that includes variables predicting a customer's race instead of actually including race in the scheme.[9] This behavior is an example of "statistical discrimination," which arises when an economic agent uses group membership to account for a variable that helps to predict business objectives but cannot be observed directly.[10]

More generally, any economic actor who wants to discriminate can design a disparate-impact discrimination scheme that yields the same, or almost the same, economic outcomes as a disparate-treatment discrimination scheme, but with a lower probability of being caught. Disparate-impact schemes rely on the correlation between membership in a protected class and customer traits that are plausibly related to customer "quality." This nation's history of oppression against certain groups has resulted in many such traits. As a result, an enforcement agency must be prepared to identify cases of disparate-impact discrimination that coincide with statistical discrimination.

Several recent studies provide evidence that discrimination persists in housing and mortgage markets.[11] One study of

---

[9] Henry Buist, Peter Linneman, & Isaac F. Megbolugbe. *Residential Lending Discrimination and Lender Compensation Policies*, 27 J. OF THE AMER. REAL EST. & URB. ECON. ASSN. 695 (1999),

[10] The concept of statistical discrimination was introduced by Kenneth J. Arrow, *The Theory of Discrimination*; *see* ORLEY ASHENFELTER & ALBERT REES, DISCRIMINATION IN LABOR MARKETS (Princeton Univ. Press 1973); Edmund S. Phelps, *The Statistical Theory of Racism and Sexism*, 62 AM. ECON. REV. 659 (1972).

[11] Studies that document ongoing discrimination in housing and mortgage markets include: Sun Jung Oh & John Yinger, *What Have We Learned from Paired Testing in Housing Markets?*, 17 CITYSCAPE: A J. OF POL'Y DEV. AND RES. 15 (2015) (reviewing extensive evidence of discrimination in housing based on the paired-testing methodology); Andrew Hanson, Zachary Hawley, Hal Martin & Bo Liu, *Discrimination in Mortgage Lending: Evidence from a Correspondence Experiment*, 92 J. OF URB. ECON. 48 (2016) (finding evidence of discrimination by mortgage loan originators based on a paired testing study); Ann Choi, Keith Herbert, & Olivia Winslow, *Long Island Divided*, NEWSDAY (Nov. 17, 2019) (reporting on housing discrimination found in paired tests conducted on Long Island); Patrick Bayer, Fernando Ferreira, & Stephen L. Ross,

particular relevance here, Bartlet et al.,[12] examines discrimination in the setting of interest rates by lenders who sell their loans to Fannie Mae or Freddie Mac (Government Sponsored Enterprises or GSEs). From the lender's point of view, the GSEs fully insure each loan for a fee that is based on the loan's loan-to-value ratio and the borrower's credit score. As a result, a higher interest rate for black or Hispanic customers controlling for these two factors is an indication of discrimination. This study estimates that discrimination in interest rate setting on home-purchase and refinancing loans costs African-American and Hispanic borrowers $765 million per year.[13] The level of discrimination is somewhat lower for lenders classified as "FinTech" than for others.

III. HOW TO DETECT DISPARATE-IMPACT DISCRIMINATION

The Bureau of Consumer Finance Protection (CFPB) are responsible for the regulations that implement the Equal Credit Opportunity Act (ECOA). These regulations, called Regulation B, provide a clear definition of disparate-impact discrimination.[14] Moreover, they define an "empirically derived, demonstrably and statistically sound, credit scoring system," which is the type of system discussed in this paper.[15] In addition, these regulations acknowledge that "neutral factors used in credit scoring systems could nonetheless be subject to challenge under the effects test" which is another name for a test to find disparate-impact discrimination.[16] Regulation B does not say anything, however, about the steps a regulator needs to take to determine whether a lender using this type of credit-scoring system is practicing disparate-impact discrimination.

The U.S. Department of Housing and Urban Development (HUD) is responsible for regulations to implement FHA.[17] These regulations also provide a detailed definition of disparate-impact

---

What Drives Racial and Ethnic Differences in High-Cost Mortgages? The Role of High-Risk Lenders, 31 REV. OF FIN. STUD. 175 (2018) (finding that African-American and Hispanic customers are more likely than equally qualified white customers to be steered to high-cost loans).

[12] Bartlett, *supra* note 5.

[13] *Id.* at 15.

[14] Equal Credit Opportunity Act – Regulation B, 12 C.F.R. § 1002 (1974).

[15] 12 C.F.R. § 1002.2 (2011).

[16] *Id.*

[17] Discriminatory Conduct Under the Fair Housing Act, 24 C.F.R. § 100.

discrimination. The 2013 version of these rules set out a three-part process for a court case in which disparate-impact discrimination was alleged.[18] (Newer rules, which are not yet implemented, are discussed below.) First, the plaintiff (i.e., the party alleging discrimination) had to establish that a practice had a disparate impact on people in a protected class. If the plaintiff was successful in making this prima facie case, the burden of proof shifted to the defendant. To offset the prima facie case, the defendant had to establish that the practice in question was necessary to achieve the defendant's legitimate, non-discriminatory objectives. If this "business necessity" was established, the burden of proof returned to the plaintiff, who could still prevail if it could identify an alternative practice that was consistent with the defendant's "business necessity" but has a smaller disparate impact on the plaintiff. This is a compelling framework but these regulations do not provide clear guidance about the methods required to test for disparate-impact discrimination in the case of a statistically based scoring scheme for credit or for housing decisions.

Although a test for disparate-impact discrimination does not appear in the federal regulations, even for the case of a data-rich environment, tests of this type are available. One type of test is presented by Ross and Yinger.[19] Another comes from a simple extension of the analysis in Bartlett et al.[20]

With the appropriate data, the five steps identified by Ross and Yinger rigorously identify disparate-treatment and disparate-impact discrimination. These steps are:

1. Gather data on applicant and transaction traits for a large sample of individuals who interacted with a given business, presumably a business suspected of practicing disparate-impact discrimination or a business that has been randomly selected for investigation.[21] These traits must include the applicant's

[18] HUD's Implementation of the Fair Housing Act's Disparate Impact Standard, 78 Fed. Reg. 11, 460 (Feb. 15, 2013). A more formal discussion of the 2013 rules can be found in Leah Powers, *The Uncertain Future of the Fair Housing Act: HUD's Recent Changes to Disparate Impact Standard*, 74 SMU L. REV. F. 29 (2021).

[19] Ross and Yinger (2002), *supra* note 3.

[20] Bartlett et al., *supra* note 5.

[21] An enforcement agency also could conduct its statistical procedures with pooled data and then extract the results for a single business. *See* Ross and Yinger (2002), *supra* note 3.

membership in a legally protected class. The data set also should include whether the individual was served by the business and/or the key terms of this service, such as the interest rate or monthly rental. Finally, the data set should follow each individual for several years to measure his or her performance, such as whether that individual was late on a rent payment.

2. The second step is to randomly divide the sample of individual into two parts, say A and B.

3. The third step is for the civil rights enforcement agency to develop its own statistical model of the way individual and transaction characteristics observed at the time of transaction influence the individual's performance for several years following the transaction. This model should be estimated for one of the two samples, say, sample A. To conduct this step the enforcement agency must invest in knowledge about the underlying behavior. What have scholars learned, for example, about the impact of factors observed at the time of loan origination on the likelihood that a customer will default within five years? The agency also needs to invest in knowledge about which statistical formulation best predicts individual behavior as a function of these factors. To avoid the omitted variable bias discussed earlier, this model must include variables to indicate a customer's membership in various protected classes.[22]

4. The fourth step is to use the results of step three, such as regression coefficients, to obtain a score for each individual in sample B. To keep the analog to disparate-treatment discrimination out of the regulator's model, this score obviously must not include the

---

[22] Barocas and Selbst, *supra* note 1 at 243, conclude that "there is no obvious way to determine how correlated a relevant attribute must be with class membership to be worrisome. Nor is there a self-evident way to determine when an attribute is sufficiently relevant to justify its consideration, despite its high correlation with class membership." When a measure of performance is available, however, an obvious way exists: include minority status in the scoring regression.

individual's membership in any protected class. In effect, regression coefficients of the protected class variables should be estimated in step 3 and then set to zero in step 4. These individuals must then be ranked according to their score.

5. The fifth step takes one of two forms, depending on the nature of the decision under investigation. First consider the case of a screening decision by a business, that is, a decision about whether to serve a given customer. The decision to approve a mortgage is an example. The enforcement agency can observe the number of individuals served in sample B, say N_B, and the minority share of those served individuals. This share can then be compared to the minority share of the N_B highest-ranking individuals according to the enforcement agency's scores, which are determined in step 4. A significantly higher minority share for the individual's ranked using the agency's scores than the minority share in actual loans is evidence of discrimination. This discrimination could be either disparate-treatment or disparate-impact. These two types of discrimination cannot be distinguished without knowledge of the business's scoring model.

The second form for this test applies when the terms of a transaction, such as an interest rate or a security deposit, are being examined. In this case the test, again based on sample B, is to regress the relevant term on minority-status variables and the agency's score for each individual. A significant coefficient for the minority status variable is evidence of discrimination.[23] Again, disparate-treatment and disparate-impact cannot be separately identified.

A key feature of these steps is that they do not require information about the business's decision rules. It is, of course, possible that the business has a better non-discriminatory model to predict customer performance than the one devised by the enforcement agency. A business should always be given the chance to show that this is true. They might be able to show, for example, that the agency's model leaves out a legitimate variable that has a

[23] Discrimination often takes the form of a higher interest rate or a higher security deposit, so a positive sign for a minority status variable indicates discrimination in these cases.

major impact on the scores of many individuals and that including this variable in the agency's model eliminates the difference in minority shares between the agency and business models—and hence the evidence of discrimination. Of course, the business must reveal its own model in order to make this type of claim.[24]

This method was applied to the pricing decisions of a large lender in California and Florida by Ghent et al.[25] This study found that the "upper bound for the effect of adverse pricing based on the borrower's race that is not due to differences in prepayment or default behavior is 29 basis points, the adverse pricing faced by blacks in 30-year ARMs. An increase in the interest rate of 29 basis points translates into an increase in the monthly payment of $57.57 or 3% of the payment" (p. 208).[26]

One complexity in contemporary credit markets is that disparate-impact discrimination could operate through "channels." As explained by White,[27] "It appears that within a large financial institution ..., mortgage prices for borrowers with similar credit scores and qualifications vary widely according to channels and products. It is as if a store charged higher prices for customers coming in the east door than those coming in the west door, and somehow directed most minority customers to the east door. While it is conceivable that a lender ... could offer cost-driven business justifications for charging different prices for loans made through different channels (e.g., broker versus retail) it is hard to see how a lender could support selling the same product-adjustable rate first

---

[24] If a firm decides to reveal its underwriting model, then an enforcement agency can separate disparate-treatment and disparate-impact discrimination.

[25] Andra C. Ghent, Rubén Hernández-Murillo & Michael T. Owyang, *Differences in Subprime Loan Pricing across Races and Neighborhoods*, 48 REG. SCI. & URBAN ECON. 199 (2014).

[26] Robert B. Avery, Kenneth Brevoort & Glenn B. Canner, *Does Credit Scoring Produce a Disparate Impact?*, 40 REAL EST. ECON. S65 (2012) (develop a model to generate individual credit scores and find that the credit scores obtained from their model are virtually the same when estimated with or without racial and ethnic fixed effects and when estimated for the white, non-Hispanic sample alone. As they point out, however, "Perhaps the most important [limitation] is that the analysis is based upon a credit history scoring model that was developed specifically for this study and not upon a commercially available score" (p. S113). In other words, they skip step 5 in the procedure developed above, and therefore cannot draw conclusions about disparate-impact discrimination in commercial credit scores).

[27] Alan M. White, *Borrowing While Black: Applying Fair Lending Laws to Risk-Based Mortgage Pricing*, 60 SOUTH CAR. L. REV. 677, 698 (2009).

lien subprime refinance mortgages-at different prices using different names by business necessity." An application of the five above steps should capture discrimination in these cases, if it exists.[28]

A straightforward extension of the method in Bartlett et al.,[29] provides an alternative test for lenders who sell their loans to a GSE. Bartlett et al. measure discrimination based on regression of borrower's interest rate on fixed effects for each risk category identified by the GSE and a variable to indicate membership in a protect class (African-American or Hispanic in their study). The coefficient of the protected class variable indicates the extent of discriminatory interest rates in the average transaction. Although Bartlett et al. do not point this out, their approach can be turned into a test for pricing discrimination by a single lender by limiting the sample to loans by a single lender or by interacting a fixed effect for each lender with the minority-status variable. Of course, these fixed effects only lead to meaningful results for large lenders.

IV. THE TRUMP ADMINISTRATION'S CHANGES TO HUD REGULATIONS

The 2013 HUD regulations concerning disparate-impact discrimination posted by the Obama Administration[30] did not sit well with the Trump Administration, and on August 19, 2019, HUD proposed a different set of rules to guide the enforcement of cases based on disparate-impact discrimination.[31] These rules include a section on the use of "algorithmic models."

HUD claimed that "[t]his section is not intended to provide a special exemption for parties who use algorithmic models, but merely to recognize that additional guidance is necessary in response to the complexity of disparate-impact cases challenging

[28] Although it does not estimate discrimination by individual lenders, the study by Bayer, Ferriera, and Ross (*supra* note 11 at 200) "identified large racial and ethnic differences in the likelihood of receiving a rate spread mortgage in the home purchase market after controlling for detailed borrower and loan attributes. Differential sorting across lenders and the differential treatment of equally qualified borrowers by the same lender both emerge as important drivers of market-wide differences."

[29] Bartlett et al., *supra* note 5.

[30] HUD's Implementation of the Fair Housing Act's Discriminatory Effects Standard, 78 Fed. Reg. 11, 460 (Feb. 15, 2013).

[31] HUD's Implementation of the Fair Housing Act's Disparate Impact Standard, 84 Fed. Reg. 42, 854 (Aug. 9, 2019).

these models."[32] From my point of view, however, the rules in this section would have made it virtually impossible for any plaintiff to prevail on a disparate-impact claim involving an "algorithmic model."

HUD summarized the proposed new rules on the use of algorithms as follows:[33]

> Paragraph (c)(2) provides that, where a plaintiff identifies an offending policy or practice that relies on an algorithmic model, a defending party may defeat the claim by: (i) Identifying the inputs used in the model and showing that these inputs are not substitutes for a protected characteristic and that the model is predictive of risk or other valid objectives; (ii) showing that a recognized third party, not the defendant, is responsible for creating or maintaining the model; or (iii) showing that a neutral third party has analyzed the model in question and determined it was empirically derived, its inputs are not substitutes for a protected characteristic, and is a demonstrably and statistically sound algorithm.

Although these rules address legitimate issues, they are meaningless as written. Almost any input is correlated with protected characteristics, but no input is literally a substitute for any one of these characteristics. A plaintiff obviously cannot succeed if success requires showing something that never occurs. Moreover, membership in a protected class is often "predictive of risk," so making evidence of this prediction a valid defense opens the door for discriminatory practices. The role of third parties is equally vague.

The procedure for identifying disparate-impact discrimination in Section 3 recognizes, of course, that potential inputs into a scoring system may be correlated with a protected characteristic and that a defendant has the right to develop a scoring system that is consistent with business necessity. However, the method in that section is based on a formal consideration of these two issues—not the inappropriate rules in these proposed amendments.

---

[32] *Id.* at 42, 859.

[33] *Id.*

HUD received 45,758 comments on these proposed rules and altered them significantly.[34] The "final" rules were posted on September 24, 2020 to be effective a month later. The 2020 rules dropped the provisions in the 2019 proposed provisions on "algorithmic models." The explanation was that

> HUD expects that there will be further development in the law in the emerging technology area of algorithms, artificial intelligence, machine learning and similar concepts. Thus, it is premature at this time to more directly address algorithms.[35]

Although the troubling provisions concerning algorithmic models were dropped from the 2020 rules, other changes undermined HUD's ability to combat disparate-impact discrimination. As Powers puts it

> Although the 2020 Rule arguably finds textual support in the language of *Inclusive Communities* for several of its five elements, HUD made the deliberate choice to go beyond the Court's language in crafting its new standard. By heightening the plaintiff's burden even further than required by the Court's ruling and providing advantageous new defenses for defendants, HUD has greatly reduced the effectiveness of disparate impact liability as a tool to fight housing discrimination.[36]

Both versions of the Trump Administrations rules also add a new paragraph: "Nothing in this part requires or encourages the collection of data with respect to race, color, religion, sex, handicap, familial status, or national origin."[37] Given the important role of group membership variables in a test for disparate-impact discrimination, this paragraph would be an unfortunate addition to the HUD regulations.

V. LESSONS FOR THE BIDEN ADMINISTRATION

The word final is in quotation marks in the previous paragraph, because "the U.S. District Court for the District of

---

[34] HUD's Implementation of the Fair Housing Act's Disparate Impact Standard, 85 Fed. Reg. 60, 288 (Sept. 24, 2020).

[35] *Id.* at 60, 290.

[36] Powers, *supra* note 18, at 54.

[37] HUD, *supra* note 18, at § 100.5(d).

Massachusetts granted a preliminary injunction staying the implementation of the 2020 Rule on October 25, 2020—just one day before the rule was scheduled to take effect."[38] Then, on January 26, 2021, President Biden issued a memorandum expressing his commitment to fair housing and calling on the Secretary of HUD to "as soon as practicable, take all steps necessary to examine the effects of" the rules posted by HUD in September, 2020.[39]

The analysis in this paper leads to several recommendations for the Biden Administration's fair housing policies. First, this analysis supports the conclusion in Powers:

> In light of the broad remedial intent of the FHA, the core mission of HUD to eradicate housing discrimination, decades of disparate impact precedent, various negative consequences of the new standard, and President Biden's recent memorandum on housing discrimination, HUD should abandon the 2020 Rule and readopt the 2013 Rule.[40]

In addition, it does not appear "premature" to me for HUD and CFPB to investigate changes in their rules designed to strengthen their disparate-impact discrimination enforcement powers given the growth in big data, algorithmic models, and FinTech methods.

Second, the available evidence suggests that federal fair housing and fair lending enforcement agencies should take the possibility of disparate-impact discrimination more seriously. The available public record does not indicate the extent to which fair housing and fair lending enforcement officials explored disparate-impact issues in their examinations and investigations. However, the annual fair lending and enforcement reports by CFPB, HUD's Office of Fair Housing and Equal Opportunity, and The Justice Department's Civil Rights Division do not mention any cases during the last five years that involve disparate-impact discrimination.[41] This lack of disparate-impact cases appears to be in conflict

---

[38] The quotation is from Powers, *supra* note 18 at 52. The case is *Mass. Fair Hous. Ctr. v. U.S. Dep't of Hous. & Urb.* Dev., No. 20-11765, 2020 U.S. Dist. LEXIS 205633, at 1 (D. Mass. Oct. 25, 2020).

[39] *Memorandum on Redressing Our Nation's and the Federal Government's History of Discriminatory Housing Practices and Policies*, 2021 DAILY COMP. PRES. DOC. 202100090, at 2 (Jan. 26, 2021).

[40] Powers, *supra* note 18, at 54.

[41] *See Fair Lending Report of the Bureau of Consumer Finance Protection*, CONSUMER FIN. PROT. BUREAU (2019),

with the scholarly evidence that racial and ethnic discrimination persists.

Third, the civil rights enforcement agencies in the federal government, particularly HUD and CFPB, must recognize that they have a unique responsibility to develop the expertise necessary to enforce FHA and ECOA in an era of big data and elaborate scoring systems. This point also has been made by Gano,[42] who writes:

> that there is simply no substitute for governmental oversight and enforcement of fair lending laws. Disparate impact cases in the mortgage lending context often involve thousands (and in some recent cases ... hundreds of thousands) of borrowers spread across the country. The evidence is complex and consists of millions of data points contained in reports submitted to and reviewed by federal financial regulators. Individual plaintiffs—and even state attorneys general and national nonprofits—quite simply lack the expertise and economic wherewithal to pursue these claims. If we are to "provide, within constitutional limitations, for fair housing throughout the United States," the federal government must continue to play a leading role.

Fourth, HUD and CFPB must recognize that the enforcement of these civil rights laws requires combining existing data

---

https://www.consumerfinance.gov/data-research/research-reports/?topics=fair-lending; *FHEO Annual Report on Fair Housing*, HUD.GOV (2019), https://www.hud.gov/program_offices/fair_housing_equal_opp/annualreport; *Fair Housing Enforcement Activity*, HUD.GOV (2020) https://www.hud.gov/program_offices/fair_housing_equal_opp/enforcement; *The Attorney General's 2019 Annual Report to Congress Pursuant to the Equal Credit Opportunity Act Amendments of 1976*, THE U.S. DEP'T. OF JUST. (2019), https://www.justice.gov/crt/fair-lending-program-0. The Federal Reserve Board does not publish an annual fair lending report. The only mention of a case involving disparate-impact discrimination that I could find on the Federal Reserve Board web site was a 2015 case involving fees and pricing, which was voluntarily settled by the lender. *See* Carol A. Evans, The Fed. Reserve Sys., *Keeping Fintech Fair: Thinking About Fair Lending and UDAP Risks.* CONSUMER COMPLIANCE OUTLOOK (2017), https://www.frbsf.org/banking/files/Fintech-Lending-Fair-Lending-and-UDAp-Risks.pdf.

[42] Alex Gano, *Disparate Impact and Mortgage Lending: A Beginner's Guide*, 88 UNIV. OF COLO. L. REV. 1109, 1116 (2017).

sets. As explained by Ross and Yinger,[43] existing data sets for both research and enforcement generally do not contain all the information needed to test for disparate-impact discrimination by a large business or a set of large businesses. Data sets designed for research, for example, often consist of a sample of individuals without identification of the businesses with which these individuals deal. The Home Mortgage Disclosure Act (HMDA) data set collected by HUD does identify lenders, but it does not follow individual borrowers over time so it cannot provide individual performance information.[44] The study by Ghent et al.[45] collected the data necessary to test for disparate-impact discrimination by one large lender. Avery et al.[46] and Bartlett[47] provide examples of the extensive efforts that are required to assemble the type of data needed to test for disparate-impact discrimination while preserving individual and business confidentiality. Any serious effort to test for disparate-impact discrimination by large businesses requires the Federal Reserve Board, CFPB, HUD, and other agencies with civil rights enforcement authority to find ways to build the required large data sets. Guidelines concerning the creation of such data sets would make a valuable addition to policies concerning disparate-impact discrimination.

Rigorous tests for disparate-treatment and disparate-impact discrimination need to have information on membership in protected classes. Federal regulations regarding ECOA and HMDA require businesses to collect this information under many circumstances, including for fair lending enforcement.[48] Any move to weaken these requirements would constitute a serious blow to the ability of civil rights enforcement agencies to detect both disparate-treatment and disparate-impact discrimination. The paragraph in HUD's 2020 rules eliminating any requirement to collect

---

[43] Ross and Yinger (2006), *supra* note 3, at 584–85.

[44] A report by the U.S. Governmental Accountability Office (*Fair Lending: Data Limitations and the Fragmented U.S. Financial Regulatory Structure Challenge Federal Oversight and Enforcement Efforts*, Report GAO-09-704 (2009)) explores the strengths and limitations of the HMDA data for fair lending enforcement. This report also discusses the weaknesses of the fair lending enforcement system before the founding of the CFPB.

[45] Ghent et al., *supra* note 25.

[46] Avery et al., *supra* note 26. The data for this study were collected under the auspices of the Federal Reserve Board.

[47] Bartlett et al., *supra* note 5.

[48] Equal Credit Opportunity Act (Regulation B) Ethnicity and Race Information Collection, 82 Fed. Reg. 45, 680 (Oct. 3, 2017); 12 C.F.R. § 1002.13 (2018).

these data should clearly be repealed. Instead, federal policy makers should make certain that data on race and ethnicity are collected whenever required for civil rights enforcement. Any data collection efforts should, of course, be accompanied by rules to make certain that businesses supplying the information do not have an opportunity to use it for discriminatory purposes.[49]

Fifth, the civil rights enforcement agencies should develop expertise in models of customer performance. A full evaluation of a business's business-necessity claims, which are inevitably based on its own customer-performance model, requires a comparison with an alternative model developed by the enforcement agency. Recall step 5 of the method in Section 3. This step determines whether outcomes for various protected classes are less favorable with the business's decision rules than with a decision rule based on the agency's scoring model. As in the case of data assembly, the studies by Ghent et al.,[50]Avery et al.,[51] and Bartlett et al.[52] reveal the type of modelling efforts in which an enforcement agency must engage. Revised policies by HUD and CFPB should validate enforcement efforts of this kind.

Overall, the federal civil rights enforcement agencies for lending and housing appear to have avoided enforcement actions involving potential disparate-impact discrimination by large businesses. The analysis in this paper indicates that these agencies should take this possibility more seriously by developing regulations and policies to facilitate such enforcement actions in the future.

Four principal arguments support this conclusion.

> (1) Given our nation's history, membership in a legally protected class often conveys negative information about a customer that a business cannot observe directly.˙ In this case, the business may have an incentive to

[49] This importance of data on group membership in a case where businesses are not required to collect it (non-mortgage loans) is examined by Winnie Taylor, *Proving Racial Discrimination and Monitoring Fair Lending Compliance: The Missing Data Problem in Nonmortgage Credit*, 31 REV. BANK. & FIN. L. 199, 244 (2010-2011). Taylor concludes that "[w]hen lending discrimination plaintiffs are unable to compare themselves to others who may be or are similarly situated, their claims will routinely fail."

[50] Ghent et al., *supra* note 25.

[51] Avery et al., *supra* note 26.

[52] Bartlett et al., *supra* note 5.

discriminate against customers in that class, that is, to practice statistical discrimination.

(2) The scholarly literature indicates that racial and ethnic discrimination have by no means disappeared in housing and mortgage markets.

(3) Disparate-impact discrimination is generally more difficult to detect than disparate-treatment and may therefore be the preferred behavior for businesses with an incentive to practice statistical discrimination—and with no qualms about doing so.

(4) Disparate-impact discrimination can arise from a scoring scheme that is based on a statistical analysis that omits variables indicating membership in a protected class. This omission may be motivated by a business's attempt to be "neutral." However, because membership in a protected class is correlated with many of the variables included in this analysis, this perceived neutrality is an illusion. Instead of preventing disparate-impact discrimination, this approach may inadvertently build it into the business's scoring scheme.