



# Fair Lending Monitorship of Upstart Network's Lending Model

---

Initial Report of the Independent Monitor

---

Pursuant to agreement by the NAACP Legal Defense and Educational Fund, the Student Borrower Protection Center, and Upstart Network, Inc.

April 14, 2021

Relman Colfax PLLC  
1225 19th St. N.W., Suite 600  
Washington D.C. 20036  
(202) 728-1888

## Table of Contents

Executive Summary .....	3
A. Fair Lending, Alternative Data, and Alternative Algorithms .....	5
1. Legal Background .....	5
2. Fair Lending Testing .....	7
a. Disparate Impact Step 1 .....	9
b. Disparate Impact Step 2 .....	10
c. Disparate Impact Step 3 .....	11
3. Credit Disparities and the Rise of Alternative Data and Alternative Algorithms .....	12
a. Disparities in Access to Credit .....	12
b. The Thin File Challenge .....	13
c. What are Alternative Data? .....	14
d. What are Alternative Algorithms? .....	15
e. Fair Lending Implications .....	15
B. Upstart’s Lending Model and Concerns Leading to this Monitoring Engagement .....	18
1. Development of Upstart’s Alternative Data Model .....	18
2. Upstart’s 2017 No Action Letter and Compliance Plan .....	19
3. SBPC’s “Educational Redlining” Report and Congressional Responses .....	21
C. LDF & SBPC Engagement with Upstart .....	23
D. Changes to the Upstart Underwriting Model .....	23
E. Upstart’s 2020 No Action Letter and Model Risk Assessment Plan .....	24
F. The Parties’ Agreement to Appoint an Independent Monitor .....	25

## Executive Summary

This is the first report of the independent fair lending Monitor regarding Upstart Network's ("Upstart") lending model. Upstart is a lending platform that relies on Machine Learning-based Artificial Intelligence ("ML" and "AI") models and non-traditional applicant data—including data related to borrowers' higher education—to underwrite and price consumer loans. The NAACP Legal Defense Fund ("LDF") is an organization dedicated to furthering racial justice and the Student Borrower Protection Center ("SBPC") is focused on protecting the rights of student borrowers. In 2020, LDF and the SBPC raised concerns with Upstart that the use of educational criteria can lead to discriminatory lending outcomes, particularly for communities of color.

Against this background, in December 2020, Upstart, LDF, and the SBPC agreed to appoint Relman Colfax, PLLC, as an independent fair lending Monitor to evaluate and make recommendations regarding the fair lending implications of Upstart's lending platform, and to issue a series of reports on its findings and recommendations. This Initial Report is the first such report. Future periodic reports will address fair lending tests conducted and recommendations specific to Upstart's model. This Initial Report, in contrast, is largely descriptive and takes no position on the parties' claims to date. Instead, this Report provides legal and historical background as context for this Monitoring engagement, including a summary of legal principles and fair lending testing, the evolution of Upstart's lending program and model, Upstart's receipt of no-action letters from the Consumer Financial Protection Bureau ("CFPB" or the "Bureau"), and studies and communications from the SBPC, LDF, and members of Congress.

Following conversations with LDF and the SBPC, Upstart changed how it utilizes educational data, including eliminating the use of average incoming SAT and ACT scores to group education institutions in its underwriting model. Upstart also established a "normalization" process for Minority Serving Institutions ("MSIs"), in an effort to ensure that students from MSIs and from non-MSIs are treated, on average, equally. Upstart emphasizes that, to date, its internal fair lending analyses (including those shared with the Bureau) have not identified fair lending violations. The fair lending testing done pursuant to this Monitorship will be of Upstart's model following adoption of these changes.

For decades, antidiscrimination laws such as the Equal Credit Opportunity Act ("ECOA") and the Fair Housing Act ("FHA") have prohibited creditors from discriminating against consumers on the basis of characteristics such as race, national origin, religion, sex, and age. As the lending landscape has evolved, these laws have proven essential to ensure that consumers are treated fairly and that, to the greatest extent possible, structural inequalities and racism are eliminated from the lending system. This ambitious goal has not been achieved. Stark disparities exist across lending markets in everything from approval rates and pricing to the availability of safe credit products and services. At the same time, there are reasons to be optimistic. Regulators, policymakers, advocates, and many market participants are invested in improving equitable access to safe credit products. Increased automation across both consumer and small

business credit markets has lessened the reliance on discretionary or judgmental decision-making, reducing the opportunities for underwriters and loan officers to directly inject animus or bias into individual credit decisions. The use of new sources of applicant data for underwriting and pricing provides an opportunity to help address the credit “Catch-22”: because prior credit history is so heavily weighted in credit determinations, it can be difficult to access traditional credit markets for those without such history.<sup>1</sup> Particularly in combination with some of these new data sources, many observers argue that sophisticated AI/ML algorithms may prove better at predicting the credit risk and optimal loan pricing for applicants who historically would have been excluded from traditional credit markets and left to pursue riskier, costlier, or predatory loan alternatives.

No market innovation should be uncritically embraced, and these developments raise new fair lending risks. The same algorithms that reduce dependence on individual discretion involve reliance on facially-neutral criteria that can imbed and replicate years of historical discrimination and racism, risking heightening barriers and widening inequities in credit markets. The opacity of automated credit decision-making means that where discrimination is occurring, it can be harder for consumers to detect, and therefore harder to challenge. The sheer number of variables used in some AI/ML algorithms increases the risk that variables that proxy for protected class status will be included, and even absent proxies, facially neutral models can exacerbate disparate impact adverse to historically disadvantaged groups. Even where there is reason to believe that particular innovations will have both consumer- and business-friendly benefits, they can have unintended consequences and must be implemented and monitored carefully.

Antidiscrimination laws have worked to combat these risks. Disparate impact, for example, prohibits facially-neutral practices (including models) that disproportionately adversely affect members of a protected class if those practices are not necessary to further business needs or if alternative practices exist that would accomplish those needs with fewer disparities. The central tenet of disparate-impact law is that entities should adopt available alternatives that can satisfy their legitimate needs with less discriminatory effect. Disparate-impact law has been critical in reducing credit inequalities. In the case of automated underwriting, for example, it has caused lenders to search for and implement model variations that predict accurately *and* reduce disparate outcomes.

An important goal of this review by the Monitor is not just to apply these and similar principles to assess Upstart’s underwriting model, but to contribute to the ongoing dialogue about the growing use of AI and alternative data, while deepening our understanding of how to promote equitable access to credit. Upstart, LDF, the SBPC, and the Monitor share the view that with the advent of these innovations, it is increasingly important that lenders take steps to ensure equitable access to responsible credit and avoid the unnecessary perpetuation of discrimination, segregation, and inequality.

---

<sup>1</sup> See, e.g., Richard Cordray, Director, CFPB, *Alternative Data Field Hearing* (Feb. 16, 2017), <https://www.consumerfinance.gov/about-us/newsroom/prepared-remarks-cfpb-director-richard-cordray-alternative-data-field-hearing/>.

## A. Fair Lending, Alternative Data, and Alternative Algorithms

Antidiscrimination laws such as the Equal Credit Opportunity Act (“ECOA”) and the Fair Housing Act (“FHA”) prohibit creditors from discriminating against consumers on the basis of characteristics such as race, national origin, religion, sex, and age. These laws prohibit intentional and overt discrimination, sometimes called “disparate treatment,” as well as an unintentional type of discrimination called “disparate impact.” For years, lenders have understood that these theories apply to their underwriting and pricing models and have developed fair lending evaluations and methodologies to address these risks. At the same time, recent innovations—such as reliance on non-traditional credit criteria and the use of sophisticated AI/ML techniques—have spurred the development of new fair lending testing methodologies and alternative optimal model discovery approaches, while creating new fair lending risks. This section introduces these concepts. It is not intended to serve as a comprehensive discussion of any of these topics. Rather, it is primarily descriptive and lays the foundation for future work related to this review by the Monitor.

### 1. Legal Background

ECOA, originally passed in 1974, is the primary federal legislation prohibiting lending discrimination. ECOA makes it unlawful for a creditor to discriminate against an applicant in “any aspect of a credit transaction” on the basis of certain protected characteristics, including race, color, religion, national origin, sex, age, or receipt of income from a public assistance program.<sup>2</sup> In addition to ECOA, residential real-estate related loans are also subject to the FHA, which prohibits discrimination based on race, color, national origin, religion, sex, disability, or familial status (meaning the presence in the household of a child under the age of 18).<sup>3</sup> The FHA covers second mortgages and home equity lines of credit in addition to standard purchase loans. These laws prohibit a lender, because of a prohibited factor, from refusing to extend credit or using different standards for determining whether to extend credit, varying the terms of credit offered—including the amount, interest rate, duration, or type of loan—or otherwise treating prospective applicants, applicants, and borrowers differently on the basis of a protected characteristic.

The antidiscrimination protections in both laws are broad, encompassing multiple forms of discrimination and discriminatory practices. Most significantly, both ECOA and the FHA prohibit explicit differential treatment or intentional discrimination (known as “disparate treatment”), as well as more subtle forms of discrimination that occur without any intent to

---

<sup>2</sup> 15 U.S.C. § 1691(a); 12 C.F.R. § 1002.2(z). In addition to the listed protected classes, ECOA protects against discrimination based on the good faith exercise of any right under the Consumer Credit Protection Act. 15 U.S.C. § 1691(a)(3). Although sexual orientation and gender identity are not listed as separate classes under ECOA, they are protected. *See* CFPB Interpretive Rule, Equal Credit Opportunity (Regulation B): Discrimination on the Bases of Sexual Orientation and Gender Identity, 86 Fed. Reg. 14363 (Mar. 16, 2021).

<sup>3</sup> 42 U.S.C. § 3605.

discriminate (known as “disparate impact”).<sup>4</sup> Although disparate treatment discrimination often involves animus or the specific intent to harm or disadvantage members the protected group, such animus is not required. A policy or practice which, on its face, treats people differently based on protected class counts as disparate treatment discrimination, even if the policy was adopted with benign intentions.<sup>5</sup>

Policies or practices which do not, on their face, differentiate based on protected class can nevertheless be illegal as a form of disparate impact discrimination. Under disparate impact, a policy or practice that is neutral on its face but disproportionately disadvantages a protected class is illegal if it either does not serve a legitimate business justification, or the legitimate justification can be served in some alternative way that results in less disadvantage to the protected class. There are three steps involved in determining whether a policy has an unlawful disparate impact:

*Step 1:* The first step is to determine whether the policy or practice disproportionately affects or disadvantages a protected class. Frequently, this is done through a quantitative analysis of the outcomes of the policy or practice.

*Step 2:* If the policy or practice does have a disproportionate effect on a protected class, the next step is to identify whether there is a legitimate business interest served by the policy. In the credit context, for example, identifying applicants who would be likely to repay the loan for which they have applied would be a legitimate business interest. If the policy or practice does not serve a legitimate business interest, it is illegal, and the inquiry ends.

*Step 3:* If the policy or practice does serve a legitimate interest, the third and final step is to determine whether there is a reasonable alternate practice or policy that would serve the same end while reducing the disproportionate impact on protected class members. If there is such a less discriminatory alternative that still fulfills the legitimate business justification, it would be unlawful to use the original policy or practice rather than adopting the alternative.

Importantly, this disparate impact framework, including the legitimate justification and less discriminatory alternative requirements, only applies in the absence of disparate treatment. A

---

<sup>4</sup> See, e.g., *Tex. Dep’t of Hous. & Cmty. Affairs v. Inclusive Communities Project*, 576 U.S. 519, 539 (2015) (holding that disparate impact claims are cognizable under the FHA); 12 C.F.R. § 1002.6(a) (ECOA regulatory codification of disparate impact); Official Staff Commentary, 12 C.F.R. pt. 1002, Supp. I, 6(a)-2 (explaining that Congress intended to apply the “effects test” to credit discrimination); HUD, DOJ, OCC, OTS, Fed. Rsrv. Bd., FDIC, FHFB, FTC, NCUA, Policy Statement on Discrimination in Lending, 59 Fed. Reg. 18266 (Apr. 15, 1994) (“Joint Policy Statement on Lending Discrimination”); *Barrett v. H&R Block, Inc.*, 652 F. Supp. 2d 104, 108 (D. Mass. 2009) (collecting cases holding that disparate impact is cognizable under ECOA).

<sup>5</sup> See, e.g., *Goodman v. Lukens Steel Co.*, 482 U.S. 656, 668–69 (1987) (explaining that animus is not required for intentional discrimination under Title VII and Section 1981), *superseded on other grounds by statute*, 28 U.S.C. § 1658(a); *Curto v. A Country Place Condominium Assoc.*, 921 F.3d 405, 410 (3d Cir. 2019) (similar for FHA claim).

policy or practice that treats people differently based on protected class status, or a close proxy for protected class status, may be unlawful without regard to whether there is a business justification.

## 2. Fair Lending Testing

Lenders have been on notice for decades that these principles, including disparate impact, apply to their lending and housing-related activities, and that federal regulatory and enforcement agencies may apply disparate impact analyses in their examinations and investigations under both the FHA and ECOA.<sup>6</sup> Accordingly, certain underlying legal principals are well-defined and, while the agencies charged with implementing these laws and regulating financial institutions have not mandated precise methodologies for fair lending testing credit models, many lenders have well-established systems for doing so. Various techniques are used, and a comprehensive survey of methodologies is beyond the scope of this Initial Report. But as a general matter, these methodologies are designed to align with traditional principles gleaned from antidiscrimination jurisprudence, tailored to the specific credit circumstances at play.<sup>7</sup> At a high level, such analyses often include: (1) ensuring that models do not include protected classes or close proxies for protected classes, for example as segmentations or variables; and (2) assessing whether a facially-neutral model is likely to disproportionately lead to negative outcomes for a protected class, and if such negative impacts exist, ensuring the model serves legitimate business needs and evaluating whether changes to the model—for example removal or substitution of variables—would result in less of a disparate effect while maintaining model performance.<sup>8</sup>

Any statistical fair lending analysis requires an awareness of consumers' likely protected class status. Age or date of birth of applicants is often available from loan files. In the mortgage context, certain protected class information is collected about applicants to satisfy legal requirements.<sup>9</sup> Outside the mortgage context, creditors usually do not have information about the race, national origin, or sex of consumers, and therefore estimation techniques are required. One

---

<sup>6</sup> Joint Policy Statement on Lending Discrimination, *supra* note 4; HUD Final rule, Implementation of the Fair Housing Act's Discriminatory Effects Standard, 78 Fed. Reg. 114460, 11476 (Feb. 15, 2013).

<sup>7</sup> There is currently a robust discussion among academics, researchers, advocates, and others on “fairness” and models, across a range of uses and markets—including contexts where disparate impact is not a basis for liability. *See, e.g.*, Deborah Hellman, “Measuring Algorithmic Fairness,” 106 Va. L. Rev. 811, 825-27 (2020) (contrasting measures of fairness). The choice of fairness standard is especially important because notions of fairness are frequently irreconcilable. *See, e.g.*, Jon Kleinberg, et al., “Inherent Trade-Offs in the Fair Determination of Risk Scores,” arXiv:1609.05807v2 [cs.LG] (Nov. 17, 2016) (demonstrating that three particular conceptions of fairness can only hold simultaneously in narrow circumstances), <https://arxiv.org/pdf/1609.05807.pdf>. While that conversation is valuable, many “fairness” proposals do not engage or align with the established three-step disparate impact analysis reflected in case law and regulatory materials. This Report describes methodologies commonly used in the financial services sector—the core of which align well with and further existing antidiscrimination jurisprudence and disparate impact principles in the lending space.

<sup>8</sup> *See* Nicholas Schmidt & Bryce Stephens, “An Introduction to Artificial Intelligence and Solutions to the Problems of Algorithmic Discrimination,” Consumer Finance Law Quarterly Report, Vol. 72, No. 2, 131, at 141–142 (2019), <https://arxiv.org/ftp/arxiv/papers/1911/1911.05755.pdf>; David Skanderson & Dubravka Ritter, Federal Reserve Bank of Philadelphia, “Fair Lending Analysis of Credit Cards” 38–40 (2014).

<sup>9</sup> *See, e.g.*, Home Mortgage Disclosure Act, 12 U.S.C. §§ 2801–2810; Regulation C, 12 C.F.R. part 1003.

common method for estimating race and national origin—which the CFPB has stated it uses—is Bayesian Improved Surname Geocoding (“BISG”), a method that relies on a combination of surnames and geography.<sup>10</sup> Sex is often estimated using first names, based on Social Security Administration data.

With limited explicit exceptions, it is a violation of the ECOA and FHA prohibitions against overt, intentional discrimination to use a protected class as a variable in a credit scoring or pricing model.<sup>11</sup> This is equally true for close proxies.<sup>12</sup> Agencies and courts have not clearly defined what qualifies as a close proxy, but it is often understood to mean a variable whose predictive value is attributable solely or largely to its correlation with a protected characteristic. A common example is local geographic location. Because of stark patterns of residential segregation, any predictive performance contributed by consideration of local geography in a model may be attributable simply to that geography’s correlation with protected status, especially at small area designations. Another example might be language preference. Language preference is highly correlated with national origin and any predictive contribution in a credit model is likely attributable to that correlation.<sup>13</sup>

Entities routinely eliminate variables that identify protected classes and variables that may be close proxies or substitutes for those classes from their models in order to mitigate risk of discriminatory treatment. In addition to qualitative reviews of whether a variable *appears* likely to serve as a proxy—which in and of itself can be relevant for assessing intentional discrimination—quantitative techniques exist for assessing whether a predictive variable may be acting as a close proxy for a protected characteristic.<sup>14</sup>

Disparate impact risk should be considered separately from whether any variable or variables act as a close proxy in a model; a model may unnecessarily cause a disparate impact

---

<sup>10</sup> See CFPB, “Using publicly available information to proxy for unidentified race and ethnicity” (Summer 2014), [https://files.consumerfinance.gov/f/201409\\_cfpb\\_report\\_proxy-methodology.pdf](https://files.consumerfinance.gov/f/201409_cfpb_report_proxy-methodology.pdf).

<sup>11</sup> 12 C.F.R. pt. 1002, Supp. I, ¶ 1002.2(p)–4 (“Besides age, no other prohibited basis may be used as a variable.”); FFIEC, “Interagency Fair Lending Examination Procedures” at 8 (Aug. 2009) (explaining that “overt discrimination” includes using “variables in a credit scoring system that constitute a basis or factor prohibited by Regulation B or, for residential loan scoring systems, the FhAct”), <https://www.ffiec.gov/pdf/fairlend.pdf>; OCC Bull. 97-24, App’x, “Safety and Soundness and Compliance Issues on Credit Scoring Models” (1997), (noting that “a creditor cannot use a credit scoring system that assigns various points based on the applicant’s race, national origin, or any other prohibited basis,” with an exception for age) <https://www.occ.treas.gov/news-issuances/bulletins/1997/bulletin-1997-24a.pdf>.

<sup>12</sup> See OCC Bull. 97-24, *supra* note 11, at 9 (“Moreover, factors linked so closely to a prohibited basis that they may actually serve as proxies for that basis cannot be used to segment the population.”).

<sup>13</sup> *Id.* (noting correlation between “those who speak Spanish as their primary language” and Hispanic national origin, and explaining that a “bank using a separate score card for persons who speak Spanish would be scrutinized by the OCC and may risk legal challenges by applicants”).

<sup>14</sup> See Skanderson & Ritter, *supra* note 8, at 37–38; Robert B. Avery, Kenneth P. Brevoort, and Glenn B. Canner, “Does Credit Scoring Produce a Disparate Impact,” Div. of Research & Statistics and Monetary Affairs, Federal Reserve Board (2010), <https://www.federalreserve.gov/pubs/feds/2010/201058/201058pap.pdf>.



and thus raise risk even if no variables in the model are close proxies.<sup>15</sup> As noted, a common approach for assessing disparate impact asks:

- (1) whether a model causes an adverse impact on a protected class;
- (2) whether the creditor has identified a legitimate business need for the model or variable; and
- (3) whether a less discriminatory alternative exists.

This disparate impact review typically starts after the entity has developed a model that has been optimized for performance and designed consistent with applicable model risk management principles, such as ensuring appropriate use, accuracy, robustness, and the like.<sup>16</sup> Credit models are also generally designed to qualify as “empirically derived, demonstrably and statistically sound, credit scoring system[s],” as that term is used in Regulation B.<sup>17</sup>

a. Disparate Impact Step 1

Once a model optimized for performance is developed, predicted outcomes of the model are reviewed to assess whether the model is likely to cause any material adverse impacts on any protected class. This is done by assessing whether each protected class disproportionately ends up with negative outcomes as compared to a control class. For example, if the model is used for underwriting, the evaluation may investigate whether there are disproportionate approval or denial rates. If the model is used for pricing, the evaluation may investigate whether there are disproportionate interest rate differences. Various metrics exist, but one common metric is the adverse impact ratio (“AIR”), which “is equal to the ratio of the proportion of the protected class that receives a favorable outcome and the proportion of the control class that receives a favorable outcome.”<sup>18</sup> AIR may be appropriate for models generating approval/denial decisions. AIR gained prominence in employment discrimination jurisprudence, and now is commonly used in other antidiscrimination scenarios as well, including financial services.<sup>19</sup> Another common metric is the standardized mean difference (“SMD”). SMD is often used to assess disparities in model outcomes in two situations. The first is when the decision being made is not discrete, but

---

<sup>15</sup> Here, we discuss analysis of purely automated decision models. Models that include some manual or judgmental decisioning include the additional risk that human discretion might introduce additional biases or discrimination.

<sup>16</sup> Particularly for supervised institutions, these principles are informed by regulatory guidance. *See, e.g.*, Bd. of Governors of the Fed. Rsrv. Sys., OCC, SR Letter 11-7, “Supervisory Guidance on Model Risk Management” (Apr. 4, 2011).

<sup>17</sup> 12 C.F.R. § 1002.2(p).

<sup>18</sup> Navdeep Gill, Patrick Hall, Kim Montgomery, and Nicholas Schmidt, “A Responsible Machine Learning Workflow with Focus on Interpretable Models, Post-hoc Explanation, and Discrimination Testing,” at 5 (2020), [https://www.bldsllc.com/publications/20200229\\_A\\_Responsible\\_Machine\\_Learning\\_Workflow.pdf](https://www.bldsllc.com/publications/20200229_A_Responsible_Machine_Learning_Workflow.pdf); *see also* 29 C.F.R. § 1607.3 (describing adverse impact test for assessing employee selection procedures under Title VII); 29 C.F.R. § 1607.16 (defining “adverse impact” as a “substantially different rate of selection in hiring, promotion or other employment decision which works to the disadvantage of” a protected class).

<sup>19</sup> The CFPB has identified other metrics for evaluating potential disparities in underwriting denial rates, including marginal effects and odds ratios. CFPB Supervisory Highlights, Issue 9, at 27 (Fall 2015), [https://files.consumerfinance.gov/f/201510\\_cfpb\\_supervisory-highlights.pdf](https://files.consumerfinance.gov/f/201510_cfpb_supervisory-highlights.pdf).

rather is a choice from a numerical range, such as an interest rate or a credit line assignment (as compared to a discrete decision, like approve/deny). The second is when the decision is based on the model output in combination with other factors.<sup>20</sup> The SMD is equal to the difference in the average protected class outcome, minus the control class outcome, divided by a measure of the standard deviation of the outcome across the overall population.<sup>21</sup>

These methods for assessing disparities align with principles from antidiscrimination case law, which instruct that the first step of the disparate impact analysis should be based on outcome disparities, comparing, for example, disparities in representation of the protected class in the total relevant pool (*e.g.*, all applicants) with their representation in the affected pool (*e.g.*, accepted applicants).<sup>22</sup> For reviews of automated models, these disparities are generally assessed without controls. That is, the disparities are measured based on the predicted outcome of the model for all protected class members compared to all control group members, without any attempt to compare outcomes of “similarly situated” protected and control group members. In other fair lending analyses—such as analyses of processes that include elements of discretion—controlling for legitimate creditworthiness criteria can help identify whether the borrowers being compared are similarly situated and hence whether disparities are being driven by legitimate factors, instead of by incidents of bias or intentional discrimination in discretionary components. In a review of a fully automated model, however, controlling for certain criteria at this stage can unnecessarily mask the disparate impact of the model and the need to search for potential less discriminatory alternatives (*e.g.*, even if a risk score is business justified, an alternative risk score may be as effective but have less disparate impact). As discussed in the next Section of this Report, traditional and commonly used credit criteria can (and often do) cause disparate impacts; controlling for these criteria can inappropriately and incorrectly assume they are justified and no less discriminatory alternatives exist.

#### b. Disparate Impact Step 2

If there are meaningful disparities adverse to a protected class, the entity should identify a legitimate business need for the model and variables. In the credit context, a model or variable is often considered to advance a valid business justification if it is predictive of a valid outcome—for example, a variable is predictive of loan performance (and its predictive relationship is not simply because it is a proxy for protected class status) or a model meets a minimum standard of accuracy for predicting default.

---

<sup>20</sup> Gill, Hall, Montgomery, and Schmidt, *supra* note 18, at 5.

<sup>21</sup> *Id.*

<sup>22</sup> See, *e.g.*, *Mt. Holly Gardens Citizens in Action, Inc. v. Twp. of Mount Holly*, 658 F.3d 375, 382 (3d Cir. 2011) (describing step 1 of the disparate impact test); *Hallmark Developers, Inc. v. Fulton Cnty., Ga.*, 466 F.3d 1276, 1286 (11th Cir. 2006); *Golden v. City of Columbus*, 404 F.3d 950, 963–64 (6th Cir. 2005); *Garcia v. Country Wide Fin. Corp.*, No. EDCV071161VAPJCRX, 2008 WL 7842104, at \*6 (C.D. Cal. Jan. 17, 2008); *Smith v. Chrysler Fin. Co., L.L.C.*, No. CIV.A. 00-6003, 2003 WL 328719, at \*5 (D.N.J. Jan. 15, 2003); *A.B. & S. Auto Serv., Inc.*, 962 F. Supp. 1056, 1060–61 (N.D. Ill. 1997); *Saldana v. Citibank, Fed. Sav. Bank*, No. 93 C 4164, 1996 WL 332451, at \*4 (N.D. Ill. June 13, 1996); *Sayers*, 522 F. Supp. 835, 839 (W.D. Mo. 1981); see also Joint Policy Statement on Lending Discrimination, *supra* note 4.

This stage might also include an assessment of whether the model or variable meets that business justification within protected class groups. Failure to perform effectively for certain protected class groups is a sign that the model or variable might be biased against a protected class or functioning as a proxy, suggesting it may not meet the legitimate business need requirement or, in certain circumstances, may raise disparate treatment concerns. In other words, whether a model predicts risk differently for different protected class groups—a form of model “bias” that is sometimes referred to as “differential validity”—can be a key indicator of fair lending risk. Those concerns are acute if a model causes disparate impact for a protected class and is also biased against that protected class. Importantly, though, confirming that a model does *not* raise bias or differential validity concerns is insufficient, standing alone, to conclude that the model does not raise disparate impact risk. As noted, under disparate impact jurisprudence, the first step in the disparate impact test is a comparison of disparities in representation of the protected class with their representation in the affected pool. Such disparities can (and often do) exist even in models that do not raise bias or differential validity concerns. A model may be reasonably accurate for different protected groups, and yet still cause disparities at step 1 of the disparate impact test. Accordingly, there may be disparate impact concerns if such disparities exist and viable less discriminatory alternatives are available, regardless of the absence of bias or differential validity in the model.

c. *Disparate Impact Step 3*

Assuming valid business justification(s), the review turns to the third step in the traditional disparate impact framework: whether less discriminatory alternatives exist. In the case of traditional statistical models, this identification process has often included a process of adding, dropping, or substituting variables in the model, with the goal of identifying variations of the model that maintain reasonable performance but that have less disparate impact on protected classes.<sup>23</sup> The increased interest in and reliance on AI/ML models has sparked development of more sophisticated methods for identifying less discriminatory alternative models.<sup>24</sup> Moreover, the availability of alternative data can offer additional avenues to explore when looking for less discriminatory alternatives to a model that has disparate impact. The foundation for many of these newer methods for identifying less discriminatory alternatives is akin to that of the traditional methods: using an awareness of the likely effects of a model on protected groups to inform a search for and development of protected-class neutral alternative models that achieve reasonable performance metrics.

Finally, the framework described here is not meant to suggest that this model testing would be sufficient, standing alone, to address fair lending risks. A complete fair lending

---

<sup>23</sup> See Schmidt & Stephens, *supra* note 8, at 141–142.

<sup>24</sup> See *id.*; Brian Hu Zhang, et al., “Mitigating Unwanted Biases with Adversarial Learning” (2018), <https://arxiv.org/abs/1801.07593>. A range of techniques for mitigating disparities are proposed in algorithmic fairness literature. Some of these proposals could raise independent fair lending risks, such as the use of different models for different protected classes or the improper use of prohibited bases as predictive variables.

program should go beyond the statistical fair lending testing sketched out here, to also include other important steps such as providing relevant personnel with appropriate fair lending training, conducting careful qualitative and quantitative reviews of credit policies and procedures, analyzing the application and use of the specific model, and maintaining robust compliance and monitoring functions. This Report focuses on statistical model testing since that forms the basis of this Monitorship.

### 3. Credit Disparities and the Rise of Alternative Data and Alternative Algorithms

#### a. *Disparities in Access to Credit*

There is a long history of credit discrimination in the United States, which underlies current disparities in wealth, rates of home and small business ownership, and current access to credit. That history provides the backdrop against which this Monitorship occurs. Mutually-reinforcing systems of discriminatory legal and public policy frameworks and private-sector discrimination have disproportionately excluded communities of color, and Black communities in particular, from access to affordable credit.<sup>25</sup> The inability to obtain mainstream credit has, among other things, limited access to the primary engines of wealth generation for white Americans, such as home- and small business-ownership.<sup>26</sup> Indeed, the net worth of a typical white family today is nearly ten times greater than that of a Black family.<sup>27</sup> These wealth disparities combine with racial segregation and credit discrimination to perpetuate inequalities that further undermine the economic well-being of communities of color.<sup>28</sup>

---

<sup>25</sup> See generally Mehrsa Baradaran, *The Color of Money: Black Banks and the Racial Wealth Gap* (2014); Keeanga-Yamahtta Taylor, *Race for Profit: How Banks and the Real Estate Industry Undermined Black Homeownership* (2019).

<sup>26</sup> Jung Hyun Choi, et al., Urban Institute, *Explaining the Black-White Homeownership Gap* at v (2019), [https://www.urban.org/sites/default/files/publication/101160/explaining\\_the\\_black-white\\_homeownership\\_gap\\_a\\_closer\\_look\\_at\\_disparities\\_across\\_local\\_markets\\_0.pdf](https://www.urban.org/sites/default/files/publication/101160/explaining_the_black-white_homeownership_gap_a_closer_look_at_disparities_across_local_markets_0.pdf); Alicia Robb, U.S. Small Business Administration, *Financing Patterns and Credit Market Experiences: A Comparison by Race and Ethnicity for U.S. Employer Firms* at 3–5 (Feb. 2018),

[https://www.sba.gov/sites/default/files/Financing\\_Patterns\\_and\\_Credit\\_Market\\_Experiences\\_report.pdf](https://www.sba.gov/sites/default/files/Financing_Patterns_and_Credit_Market_Experiences_report.pdf).

<sup>27</sup> Kriston McIntosh, et al., *Examining the Black-white wealth gap*, Brookings, (Feb. 27, 2020), <https://www.brookings.edu/blog/up-front/2020/02/27/examining-the-black-white-wealth-gap/>; see also Urban Institute, *Nine Charts about Wealth Inequality in America*, (Oct. 5, 2017), <https://apps.urban.org/features/wealth-inequality-charts/>; Dionissi Alilprantis & Daniel R. Carroll, *What is Behind the Persistence of the Racial Wealth Gap?*, Federal Reserve Bank of Cleveland (Feb. 28, 2019), <https://www.clevelandfed.org/newsroom-and-events/publications/economic-commentary/2019-economic-commentaries/ec-201903-what-is-behind-the-persistence-of-the-racial-wealth-gap.aspx>.

<sup>28</sup> Baradaran, *The Color of Money*, *supra* note 25, at 90–96 (describing how racial segregation, wealth disparities, and credit discrimination combine to prevent the generation of wealth in Black communities); see also Jacob S. Rugh & Douglas S. Massey, *Racial Segregation and the American Foreclosure Crises*, 75 Am. Soc. Rev. 629 (Oct. 2010) (describing the racialized process leading up to the subprime lending and foreclosure crises), <https://www.asanet.org/sites/default/files/savvy/images/journals/docs/pdf/asr/Oct10ASRfeature.pdf>.

b. The Thin File Challenge

Credit history is often a primary determinant for whether consumers will be offered credit. Consumers without such history—those that are “credit invisible” or “thin file”—are commonly denied, which limits their ability to develop a history. This phenomenon is the credit “Catch-22,” and it affects a significant number of American consumers. In 2017, the CFPB “estimate[d] that 26 million Americans are ‘credit invisible,’ meaning they have no credit history at all,” and “another 19 million people have credit histories that are too limited or have been inactive for too long to generate a credit score” under the most widely-used credit scoring models.<sup>29</sup> The issue is pressing because it is hard to build a credit history without access to credit in the first place; without access to affordable credit, individuals remain more financially vulnerable and less able to improve their financial circumstances.<sup>30</sup>

Disparities exist both with respect to the populations likely to have credit histories, and with respect to how histories are acquired. There is a high correlation between income and having a credit record: “[a]most 30 percent of consumers in low-income neighborhoods are credit invisible and an additional 15 percent have unscored records. These percentages are notably lower in higher-income neighborhoods.”<sup>31</sup> Black and Hispanic Americans are also more likely than white or Asian Americans to be credit invisible or to have unscored records.<sup>32</sup> Moreover, traditional techniques for building strong credit files—for example, “[t]he use of co-borrowers and authorized user account status [are] notably less common in lower-income neighborhoods.”<sup>33</sup> Consumers in lower-income neighborhoods are more likely to acquire credit records from “non-loan items,” like third-party collection accounts or public records.<sup>34</sup>

Many individuals with thin files or who are credit invisible may in fact be good credit risks. The challenge for lenders is how to make underwriting and loan pricing decisions for consumers who do not already have an established credit history. Over the past several years, researchers, lenders, policymakers, and advocates have explored and debated whether two developments in credit underwriting—“alternative data” and “alternative algorithms”—can help address this challenge.

The issue of thin files and credit invisibility is far from the only issue driving continued racial disparities in access to high-quality credit. Indeed, among those consumers who *do* have enough credit history to generate a credit score, traditional credit scores demonstrate a strong and

---

<sup>29</sup> Cordray, *supra* note 1; see also Kenneth P. Brevoort, et al., CFPB Office of Research, “Data Point: Credit Invisibles,” at 12 (May 2015), [https://files.consumerfinance.gov/f/201505\\_cfpb\\_data-point-credit-invisibles.pdf](https://files.consumerfinance.gov/f/201505_cfpb_data-point-credit-invisibles.pdf).

<sup>30</sup> Cordray, *supra* note 1.

<sup>31</sup> CFPB “Data Point: Credit Invisibles,” *supra* note 29, at 6.

<sup>32</sup> *Id.*

<sup>33</sup> Kenneth P. Brevoort & Michelle Kambara, CFPB Office of Research, “Data Point: Becoming Credit Visible” at 6 (June 2017), [https://files.consumerfinance.gov/f/documents/BecomingCreditVisible\\_Data\\_Point\\_Final.pdf](https://files.consumerfinance.gov/f/documents/BecomingCreditVisible_Data_Point_Final.pdf).

<sup>34</sup> *Id.* at 11–12.

persistent disparate impact against people of color.<sup>35</sup> Nevertheless, the racial disparities in credit invisibility are one important piece of the puzzle.

c. What are Alternative Data?

“Alternative data” generally describes information that lenders may use for credit decisions but that is “not typically found in the consumer’s credit files of the nationwide consumer reporting agencies or customarily provided by consumers as part of applications for credit.”<sup>36</sup> This definition is imperfect, in part because the types of information typically found in the nationwide consumer reporting agency files are a moving target. Regardless, examples of alternative data include: data regarding payments for utilities, rent, or other recurring payments; cash flow data regarding deposit accounts; and data regarding education, such as information about an individual’s school or degrees.<sup>37</sup>

Regulators and others have argued that, by making underwriting feasible in a broader set of circumstances, alternative data have the potential to expand access to credit.<sup>38</sup> An interagency statement from the Federal Reserve Board, CFPB, FDIC, NCUA, and OCC observed that the “agencies recognize that use of alternative data may improve the speed and accuracy of credit decisions and may help firms evaluate the creditworthiness of consumers who currently may not obtain credit in the mainstream credit system.”<sup>39</sup> “[E]nhanced assessments of repayment capacity,” may, in turn, “enable consumers to obtain additional products and/or more favorable pricing/terms.”<sup>40</sup> Accordingly, market participants, including lenders, consumer reporting agencies, and other market participants, have been particularly focused on leveraging alternative data in ways meant to augment traditional credit information.<sup>41</sup>

At the same time, the use of alternative data raises potential risks, including the introduction of “[i]naccurate or incomplete information,” the creation of new hurdles to improving credit standing, and “the potential for discrimination.”<sup>42</sup>

---

<sup>35</sup> Lisa Rice & Deidre Swesnik, *Discriminatory Effects of Credit Scoring on Communities of Color*, 46 *Suffolk U. L. Rev.* 935, 952-959 (2013); Aaron Klein, Brookings Institution, “Reducing bias in AI-based financial services” (July 10, 2020), <https://www.brookings.edu/research/reducing-bias-in-ai-based-financial-services>.

<sup>36</sup> See Bd. of Governors of the Fed. Rsrv. Sys., CFPB, FDIC, NCUA, OCC, “Interagency Statement on the Use of Alternative Data in Credit Underwriting,” at 1 n.1 (Jan. 2019), <https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20191203b1.pdf>.

<sup>37</sup> See, e.g., CFPB, Request for Information Regarding Use of Alternative Data and Modeling Techniques in the Credit Process, 82 *Fed. Reg.* 11183, 11185 (Feb. 21, 2017).

<sup>38</sup> Cordray, *supra* note 1; see also CFPB, “CFPB Explores Impact of Alternative Data on Credit Access for Consumers Who Are Credit Invisible” (Feb. 16, 2017), <https://www.consumerfinance.gov/about-us/newsroom/cfpb-explores-impact-alternative-data-credit-access-consumers-who-are-credit-invisible/>.

<sup>39</sup> See Interagency Statement on Alternative Data, *supra* note 36, at 1 n.1.

<sup>40</sup> *Id.*

<sup>41</sup> See, e.g., FinRegLab, “The Use of Cash-Flow Data in Underwriting Credit” at 8 (July 2019), [https://finreglab.org/wp-content/uploads/2019/07/FRL\\_Research-Report\\_Final.pdf](https://finreglab.org/wp-content/uploads/2019/07/FRL_Research-Report_Final.pdf).

<sup>42</sup> Brian Kreiswirth, et al., CFPB Blog, “Using alternative data to evaluate creditworthiness,” (Feb. 16, 2017), <https://www.consumerfinance.gov/about-us/blog/using-alternative-data-evaluate-creditworthiness/>; see also Solon Barocas and Andrew Selbst, “Big Data’s Disparate Impact,” 104 *Cal. L. Rev.* 671 (June 2016).

d. What are Alternative Algorithms?

Alternative algorithms (e.g., ML, Semi-Parametric Methods, Symbolic Regression, Evolutionary Search methods, etc.) in contrast to traditional statistical modeling, refers to a set of methods for constructing predictive models and optimization models from historical data (also referred to as “training data”) without requiring human modelers to explicitly specify relationships among the variables that can be used in the model.<sup>43</sup> Such alternative algorithms can automatically encode complex interactions among input variables and the outcome being predicted.<sup>44</sup>

Commenters have argued that the use of alternative algorithms (and, in particular ML algorithms), like alternative data, “has the potential to expand credit access by enabling lenders to evaluate the creditworthiness of some of the millions of consumers who are unscorable using traditional underwriting techniques.”<sup>45</sup> Here too, a robust debate exists regarding the risks and benefits of using alternative algorithms, including risks regarding explainability, interpretability, model drift, privacy, and discrimination.<sup>46</sup>

e. Fair Lending Implications

As noted, disparities and barriers to access along protected class lines are endemic to existing credit markets. Relevant here, studies have found significant racial disparities in credit scoring, and there is evidence that factors on which traditional scores rely may result in disparities in scores.<sup>47</sup> Some argue that existing credit-scoring systems have a disparate impact on communities of color because they are rooted in our Nation’s history of intentional, and often state-sponsored, discrimination.<sup>48</sup> Importantly, then, fair lending and disparate impact risks exist with respect to traditional data and models; the use of alternative data and models operates against a landscape that produces disparate outcomes, and that landscape could be improved. A

---

<sup>43</sup> See Schmidt & Stephens, *supra* note 8, at 133.

<sup>44</sup> *Id.* at 134

<sup>45</sup> Patrice Alexander Ficklin, et al., CFPB Blog, “Innovation spotlight: Providing adverse action notices when using AI/ML models” (July 7, 2020), <https://www.consumerfinance.gov/about-us/blog/innovation-spotlight-providing-adverse-action-notices-when-using-ai-ml-models/>; Fed. Rsr. Bd. Governor Lael Brainard, *Supporting Responsible Use of AI and Equitable Outcomes in Financial Services*, Speech at the AI Academic Symposium hosted by the Bd. of Governors of the Fed. Rsr. Sys. (Jan. 21, 2021), <https://www.federalreserve.gov/newsevents/speech/brainard20210112a.htm>.

<sup>46</sup> See, e.g., FinRegLab, “AI in Financial Services: Key Concepts” (May 2020), [https://finreglab.org/wp-content/uploads/2020/05/FinRegLab\\_AIFAQ\\_Key-Concepts\\_AI-in-Financial-Services.pdf](https://finreglab.org/wp-content/uploads/2020/05/FinRegLab_AIFAQ_Key-Concepts_AI-in-Financial-Services.pdf); Caroline Wang et al., “In Pursuit of Interpretable, Fair and Accurate Machine Learning for Criminal Recidivism Prediction” (May 2020), <https://arxiv.org/abs/2005.04176>.

<sup>47</sup> See, e.g., Rice & Swesnik, *supra* note 35, at 936; see Klein, *supra* note 35; Bd. of Governors of the Fed. Reserve Sys., Report to the Congress on Credit Scoring and Its Effects on the Availability and Affordability of Credit (2007); CFPB, “Analysis of Differences Between Consumer- and Creditor-Purchased Credit Scores (Sept. 2012); National Consumer Law Center, Credit Discrimination § 6.4.1.1, “Studies Showing that Minorities Have Lower Credit Scores as a Group” (7th ed. 2018).

<sup>48</sup> See, e.g., Rice & Swesnik, *supra* note 35.

corresponding point is that fair lending testing of credit models should not treat traditional models as grandfathered benchmarks for fairness, immunizing models using alternative data or algorithms from scrutiny simply if the disparities they produce are no worse than those caused by traditional models or variables. To do so would assume away the fair lending implications of traditional models, while failing to properly evaluate the impact of non-traditional approaches.

That said, what effects alternative data and AI/ML will have on discrimination risks cannot be answered in the abstract; it is a context-specific and largely quantitative question. Because alternative data can be used to underwrite and price consumers with a thin file, or no credit history at all, these data have the potential to improve both fair lending and business outcomes. Whether this possibility bears out, however, depends on criteria such as the type of alternative data at issue, how those data are deployed, the type of alternative algorithm used, and characteristics of the likely applicant pool. And many have noted the potential for risks and pitfalls.<sup>49</sup> For example, the use of biased or unrepresentative data could undermine the potential of alternative data to improve fair lending outcomes.<sup>50</sup> Either over- or under-representation of protected class members in a dataset may result in adverse effects.<sup>51</sup> Similarly, if the data on which a model is trained reflect existing discriminatory patterns or biases, the model may perpetuate similar problems.<sup>52</sup> It can also be “difficult to understand how learning algorithms reach the results they do,”<sup>53</sup> including how AI/ML models process variables, which adds to concerns that they may rely on or contribute to protected class disparities in subtle ways, or that they may otherwise unnecessarily perpetuate disparate impacts.<sup>54</sup>

Similarly, fair lending concerns are not resolved solely because a practice increases access to credit. Alternative algorithms and consideration of alternative data elements could result in additional approvals for all groups, but benefit non-Hispanic white men to a greater extent than other protected class members, thus exacerbating disparities. This is a key point:

---

<sup>49</sup> See, e.g., Governor Brainard Speech, *supra* note 45 (“To harness the promise of machine learning to expand access to credit, especially to underserved consumers and businesses that may lack traditional credit histories, it is important to be keenly alert to potential risks around bias and inequitable outcomes.”).

<sup>50</sup> See generally FTC, “Big Data: A Tool for Inclusion or Exclusion? Understanding the Issues” (Jan. 2016), <https://www.ftc.gov/reports/big-data-tool-inclusion-or-exclusion-understanding-issues-ftc-report>; see also Barocas & Selbst, *supra* note 42, at 680-81.

<sup>51</sup> Barocas & Selbst, *supra* note 42, at 686 (“If a sample includes a disproportionate representation of a particular class (more or less than its actual incidence in the overall population), the results of an analysis of that sample may skew in favor of or against the over- or underrepresented class.”).

<sup>52</sup> Carol A. Evans, Federal Reserve, “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>; Barocas & Selbst, *supra* note 42, at 674.

<sup>53</sup> See, e.g., Cary Coglianese & David Lehr, “Transparency and Algorithmic Governance,” 71 Admin. L. Rev. 1, 4 (2019); Cary Coglianese & David Lehr, “Regulating by Robot: Administrative Decision Making in the Machine-Learning Era,” 105 Geo. L.J. 1147, 1159 (2017).

<sup>54</sup> Evans, *supra* note 52; FTC Big Data, *supra* note 50; Exec. Off. of the President, “Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights” (2016), [https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016\\_0504\\_data\\_discrimination.pdf](https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf); Mark MacCarthy, “Standards of Fairness for Disparate Impact Assessment of Big Data Algorithms,” 48 Cum. L. Rev. 67, 75-76 (2018).



increases in access to even safe or favorable credit or financial products can drive persistent inequality—and disparate impact—if distributed unequally.<sup>55</sup>

A perceived risk central to this monitoring project is the concern that use of certain information related to higher education may contribute to discriminatory outcomes that disproportionately affect communities of color, including students who attend minority-serving institutions, such as Historically Black Colleges and Universities (“HBCUs”) and Hispanic Serving Institutions (“HSIs”). Educational variables, particularly those that are gathered at the cohort level rather than individually, have been singled out by some regulators and advocates for criticism based on a concern that these variables reflect and could perpetuate societal disparities in educational access. For example, the use of cohort default rate (“CDR”), which measures the average loan default rate for students at a particular school, has been widely seen as problematic from a fair lending perspective. As far back as 2007, Andrew Cuomo, then the New York Attorney General, compared the use of CDRs in lending to mortgage redlining, stating, “[j]ust as lenders in the mortgage industry once made judgments about credit lending in entire neighborhoods as a whole, so too are lenders making generalized judgments about student and parent credit risk based on a student’s ‘school neighborhood.’”<sup>56</sup> The CFPB has publicly criticized the use of CDRs in lending as a fair lending concern due to the potential for disparate impact on minority students, and the FDIC entered into a consent order with Sallie Mae in 2014 prohibiting it from continuing to use CDRs, finding it to be an ECOA violation.<sup>57</sup> Other cohort-level educational data, such as school attended or major/program of study, have also been noted as possible sources of fair lending risk, given the race, ethnicity, and gender differences between schools, programs, and majors.<sup>58</sup>

---

<sup>55</sup> As an extreme example, the government programs adopted during the New Deal, such as the Home Owners’ Loan Corporation and the Federal Housing Administration, significantly expanded home buying opportunities for a large number of American households. These programs dramatically reshaped homeownership and wealth opportunities, but infamously did so in an explicitly discriminatory manner, laying the foundation for today’s racial wealth gaps and existing segregation. Although those programs provided some limited number of minority households opportunities that would not have otherwise been available, the consequences of the disparities generated by those programs—with the lion’s share of benefits allocated to white households—dwarf those few opportunities for minority households. *See, e.g.,* Jacob W. Faber, “We Built This: Consequences of New Deal Era Intervention in America’s Racial Geography,” *American Sociological Review* 85(5), 740 (2020) <https://journals.sagepub.com/doi/pdf/10.1177/0003122420948464>; Daniel Aaronson, “The Long-run effects of the 1930s HOLC ‘redlining’ maps on place-based measures of economic opportunity and socioeconomic success” (Jan. 2021), [https://www.sciencedirect.com/science/article/abs/pii/S0166046220303070?dgcid=rss\\_sd\\_all](https://www.sciencedirect.com/science/article/abs/pii/S0166046220303070?dgcid=rss_sd_all).

<sup>56</sup> George M. Walsh, *Cuomo: Lenders ‘redlining’ student loans*, Associated Press (June 19, 2007), [http://archive.boston.com/news/education/higher/articles/2007/06/19/cuomo\\_lenders\\_redlining\\_student\\_loans/](http://archive.boston.com/news/education/higher/articles/2007/06/19/cuomo_lenders_redlining_student_loans/).

<sup>57</sup> CFPB, *Private Student Loans Report* (2012), [https://files.consumerfinance.gov/f/201207\\_cfpb\\_Reports\\_Private-Student-Loans.pdf](https://files.consumerfinance.gov/f/201207_cfpb_Reports_Private-Student-Loans.pdf); *see also* CFPB, *Fair Lending Report* (2012), [https://files.consumerfinance.gov/f/201212\\_cfpb\\_fair-lending-report.pdf](https://files.consumerfinance.gov/f/201212_cfpb_fair-lending-report.pdf); Sallie Mae Bank, Fed. Deposit Ins. Corp. 13-0366b (2014) (Consent Order, Order for Restitution, and Order to Pay Civil Money Penalty), <https://www.fdic.gov/news/news/press/2014/salliemae.pdf>.

<sup>58</sup> *See, e.g.,* Dowse B. (Brad) Rustin IV, et al., *Pricing Without Discrimination: Alternative Student Loan Pricing, Income Share Agreements, and the Equal Credit Opportunity Act*, American Enterprise Institute (Feb. 2017) at 10, <https://www.aei.org/wp-content/uploads/2017/02/Pricing-Without-Discrimination.pdf>.

If educational data were used in a manner that disadvantaged students at and graduates of HBCUs and HSIs, it would be of particular concern given the significance of those institutions—both historically and contemporarily—to communities of color. HBCUs, for example, have historically played, and continue to play, an outsized role in economic mobility and the cultural identity of Black communities.<sup>59</sup> Despite representing just three percent of four-year colleges and universities in the United States, HBCUs graduate twenty-seven percent of the nation’s Black undergraduate STEM majors and roughly a fourth of Black education majors.<sup>60</sup> HBCU alumni further account for roughly *eighty* percent of Black judges and fifty percent of Black lawyers and doctors.<sup>61</sup> And they do so despite long histories—continuing to the present day—of undercapitalized endowments, fewer resources, and in the case of public HBCUs, disinvestment and discrimination by state governments.<sup>62</sup>

## **B. Upstart’s Lending Model and Concerns Leading to this Monitoring Engagement**

### **1. Development of Upstart’s Alternative Data Model**

Upstart initially introduced its lending platform in 2014, stating that its goal was to use non-traditional variables and advanced modeling techniques to more effectively underwrite and approve borrowers, particularly those with short credit histories.<sup>63</sup>

Upstart’s minimum credit characteristics included, for applicants without a credit score, the option to qualify based on educational enrollment or attainment such as enrollment in or graduation from an accredited associate’s degree or bachelor’s degree program.<sup>64</sup> In addition, Upstart’s underwriting and pricing model incorporated educational and employment variables, such as school attended, degree attained, and current employment, along with more traditional credit variables.<sup>65</sup> These variables were used to train an AI model via a set of ML algorithms that evaluated these variables individually and in combination with Upstart’s loan performance data to predict an applicant’s credit risk.<sup>66</sup>

Initially, all Upstart loans were originated by one financial institution, Cross River Bank, with applications coming through Upstart’s consumer-facing website and decided using Upstart’s underwriting and pricing model, with the underlying loans made available for

---

<sup>59</sup> See, e.g., Arne Duncan, U.S. Sec’y of Educ., Opening Remarks at the 2014 National HBCU Conference (Sept. 23, 2014), <https://www.ed.gov/news/speeches/hbcu-value-proposition>.

<sup>60</sup> *Id.*

<sup>61</sup> Adam Harris, “Why America Needs Its HBCUs,” *The Atlantic* (May 16, 2019), <https://www.theatlantic.com/education/archive/2019/05/howard-universitys-president-why-america-needs-hbcus/589582/>.

<sup>62</sup> See, e.g., *id.*; Fourth Amended Complaint, *The Coalition for Equity and Excellence in Maryland Higher Education, Inc. v. Maryland Higher Education Commission*, Civil No. MJG-06-2773 (March 30, 2010), <https://lawyerscommittee.org/wp-content/uploads/2015/06/Docket-165-Fourth-Amended-Complaint.pdf>.

<sup>63</sup> Upstart’s 2017 NAL application, 1, 3 (hereinafter “2017 NAL Request”).

<sup>64</sup> *Id.* at 1–2.

<sup>65</sup> *Id.* at 4.

<sup>66</sup> *Id.*

investment to accredited investors.<sup>67</sup> Since then, the business has expanded and Upstart now also partners with a number of other banks and financial institutions to provide the underwriting and backend services for loan applications coming through the bank partner’s own web portals.<sup>68</sup>

Upstart represents that its technology helps expand access to credit beyond traditional credit scoring techniques. According to Upstart, as of June 2020, 45.5% of Upstart borrowers were low- and moderate-income individuals.<sup>69</sup> Upstart also represents that, as of Q4 2019, about 36% of Upstart borrowers had no college degree and less than 5% of Upstart borrowers had a credit score that was greater than 750. As explained in the next section, a CFPB blog post from August 2019 discusses additional Upstart representations regarding reported increases in access to affordable credit for certain populations.

## 2. Upstart’s 2017 No Action Letter and Compliance Plan

In 2017, Upstart submitted an application to the CFPB under the Bureau’s No Action Letter (“NAL”) program. The NAL program, finalized in a policy statement issued in February 2016, permits institutions to apply to the Bureau for a non-binding letter expressing that the agency has no present intention to bring enforcement or supervisory actions against the institution related to the product or service in question. The program was instituted to facilitate innovative financial products or services that are likely to lead to consumer benefits but might otherwise be hindered by legal or regulatory uncertainty.<sup>70</sup>

At the time of Upstart’s NAL application, it indicated that it had originated over 80,000 loans totaling over \$1 billion through its platform, with interest rates ranging from 4% to 25.9% and an average loan amount of \$12,000.<sup>71</sup> In its application materials, Upstart represented that it could expand credit access and provide better interest rates to “traditionally underserved borrowers,” such as those with shorter credit histories, including younger borrowers and recent immigrants.<sup>72</sup>

Upstart’s application described testing it had done to compare outcomes for various groups of consumers under Upstart’s pricing model to outcomes those borrowers could expect under a hypothetical model it designed to imitate more traditional models using typical credit history-based variables.<sup>73</sup> It reported that using this hypothetical model on a sample of borrowers who had received Upstart loans resulted in an average APR of 23.4%, compared to the average APR indicated by Upstart’s own model of 16.7%.<sup>74</sup> The testing results included in the public application compared how applicants in general fared under the Upstart model as compared to

---

<sup>67</sup> *Id.*

<sup>68</sup> Upstart Website (last visited Feb. 25, 2021), <https://www.upstart.com/for-banks/>.

<sup>69</sup> Upstart Blog, “Upstart by the numbers,” <https://www.upstart.com/blog/upstart-by-the-numbers>.

<sup>70</sup> CFPB, Policy on No-Action Letters; Information Collection, 81 Fed. Reg. 8686, 8686 (Feb. 22, 2016).

<sup>71</sup> 2017 NAL Request at 4.

<sup>72</sup> *Id.* at 6.

<sup>73</sup> *Id.* at 7–8.

<sup>74</sup> *Id.* at 8.

the traditional model, but did not include any assessments distinguishing between protected class and reference group members or of any alternative models. The public application reported that Upstart had conducted disparate impact testing, which it believed demonstrated that its underwriting and pricing methodology had not resulted in any unlawful disparate impact.<sup>75</sup> Results of that testing were provided confidentially to the Bureau but were not disclosed in the public application materials.<sup>76</sup>

As part of the initial NAL application, Upstart committed to implementing a compliance plan and providing the Bureau with results of its fair lending testing and providing additional information, including underlying data, upon request.<sup>77</sup> The compliance plan, which was the result of an agreement reached with Bureau staff, was not made public.

The CFPB granted Upstart's NAL request on September 14, 2017, indicating that the Bureau had "no present intention" to recommend supervisory or enforcement fair lending action against Upstart related to its automated underwriting model.<sup>78</sup> The NAL was effective for a three-year period.

In August 2019, the CFPB issued an update regarding information about the testing results Upstart had conducted as part of its compliance plan.<sup>79</sup> The Bureau shared a summary of results provided by Upstart of Upstart's access to credit testing, although the Bureau also noted that Upstart's fair lending testing results "show no disparities that require further fair lending analysis under the compliance plan."<sup>80</sup> The update stated that Upstart's access to credit analyses had shown that Upstart's model approved 27% more applicants than the traditional model, and resulted in 16% lower APRs for approved loans. It also noted that those access to credit analyses had shown similar results "across all tested race, ethnicity, and sex segments," and that the increase in approval rates varied between 23% and 29% for these segments, while the decrease in average APR varied between 15% and 17%.<sup>81</sup> According to Upstart's results, "[n]ear prime' consumers with FICO scores from 620 to 660 [were] approved approximately twice as frequently."<sup>82</sup> The CFPB did not separately replicate these analyses.<sup>83</sup>

---

<sup>75</sup> *Id.* at 9.

<sup>76</sup> *Id.*

<sup>77</sup> *Id.* at 12.

<sup>78</sup> CFPB, Letter Response to 2017 NAL Request (Sept. 14, 2017),

[https://files.consumerfinance.gov/f/documents/201709\\_cfpb\\_upstart-no-action-letter.pdf](https://files.consumerfinance.gov/f/documents/201709_cfpb_upstart-no-action-letter.pdf).

<sup>79</sup> Patrice Alexander Ficklin & Paul Watkins, "An update on credit access and the Bureau's first No-Action Letter" (Aug. 6, 2019), <https://www.consumerfinance.gov/about-us/blog/update-credit-access-and-no-action-letter/>.

<sup>80</sup> *Id.*

<sup>81</sup> *Id.*

<sup>82</sup> *Id.*

<sup>83</sup> *Id.*

### 3. SBPC’s “Educational Redlining” Report and Congressional Responses

In February 2020, the SBPC released a report, entitled “Educational Redlining,” examining fair lending issues related to student lending and the refinancing of educational loans.<sup>84</sup> One section of that report focused on SBPC’s concerns that the use of educational data by fintech lenders such as Upstart might disadvantage students and graduates of HBCUs and HSIs.

As part of that inquiry, SBPC presented a case study that explored Upstart’s use of educational institution information. To conduct this case study, SBPC submitted multiple rate requests using Upstart’s online rate quote tool. Each request was submitted on behalf of a 24-year-old New York City resident with a bachelor’s degree, but the rate request submissions varied the college or university from which the applicant had graduated while keeping other inputs such as employment status, income, and savings constant.<sup>85</sup> The report presented the results of three of these requests, representing three different educational institutions: New York University; Howard University, an HBCU; and New Mexico State University, an HSI. In each case the applicant requested a quote for a \$30,000 loan to refinance existing student loans.

The report found that the rates and origination fees returned by the Upstart rate quote tool differed between these three schools, with the NYU graduate paying \$1,724 less than the New Mexico State graduate, and \$3,499 less than the Howard graduate, over the life of the loan.<sup>86</sup> SBPC indicated that although it was highlighting these three institutions, the “findings were consistent across hypotheticals.”<sup>87</sup> Based on these results, SBPC concluded that, according to Upstart’s own data, otherwise similarly-situated applicants who had attended these minority-serving institutions were charged more because of their choice of school.<sup>88</sup>

Upstart publicly responded to the SBPC report. First, Upstart noted that the rate quotes, of which there were 26 in total, were submitted using the same individual’s credit report over a two-and-a-half-month period.<sup>89</sup> Upstart’s response stated that both the Upstart model and the individual’s credit report had changed during this time, and that these changes were responsible for about half of the pricing differences noted in the SBPC study. Second, it indicated that the hypothetical Howard graduate had received a better rate than approximately half of the 26 hypothetical applicants.<sup>90</sup> The response also stated that the Upstart model contains over 1,500 variables or factors, such that in the real world no two applicants would match on so many variables.

---

<sup>84</sup> SBPC, “Educational Redlining” at 15 (Feb. 2020), <https://protectborrowers.org/wp-content/uploads/2020/02/Education-Redlining-Report.pdf>.

<sup>85</sup> *Id.*

<sup>86</sup> *Id.* at 18.

<sup>87</sup> *Id.* at 16.

<sup>88</sup> *Id.* at 17.

<sup>89</sup> Upstart, Response to SBPC report (on file with Monitor).

<sup>90</sup> The response did not provide any further information about how the hypothetical New Mexico State University graduate fared in comparison to the other hypothetical applicants tested.

The SBPC, for its part, contends that any relevant changes in credit score did not take place during the period applicable to the report, and do not change the nature of its findings.<sup>91</sup>

Shortly after the SBPC study was released, Democratic members of the Senate Committee on Banking, Housing, and Urban Affairs sent letters to Upstart and other lenders and service providers utilizing education-related data for lending purposes.<sup>92</sup> The letters expressed concern about the use of non-individualized education factors, which they noted the CFPB, the FDIC, and the New York Attorney General’s Office had all found could, in at least some circumstances, result in ECOA violations.

In its response to the Senate letter, Upstart indicated that its model did not incorporate the specific educational institution attended, but rather grouped schools into categories, at that time based on the institution’s average incoming standardized test scores (such as SAT or ACT), and incorporated the category as a variable.<sup>93</sup> The letter also noted that traditional credit scores disproportionately classify Black and Hispanic borrowers into the lowest deciles of credit scores, and argued that Upstart’s model helps address biases in traditional lending models by expanding access to credit.<sup>94</sup>

In July 2020, Democratic Committee members released a report assessing the responses received from Upstart and the other industry participants.<sup>95</sup> In particular, the report expressed concern with practices identified in the industry responses that it concluded “create significant risk of unlawful discrimination in violation of federal fair lending laws.”<sup>96</sup> The report specifically identified Upstart’s use of the school an applicant had attended and found problematic the use of non-individualized cohort-level educational data. Committee members forwarded this report to the CFPB along with a request that the CFPB take action to address the concerns.

---

<sup>91</sup> The Monitor does not take any position on the SBPC’s study or the parties’ differing conclusions about its findings. The methodology used in the SBPC report is not the same fair lending methodology that will be used by the Monitor when conducting fair lending testing with more complete access to Upstart’s model and underlying data.

<sup>92</sup> Letter from Senators Sherrod Brown, Elizabeth Warren, Robert Menendez, Cory Booker, and Kamala D. Harris to Dave Girouard (Feb. 13, 2020), <https://www.brown.senate.gov/imo/media/doc/2020-02-13%20Senate%20letter%20to%20Upstart.pdf>.

<sup>93</sup> Response to Letter Dated February 13, 2020 (Feb. 28, 2020), <https://www.banking.senate.gov/imo/media/doc/Educational%20data%20-%20appendix.pdf>.

<sup>94</sup> *Id.*

<sup>95</sup> “Use of Educational Data to Make Credit Determinations,” <https://www.banking.senate.gov/imo/media/doc/Review%20-%20Use%20of%20Educational%20Data.pdf>.

<sup>96</sup> Letter from Senators Sherrod Brown, Elizabeth Warren, and Kamala D. Harris to Kathleen Kraninger, Director, CFPB at 1 (July 30, 2020), [https://www.banking.senate.gov/imo/media/doc/2020-07-30\\_Letter%20to%20CFPB%20re%20use%20of%20educational%20data.pdf](https://www.banking.senate.gov/imo/media/doc/2020-07-30_Letter%20to%20CFPB%20re%20use%20of%20educational%20data.pdf).

### C. LDF & SBPC Engagement with Upstart

Around this time, Upstart, the SBPC, and LDF engaged in conversations regarding SBPC and LDF's concerns about Upstart's use of cohort-level educational data, which ultimately led to this monitoring engagement.<sup>97</sup> In July 2020, the two organizations wrote to Upstart stating that their further investigation, including a review of the responses to the congressional inquiry, had indicated that "the way in which Upstart uses education data in determining an applicant's creditworthiness causes an unjustified disparate impact on Black borrowers."<sup>98</sup> In particular, the organizations expressed concern with Upstart's practice—described in its responses to the Senators' request for information—of sorting roughly 2,700 schools into eight tranches "primarily" based on average standardized test results of incoming students, and then using the tranche to which an applicant's school belonged as a variable in its underwriting model.<sup>99</sup> The letter noted that "a wealth of social science research reveals racial disparities in scores on standardized tests" that are unrelated to any differences in academic merit—and that test scores have virtually no relationship with outcome measures such as grade point average, graduation rate, or postgraduate performance.<sup>100</sup> Given these racial disparities, the organizations noted their strong concern that Upstart's use of cohort-level standardized test data imported the disparate impact of these tests into the lending context—a concern that appeared to correlate with Upstart's statements to the Senators confirming that over 95% of HBCUs fell into the bottom half of Upstart's rankings.<sup>101</sup>

### D. Changes to the Upstart Underwriting Model

In responses to those conversations, as well as the congressional inquiry, Upstart made certain changes to how its underwriting model utilized educational data. Most notably, it abolished the use of average incoming SAT and ACT scores to group education institutions in its underwriting model. While Upstart's model continues to incorporate information about the educational institution attended, it switched to grouping schools based on average post-graduation income.<sup>102</sup>

---

<sup>97</sup> See, e.g., Letter from Ashok Chandran and Seth Frotman to Dave Girouard and Alison Nicoll re: Potential Discrimination in Upstart Lending Platform (July 30, 2020), <https://www.naacpldf.org/wp-content/uploads/2020-07-30-FINAL-Demand-Letter.pdf>.

<sup>98</sup> *Id.* at 1.

<sup>99</sup> *Id.* at 1–2.

<sup>100</sup> *Id.* at 2 (citing Richard V. Reeves and Dimitrios Halikias, *Race gaps in SAT scores highlight inequality and hinder upward mobility*, Brookings (Feb. 1, 2017), <https://www.brookings.edu/research/race-gaps-in-sat-scores-highlight-inequality-and-hinder-upward-mobility/>; William C. Kidder & Jay Rosner, "How the SAT Creates "Built-In Headwinds": An Educational and Legal Analysis of Disparate Impact," 43 Santa Clara L. Rev. 131, 157 (2002); Jay Rosner, "The SAT: Quantifying the Unfairness Behind the Bubbles," in *SAT Wars: The Case for Test-Optional College Admissions* 134 (Joseph A. Soares 2015); Saul Geiser, "Norm-Referenced Tests and Race-Blind Admissions: The Case for Eliminating the SAT and ACT at the University of California" (Dec. 2017), <https://cshe.berkeley.edu/publications/norm-referenced-tests-and-race-blind-admissions-case-eliminating-sat-and-act-university>).

<sup>101</sup> Chandran & Frotman, *supra* note 97, at 4.

<sup>102</sup> Upstart presentation to SBPC and LDF, August 12, 2020 (on file with Monitor).

Upstart also established a “normalization” process for “Minority Serving Institutions” (“MSIs”)—which Upstart defines as schools where 80 percent or more of the student body are members of the same racial demographic group.<sup>103</sup> Under that process, Upstart normalized MSIs as a group to have equal graduate incomes as non-MSIs by calculating and using the distance, as a percentage, between a school’s graduate incomes and its respective school group average (*i.e.*, MSIs, non-MSIs). This process results in MSIs and non-MSIs being on average equal. Put another way, above average MSIs (in terms of graduate income) are treated above average overall by as much as they are above the MSI average. Any decisioning, including tranching, is then performed on this normalized information.

These changes are in place now and the fair lending testing conducted pursuant to this Monitorship will be of Upstart’s platform with these changes incorporated. Upstart emphasizes that it voluntarily adopted these changes and that none of Upstart’s fair lending tests—which are reported to the CFPB—have identified unlawful bias against any protected class, including any racial group.

#### **E. Upstart’s 2020 No Action Letter and Model Risk Assessment Plan**

In 2020, following an extension of its original NAL (and changes by the Bureau to its No-Action Letter Policy), Upstart received a second No-Action Letter from the CFPB.<sup>104</sup> Upstart’s application expressed an intent to “expand its reporting obligation to the Bureau and engage in additional fair lending testing.”<sup>105</sup> The NAL was conditioned on Upstart’s implementation of a Model Risk Assessment Plan, which, among other things, requires Upstart to provide the Bureau with model documentation; test its model and variables or groups of variables for adverse impact and predictive accuracy by group; report results to the Bureau; and research approaches that may produce less discriminatory alternative models that meet legitimate business needs.<sup>106</sup> The renewed 2020 No-Action Letter extends through November 2023.<sup>107</sup>

---

<sup>103</sup> Upstart’s definition is not the same as federal definitions of “minority-serving institution.” Seven categories of MSIs are defined in federal law. *See* 20 U.S.C. § 1067q(a). One is HBCUs (referred to as “part B institutions”), which are historically Black colleges or universities established prior to 1964, whose “principal mission was, and is, the education of Black Americans” and that are properly accredited. 20 U.S.C. § 1061(2). Another category is Hispanic-serving institutions, which are eligible institutions with enrollments that are at least 25% Hispanic students. 20 U.S.C. § 1101a(a)(5). Other categories include: Tribal College or University, Alaska Native- or Native Hawaiian-serving institution, Predominantly Black Institution, Asian American and Native American Pacific Islander-serving institution, and Native American-serving nontribal institution. 20 U.S.C. § 1067q(a). According to Upstart, it used its MSI definition because the relatively low thresholds for certain categories under the statutory definitions would have resulted in the inclusion of too many institutions being normalized; for example, many of the University of California schools would qualify as HSIs and/or AANPIs. Upstart represents that over 400 schools qualify as MSIs under its definition.

<sup>104</sup> CFPB, Letter Response to 2020 NAL Request at 1 (Nov. 30, 2020), [https://files.consumerfinance.gov/f/documents/cfpb\\_upstart-network-inc\\_no-action-letter\\_2020-11.pdf](https://files.consumerfinance.gov/f/documents/cfpb_upstart-network-inc_no-action-letter_2020-11.pdf).

<sup>105</sup> *Id.* at 1.

<sup>106</sup> *Id.* at 3.

<sup>107</sup> *Id.* at 2.



## **F. The Parties' Agreement to Appoint an Independent Monitor**

On December 1, 2020, LDF, Upstart, and the SBPC entered into an agreement under which Relman Colfax would act as a neutral, independent, third-party Monitor charged with conducting ongoing fair lending testing of Upstart's platform, including, but not limited to, Upstart's use of educational variables. As Monitor, our task is to assess the lending outcomes from Upstart's underwriting model to determine if the model causes or results in an adverse impact on any protected class and, if so, whether there are less discriminatory alternative practices that maintain the model's predictiveness. In our capacity as Monitor, we have engaged Sentrana, led by Syeed Mansur, to serve as a consultant to assist with these analyses. Sentrana is a leading firm in the field of automation systems that rely on advanced analytics and artificial intelligence. We have also engaged Bernard Siskin, of BLDS, to serve as a statistical consultant. Dr. Siskin is a leading expert on the use of statistical analyses to measure discrimination in the financial services industry.

As Monitor, we are also charged with preparing and making public certain reports. This Initial Report is the first such report. We are responsible for preparing quarterly reports for the parties detailing the tests we conduct as part of this review and the results of those tests; any trends or changes observed from prior reports; recommendations for any less discriminatory alternatives that maintain the model's predictiveness and meet other legitimate business needs; and whether Upstart appropriately addressed prior recommendations. A public version of each quarterly report will summarize the general findings of the report, best practices identified, and any aspects of the model that raise particular fair lending concerns or implicate novel insights on educational equity that serve the public interest. Upstart has agreed to adopt, implement, and maintain for the term of the agreement the recommendations in each periodic report.

Periodic Reports will be published as follows:

April 14, 2021:	Initial Report (this Report)
October 14, 2021:	Periodic Report 1
February 14, 2022:	Periodic Report 2
June 15, 2022:	Periodic Report 3
October 14, 2022:	Periodic Report 4
February 15, 2023:	Periodic Report 5
June 14, 2023:	Periodic Report 6
October 16, 2023:	Periodic Report 7