



A Holistic Approach to Addressing Structural Racism in Insurance

NAIC Consumer Liaison Committee

August 12, 2022

Birny Birnbaum
Center for Economic Justice

The Center for Economic Justice

CEJ is a non-profit consumer advocacy organization dedicated to representing the interests of low-income and minority consumers as a class on economic justice issues. Most of our work is before administrative agencies on insurance, financial services and utility issues.

On the Web: www.cej-online.org

About Birny Birnbaum

Birny Birnbaum is the Director of the Center for Economic Justice, a non-profit organization whose mission is to advocate on behalf of low-income consumers on issues of availability, affordability, accessibility of basic goods and services, such as utilities, credit and insurance.

Birny, an economist and former insurance regulator, has worked on racial justice issues for 30 years. He performed the first insurance redlining studies in Texas in 1991 and since then has conducted numerous studies and analyses of racial bias in insurance for consumer and public organizations. He has served for many years as a designated Consumer Representative at the National Association of Insurance Commissioners and is a member of the U.S. Department of Treasury's Federal Advisory Committee on Insurance, where he co-chairs the subcommittee on insurance availability.

Birny served as Associate Commissioner for Policy and Research and the Chief Economist at the Texas Department of Insurance. At the Department, Birny developed and implemented a robust data collection program for market monitoring and surveillance.

Birny was educated at Bowdoin College and the Massachusetts Institute of Technology. He holds Master's Degrees from MIT in Management and in Urban Planning with concentrations in finance and applied economics. He holds the AMCM certification.

Why CEJ Works on Insurance Issues

Insurance Products Are Financial Security Tools Essential for Individual and Community Economic Development:

CEJ works to ensure ***fair access*** and ***fair treatment*** for insurance consumers, particularly for low- and moderate-income consumers.

Insurance is the Primary Institution to Promote Loss Prevention and Mitigation, Resiliency and Sustainability:

CEJ works to ensure insurance institutions maximize their role in efforts to reduce loss of life and property from catastrophic events and to ***promote resiliency and sustainability*** of individuals, businesses and communities.

Statutory Foundation:

Fair and Unfair Discrimination in Insurance

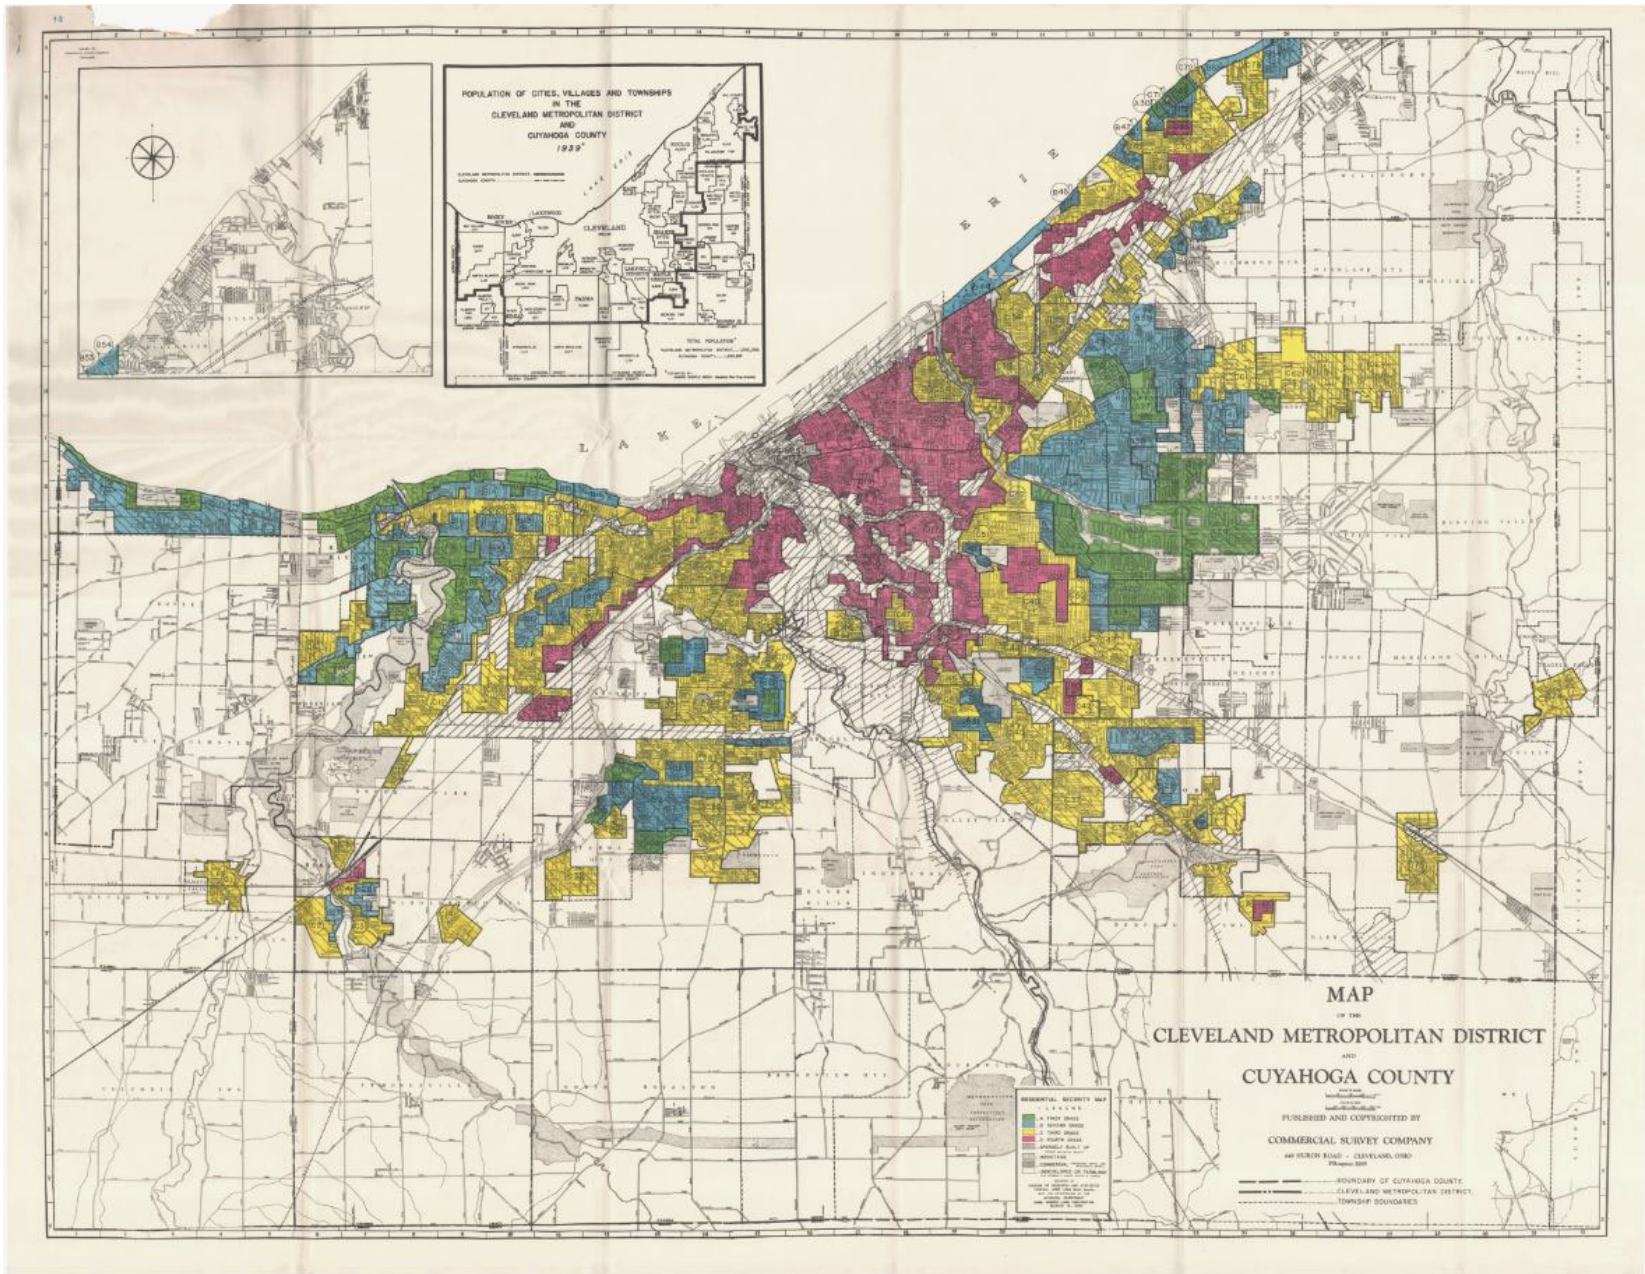
Provisions regarding fair and unfair discrimination are generally found in two parts of U.S. insurance statutes: rating and unfair trade practices.

We find two types of unfair discrimination:

- Actuarial – there must be an actuarial basis for distinction among groups of consumers; and
- Protected Classes – distinctions among groups defined by certain characteristics – race, religion, national origin – prohibited regardless of actuarial basis.

What Information Does This Map of Cleveland Present?

- a. Concentration of Minority Population
- b. Eviction Rates
- c. COVID Infections and Deaths Rates
- d. Flood Risk
- e. Environment-related Illnesses
- f. Intensity of Policing
- g. Predatory Lending
- h. Federal Home Loan Eligibility 1930's to 1960's



What is Structural or Systemic Racism?

Structural Racism: A system in which public policies, institutional practices, cultural representations, and other norms work in various, often reinforcing ways to perpetuate racial group inequity. It identifies dimensions of our history and culture that have allowed privileges associated with “whiteness” and disadvantages associated with “color” to endure and adapt over time.

Aspen Institute at <https://www.aspeninstitute.org/wp-content/uploads/files/content/docs/rcc/RCC-Structural-Racism-Glossary.pdf>

Why Are Race and Other Protected Class Characteristics Carved Out of Fair Actuarial Discrimination?

The existence of historical, intentional discrimination based on these characteristics – discrimination that violates state and federal constitutions. But, also, the recognition that the historical discrimination has long-lasting effects that disadvantage those groups. Stated differently, you can't enslave a population for two hundred years and then expect the legacy of that enslavement will disappear overnight.

We continue to see those legacies of historical discrimination – systemic racism -- today both directly and indirectly in policing and criminal justice, housing, and the impacts of the Covid-19 pandemic.

How Can Structural Racism Manifest Itself in Insurance – Whether for Marketing, Pricing or Claims Settlement?

1. ***Disparate Intent:*** Intentional Use of Race
2. ***Proxy Discrimination:*** Disproportionate Racial Outcomes Tied to Use of Proxies for Race, Not to Outcomes
3. ***Disparate Impact:*** Disproportionate Racial Outcomes Tied to Historic Discrimination and Embedded in Insurance Outcomes

We'll assume insurers don't intentionally discriminate on the basis of race. Addressing proxy discrimination is easy – the data are not predicting insurance outcomes so they violate both the actuarial and protected class requirements for fair discrimination. Addressing disparate impact requires empirical analysis and public policy considerations.

Practices That Raise Concerns About Disparate Impact and Proxy Discrimination on the Basis of Race

Price Optimization and Consumer Lifetime Value Scores

By definition, these algorithms used by insurers utilize non-cost factors to differentiate among consumers and the factors and data reflect bias against communities of color.

Credit-Based Insurance Scores

The consumer credit information factors used in CBIS are highly correlated with race. The Missouri Department of Insurance found that the single best predictor of the average CBIS in a ZIP Code was minority population.

Criminal History Scores

Here, the problem is not just the legacy of historical discrimination, but ongoing discrimination in policing and criminal justice.

How Can an Insurer or Regulator Identify Proxy Discrimination and Disparate Impact?

Fortunately, there are statistical techniques that enable an analyst to determine whether a particular type of data is predicting race or the insurance outcome or both. And, if both, how much of a proxy for race and how much predictive of the insurance outcome.

The slides at the end of the presentation explain in more detail.

Holistic Approach to Addressing Structural Racism in Insurance

1. Testing by insurers of their algorithms and actual consumer outcomes for proxy discrimination and disparate impact based on protected class characteristics – principles-based model governance is not sufficient.
2. Testing of algorithms used in all consumer-facing parts of the insurance life-cycle – marketing, underwriting, pricing, claims settlement and anti-fraud.
3. Testing of algorithm factors simultaneously, not in isolation.
4. Regulatory collection of granular consumer outcome data from insurers, including applications and related quotes as well as policy and claims outcomes to support regulatory testing.
5. Regulatory guidance for bias thresholds and equity trade-offs.

Insurer Testing of Algorithms / Actual Consumer Outcomes

Some have suggested an algorithmic model governance approach to addressing structural racism in insurance similar to the approach used for ORSA and preventing cyber breaches.

Model governance is essential, but not sufficient. Testing of actual consumer outcomes is reasonable and necessary because there are literally millions of such outcomes in every phase of the insurance life cycle that be analyzed.

Insurers test these outcomes as they develop the algorithms for marketing, pricing, claims settlement and anti-fraud. Testing for spurious correlations (proxy discrimination) and disparate impact on the basis of protected class characteristics should simply be part of model development.

Uniform Methods of Testing and Evaluation across Insurers

A “principles-based approach” to address structural racism is not necessary or desirable, because uniform methods of testing and evaluation across insurers is possible because all insurers share the same types of consumer outcomes, regardless of business model or product:

- Did the insurer receive an application?
- Did the application result in a policy?
- If a policy was issued, what was the premium and coverage provided?
- Was a claim filed?
- Was the claim denied or paid?
- If the claim was paid, how much?

Why Test for Disparate Impact and Proxy Discrimination in All Aspects of Insurers' Operations?

While pricing / rating has gotten the most regulatory attention in terms of complex model scrutiny by regulators, it's imperative for insurers and regulators to test algorithms used in all aspects of the insurance life-cycle for racial bias.

Antifraud algorithms are particularly susceptible to reflecting and perpetuating historic racism because antifraud algorithms can identify suspicious claims. If the identification of suspicious claims is racially-biased, so will the identification of claims as fraudulent – a claim that's not investigated will not be identified as fraud.

Marketing algorithms also raise great concern – the new data sources and algorithms used to micro-target consumers have become the *de facto* gateway for access to insurance.

Testing of Marketing Algorithms

Consider the following quotes from 2005 to present. In 2005, in a meeting with investment analysts, the CEO of a major publicly-traded insurer was effusive about the benefits of the then relatively new use of consumer credit information – referred to as tiered pricing.

Tiered pricing helps us attract higher lifetime value customers who buy more products and stay with us for a longer period of time. That's Nirvana for an insurance company.

This year, we've expanded from 7 basic price levels to 384 potential price levels in our auto business.

Tiered pricing has several very good, very positive effects on our business. It enables us to attract really high quality customers to our book of business.

The key, of course, is if 23% or 20% of the American public shops, some will shop every six months in order to save a buck on a six-month auto policy. ***That's not exactly the kind of customer that we want.*** So, the key is to use our drawing mechanisms and our tiered pricing to find out of that 20% or 23%, to find those *that are unhappy with their current carrier, are likely to stay with us longer, likely to buy multiple products and that's where tiered pricing and a good advertising campaign comes in.*

Now fast forward to 2017, when the new CEO of that insurer told investment analysts:

The insurer's "universal consumer view" keeps track of information on 125 million households, or 300 million-plus people, Wilson said.

"When you call now they'll know you and know you in some ways that they will surprise you, and give them the ability to provide more value added, so we call it the trusted adviser initiative"

Focus on Holistic Testing, Not Individual Factors in Isolation

Over the last several decades, much of the focus on efforts to address racial bias in insurance has been on data sources that are highly correlated with race with calls to ban those factors.

While insurers should surely not be using data sources and factors that are proxies for race and not predictive of insurance outcomes, testing for racial bias must be of the entire algorithm and all the data sources used in the algorithm simultaneously.

- Eliminating one factor may simply shift the racial bias to another factor instead of eliminating the racial bias. Testing of the algorithm is designed to eliminate proxy discrimination and identify disparate impact of the entire algorithm.
- Multi-variate testing can remove eliminate correlations with race and reveal the factor's true contribution to explaining the insurance outcome and provide a statistical basis for addressing disparate impact.

Modernizing Data Reporting for Market Regulation is Essential

The current regulatory data collection is woefully outdated and doesn't serve the needs of regulators and policymakers generally. In particular, testing for protected class bias requires the reporting of granular consumer outcome data by insurers and analyses of those data by regulators. Absent this type of empirical analysis by regulators, we will not be able to move beyond the historical debates about race and insurance and not be able to ground our anti-racism efforts in the risk-based foundation of insurance.

The collection of granular consumer outcome data must include individual applications for insurance that don't end up in policy issuance. As mentioned, marketing algorithms have become the new gatekeeper for insurance access – analysis of application data is essential to see if those algorithms systematically deny communities of color such access.

Regulatory Standards for Bias Thresholds and Equity Trade-Offs

While there may be some data sources and factors that lie at the extremes – pure proxies for protected classes or pure predictors of risk-based insurance outcomes – the nature of structural racism means that the vast majority of data sources will likely result in some racial disparities.

Insurers need guidance on, for example, on

- What degree of proxy discrimination should lead to prohibiting the use of that data source or factor from the deployed algorithm?
- How can an insurer utilize alternate data sources to maintain the algorithm's efficiency while reducing disparate impact?
- What trade-off between reducing disparate impact and weakening the algorithm's efficiency is reasonable? If we could change an algorithm to eliminate 95% of disparate impact at a cost of 5% of statistical predictive strength, would that be a fair trade?

Testing for Disparate Impact and Proxy Discrimination:

A Natural Extension of Typical Insurer Practices

While proxy discrimination and disparate impact are different forms of unfair discrimination, there is a common methodology to test for both.

There is a long history of and many approaches to identifying and minimizing disparate impact in employment, credit and insurance. But, the general principle is to identify and remove the correlations between the protected class characteristic and the predictive variables by explicit consideration of the protected class characteristic.

The techniques to analyze proxy discrimination and disparate impact are the same techniques insurers use in developing predictive models for all aspects of the insurance life cycle. See below for more technical explanation.

Risk Segmentation is not the Purpose of Insurance

Insurer trades argue that anything that restricts their ability to segment the population for any aspect of the insurance life cycle will destroy the cost-based foundation of insurance, will lead to “good risks” subsidizing “bad risks” and lead to insurer financial ruin.

In fact, the existence of protected class characteristics demonstrates that risk segmentation – “predicting risk” – is not the goal of insurance but a tool to help achieve the real goal of insurance – a risk pooling mechanism providing financial security for as many as possible and particularly for those with modest resources. Insurers’ arguments for unfettered risk classifications are inconsistent with the goal of insurance.

While some risk segmentation is necessary to avoid adverse selection, the logical extension of that argument is not unlimited risk segmentation. In fact, if unlimited risk segmentation was necessary, we would see all insurers using all risk characteristics – they don’t – and collapsing markets in states where some limitations on risk characteristics exist – they aren’t.

Disparate Impact Analysis Improves Cost-Based Pricing

With proxy discrimination, an insurer is using a factor – a characteristic of the consumer, vehicle, property or environment – that is predicting race and not the insurance outcome. Proxy discrimination is, therefore, a spurious correlation and eliminating such spurious correlation improves cost-based pricing. Since proxy discrimination is indirect racial discrimination, it is currently a prohibited practice. Testing would therefore both improve risk-based pricing and stop unintentional or intentional racial discrimination.

There is a long history and many approaches to identifying and minimizing disparate impact in employment, credit and insurance. But, the general principle is to identify and remove the correlations between the protected class characteristic and the predictive variables. Testing identifies true disparate impact that may require a public policy that recognizes equity – such as the prohibition against using race itself as a factor.

Why is it Reasonable and Necessary to Recognize Disparate Impact as Unfair Discrimination in Insurance?

1. It makes no sense to permit insurers to do indirectly what they are prohibited from doing directly. If we don't want insurers to discriminate on the basis of race, why would we ignore practices that have the same effect?
2. It improves risk-based and cost-based practices.
3. In an era of Big Data, systemic racism means that there are no “facially-neutral” factors.

The Evolution of Insurers' Analytics: Univariate to Multivariate Analysis

In the past 30 years, insurers have moved away from univariate analysis to multivariate analysis – from analyzing the effects of one risk characteristic at a time to simultaneous analysis of many risk characteristics.

What the problem with univariate analysis?

If I analyze the relationship of age, gender and credit score – each individually – to the likelihood of a claim, the individual results for each risk characteristic are likely capturing some of the effects of the other risk characteristics – ***because age, gender and credit score (or other risk classifications) may be correlated to each other*** as well as to the outcome variable.

How does multi-variate analysis address this problem?

Testing for Disparate Impact and Proxy Discrimination:

A Natural Extension of Typical Insurer Practices

Here's a simple illustration of a multivariate model. Let's create a simple model to predict the likelihood of an auto claim:

$$b_0 + b_1X_1 + b_2X_2 + b_3X_3 + e = y$$

$X_1, X_2 + X_3$ are the predictive variables trying to predict y .

Say that $X_1, X_2 + X_3$ are age, gender and credit score and we are trying to predict y – the likelihood of an auto insurance claim

Let's assume that all three X s are statistically significant predictors of the likelihood of a claim and the b values are how much each X contributes to the explanation of claim. The b values can be tested for statistical significance – how reliable are these estimates of the contribution of each X ?

By analyzing these predictive variable simultaneously, the model removes the correlation among the predictive variables.

Use of Control Variables in Multivariate Insurance Models

Suppose an insurer want to control for certain factors that might distort the analysis? For example, an insurer developing a national pricing model would might want to control for different state effects like different age distributions, different occupation mixes or differences in jurisprudence. An insurer would add one or more control variables.

$$b_0 + b_1X_1 + b_2X_2 + b_3X_3 + b_4C_1 + e = y$$

C_1 is a control variable – let's say for State. By including State as a control variable, the correlation of the Xs to State is statistically removed and the new b values are now the contribution of the Xs, independent of their correlation to State, to explaining the likelihood of a claim. When the insurer deploys the model, it still only uses the X variables, but now with more accurate b values.

Disparate Impact as Both a Standard and a Methodology

Let's go back to multi-variate model, but now use Race as a control variable:

$$b_0 + b_1X_1 + b_2X_2 + b_3X_3 + b_4R_1 + e = y$$

R_1 is a control variable – by including race in the model development, the correlation of the X s to race is statistically removed and the new b values are now the contribution of the X s, independent of their correlation to race, to explaining the likelihood of a claim

How Do We Interpret the Disparate Impact Analysis?

$$b_0 + b_1X_1 + b_2X_2 + b_3X_3 + b_4R_1 + e = y$$

Result: No Proxy Discrimination or Disparate Impact

Outcome	Interpretation	Indicated Action
R is not statistically significant and there is little change to b1, b2 and b3.	There is little correlation between X1, X2 and X3 and race, little or no disparate impact or proxy discrimination	None, utilize the model.

How Do We Interpret the Disparate Impact Analysis?

$$b_0 + b_1X_1 + b_2X_2 + b_3X_3 + b_4R_1 + e = y$$

Result: Proxy Discrimination

Outcome	Interpretation	Indicated Action
R is statistically significant and b1 has lost its statistical significance	X1 was largely a proxy for race and the original predictive value of X1 was spurious. This is an example of proxy discrimination	Remove X1 from the marketing, pricing, claims settlement or anti-fraud model.

How Do We Interpret the Disparate Impact Analysis?

$$b_0 + b_1X_1 + b_2X_2 + b_3X_3 + \mathbf{b_4R_1} + e = y$$

Result: Disparate Impact

Outcome	Interpretation	Indicated Action
R is statistically significant and has a large impact on the outcome, but b1, b2 and b3 remain largely unchanged and statistically significant	This is an example of disparate impact.	Are X1, X2 or X3 essential for the insurer's business purposes? Are there less discriminatory approaches available? Would eliminating a predictive variable significantly reduce the disparate impact but not materially affect the efficiency or productiveness of the model?

How Do We Interpret the Disparate Impact Analysis?

$$b_0 + b_1X_1 + b_2X_2 + b_3X_3 + \mathbf{b_4R_1} + e = y$$

Result: Some Proxy Discrimination, Some Disparate Impact

Outcome	Interpretation	Indicated Action
R is statistically significant, but b1, b2 and b3 remain statistically significant with different values from the original.	X1, X2 and X3 are correlated to race, but also predictive of the outcome, even after removing the variables' correlation to race. This is an example of some proxy discrimination and some disparate impact.	Depending on the significance of the racial impact, utilize the model with the revised predictive variable coefficients, consider prohibiting a variable on the basis of equity or both.

Insurers Don't Collect Applicant's Race – How Can an Actuary Get Data on Race to Perform a Disparate Impact Analysis?

1. Assign a racial characteristic to an individual based on racial characteristic of a small geographic area – Census data at the census block level.
2. Utilize the Bayesian Improved Surname Geocoding Method, based on census geography and surname data.¹
3. Reach out to data brokers and vendors for a new data service.

¹ See Consumer Financial Protection Bureau, "Using publicly available information to proxy for unidentified race and ethnicity."
<https://www.consumerfinance.gov/data-research/research-reports/using-publicly-available-information-to-proxy-for-unidentified-race-and-ethnicity/>
and Yin Zhang, "Assessing Fair Lending risks Using Race/Ethnicity Proxies."

Ethical Algorithms -- Sources

Pauline T. Kim, "Auditing Algorithms for Discrimination"

<https://www.pennlawreview.com/online/166-U-Pa-L-Rev-Online-189.pdf>

Claire Whitaker, "Ethical Algorithms"

<https://www.kdnuggets.com/2019/03/designing-ethical-algorithms.html>

Erin Russel, "The Ethical Algorithm"

<https://www.cognitivetimes.com/2019/01/the-ethical-algorithm/>

Barocas and Selbst

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2477899

Kroll, et al, "Accountable Algorithms:

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2765268

Virginia Eubanks, *Automating Inequality: How High Tech Tools Profile, Police and Punish the Poor*

Selbst and Barocas, "The Intuitive Appeal of Explainable Machines

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3126971

Levy and Barocas, "Designing Against Discrimination in Online Markets

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3084502

New York Times, "Algorithms and Bias, Q and A with Cynthia Dwork," 10 August 2015

<https://www.nytimes.com/2015/08/11/upshot/algorithms-and-bias-q-and-a-with-cynthia-dwork.html>

Martin, Kirsten E. M., What Is an Ethical Algorithm (And Who Is Responsible for It?) (October 21, 2017). Available at SSRN:

<https://ssrn.com/abstract=3056692> or <http://dx.doi.org/10.2139/ssrn.3056692>

Kirsten Martin, "Ethical Implications and Accountability of Algorithms"

<http://kirstenmartin.net/wp-content/uploads/2013/11/Martin-JBE-Ethics-and-Accountability-of-Algorithms.pdf>

Kirsten Martin, DATA AGGREGATORS, BIG DATA, & RESPONSIBILITY ONLINE

<http://kirstenmartin.net/wp-content/uploads/2013/11/AAPOR-Martin-Info-Value-Chain-v2.pdf>

AIandBigData: A blueprint for human rights, social and ethical impact assessment Alessandro Mantelero

<https://reader.elsevier.com/reader/sd/pii/S0267364918302012?token=3836947F0CAD3C145A1F273E3CBE6C38F67E777DD7E4D590548F481916130DAACA8D57BED4667BD1FE1F4D8FC80E7C56>