OUR MEETING WILL BEGIN SHORTLY

WELCOME TO THE

JOINT MEETING OF THE INNOVATION, CYBERSECURITY, AND TECHNOLOGY (H)COMMITTEE AND THE NAIC/CONSUMER LIAISON COMMITTEE

October 14, 2022

- ***Committee members (H Committee and Consumer Liaison)**, please post your name and state in the "Chat" upon joining the meeting*
- *All audio will be muted upon entry*
- *Enter with video on or off (your choice)*
- *If attendees would like to speak, please use the "Raise Hand" feature*
- *Use the "Chat" feature for comments or assistance*
- *If you have joined by phone, to mute and unmute your line, press *6*
- *For additional help, please contact NAIC Technical Support team at MeetingTechHelp@naic.org or call (866) 874-4905*

NAIC

NATIONAL ASSOCIATION OF INSURANCE COMMISSIONERS

NAIC NATIONAL ASSOCIATION OF INSURANCE COMMISSIONERS

Date: 10/11/22

JOINT MEETING OF THE INNOVATION, CYBERSECURITY, AND TECHNOLOGY (H) COMMITTEE AND THE NAIC/CONSUMER LIAISON COMMITTEE
Friday, October 14, 2022
2:00 – 4:00 p.m. ET / 1:00 – 3:00 p.m. CT / 12:00 – 2:00 p.m. MT / 11:00 a.m. – 1:00 p.m. PT

ROLL CALL

INNOVATION, CYBERSECURITY, AND TECHNOLOGY (H) COMMITTEE

Kathleen A. Birrane, Chair	Maryland	Troy Downing	Montana
Evan G. Daniels, Co-Vice Chair	Arizona	Adrienne A. Harris	New York
Dana Popish Severinghaus, Co-Vice Chair	Illinois	Jon Godfread	North Dakota
		Judith L. French	Ohio
Karima M. Woods	District of Columbia	Elizabeth Kelleher Dwyer	Rhode Island
John F. King	Georgia	Carter Lawrence	Tennessee
Amy L. Beard	Indiana	Kevin Gaffney	Vermont
Chlora Lindley-Myers	Missouri	Mike Kreidler	Washington

NAIC Support Staff: Denise Matthews/Scott Morris

NAIC/CONSUMER LIAISON COMMITTEE

Andrew R. Stolfi, Chair	Oregon	Anita G. Fox	Michigan
Grace Arnold, Vice Chair	Minnesota	Mike Chaney	Mississippi
Mark Fowler	Alabama	Chlora Lindley-Myers	Missouri
Lori K. Wing-Heier	Alaska	Eric Dunning	Nebraska
Peni Itula Sapini Teo	American Samoa	Barbara D. Richardson	Nevada
Evan G. Daniels	Arizona	Chris Nicolopoulos	New Hampshire
Alan McClain	Arkansas	Adrienne A. Harris	New York
Ricardo Lara	California	Mike Causey	North Carolina
Andrew N. Mais	Connecticut	Jon Godfread	North Dakota
Trinidad Navarro	Delaware	Judith L. French	Ohio
Karima M. Woods	District of Columbia	Michael Humphreys	Pennsylvania
David Altmaier	Florida	Cassie Brown	Texas
Colin M. Hayashida	Hawaii	Jon Pike	Utah
Dean L. Cameron	Idaho	Scott A. White	Virginia
Vicki Schmidt	Kansas	Mike Kreidler	Washington
James J. Donelon	Louisiana	Allan L. McVey	West Virginia
Kathleen A. Birrane	Maryland	Nathan Houdek	Wisconsin

NAIC Support Staff: Lois E. Alexander

© 2022 National Association of Insurance Commissioners

AGENDA

1. Hear Presentations on Algorithmic Bias and Approaches Insurance Companies Are or Can Implement to Manage and Mitigate the Risk of Unintended Bias and Illegal Discrimination When Developing and Using Artificial Intelligence (AI)/Machine Learning (ML)—*Tony Cotto (National Association of Mutual Insurance Companies—NAMIC), Karen Melchert (American Council of Life Insurers—ACLI), and David F. Snyder (American Property Casualty Insurance Association—APCIA)*

2. Hear Presentations on Algorithmic Bias and a Holistic Approach to Confronting Structural Racism in Insurance—*Birny Birnbaum (Center for Economic Justice—CEJ), Dr. Michael Akinwumi (National Fair Housing Alliance—NFHA), Peter Kochenburger (Southern University Law Center and University of Connecticut School of Law), Morgan Williams (NFHA)*

3. Receive Comments from Interested Parties—*Commissioner Kathleen A. Birrane (MD)*

4. Discuss Any Other Matters Brought Before the Committees —*Commissioner Kathleen A. Birrane (MD)*

5. Adjournment

© 2022 National Association of Insurance Commissioners

Presentations on Algorithmic Bias and Approaches Insurance Companies Are or Can Implement When Developing and Using Artificial Intelligence (AI)/Machine Learning (ML)

- Tony Cotto (National Association of Mutual Insurance Companies—NAMIC)
- Karen Melchert (American Council of Life Insurers—ACLI)
- David F. Snyder (American Property Casualty Insurance Association—APCIA)

NAIC
NATIONAL ASSOCIATION OF INSURANCE COMMISSIONERS

AI/ML and Insurance Regulation

NAIC

October 14, 2022

David Snyder

APCIA

Copyright © 2022 by the American Property Casualty Insurance Association

Commitment

APCIA agrees with the NAIC on the need to collectively determine how to tackle concerns related to fairness and preventing unlawful discrimination while seeking improvements that strengthen competitive markets and addressing potential inequities while preserving the risk-based foundation of insurance.

Copyright © 2022 by the American Property Casualty Insurance Association

Benefits of Insurers' Use of AI/ML

- Better and more efficiently meeting customer expectations in a digital world
- More rapidly responding to, settling and paying claims—(e.g., auto physical damage claims and exterior and interior property claims)
- More accurately and objectively assessing risk—telematics in auto insurance
- Analyzing company performance to improve product offerings, customer service and compliance

Copyright © 2022 by the American Property Casualty Insurance Association

Need for Definitions

- Unclear definitions of AL, ML and other basic terms make it difficult to respond most constructively
- General definitions of “bias” do not consider long-standing legislated insurance regulatory and judicial standards

Copyright © 2022 by the American Property Casualty Insurance Association

Importance of Robust Governance

- Should be a priority—critical foundational element
- Rapidly evolving, with many different stakeholders—NAIC principles and NIST draft AI Risk Management Framework
- Prominent role of human review and decision-making
- Should be flexible, proportionate, scalable and explainable

Copyright © 2022 by the American Property Casualty Insurance Association

Insurers' Use of AI/ML Already Subject to Legislated Regulatory and Judicial Standards

- When used for regulated activities—underwriting, pricing and claims settlement, AI/ML already extensively regulated
- Use of AI/ML should neither expand nor contract regulatory standards and enforcement when used for activities previously conducted purely by humans
- Disparate impact well defined by the courts—
  - Texas Dept of Housing and Community Affairs v. The Inclusive Communities Project (576 U.S. 519 (2015)) – requires consideration of valid interests as well as adverse outcomes

Copyright © 2022 by the American Property Casualty Insurance Association

Conclusion

- APCIA committed to working constructively with NAIC and state regulators in their work on AI/ML
- Need for clear definitions
- Governance is paramount
- Need to avoid inhibiting beneficial uses of AI/ML
- Legislated regulatory standards and judicial decisions should be applied to AI/ML

Copyright © 2022 by the American Property Casualty Insurance Association

Presentations on Algorithmic Bias and a Holistic Approach to Confronting Structural Racism in Insurance

- Birny Birnbaum (Center for Economic Justice—CEJ)
- Dr. Michael Akinwumi (National Fair Housing Alliance—NFHA)
- Peter Kochenburger (Southern University Law Center and University of Connecticut School of Law)
- Morgan Williams (National Fair Housing Alliance—NFHA)

NAIC
NATIONAL ASSOCIATION OF INSURANCE COMMISSIONERS

CENTER for ECONOMIC JUSTICE | FAIR ACCESS FAIR TREATMENT

A Holistic Approach to Addressing Structural Racism in Insurance

Joint Meeting of NAIC Consumer Liaison and Innovation, Cybersecurity and Technology Committees

October 14, 2022

Birny Birnbaum

The Center for Economic Justice

CEJ is a non-profit consumer advocacy organization dedicated to representing the interests of consumers – particularly consumers of modest means and of communities of color – 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 is finance and applied economics. He holds the AMCM certification.

CEO Comments after the George Floyd Murder

“In the coming days, I encourage each of us to step outside of our comfort zones, seek to understand, engage in productive conversations and hold ourselves accountable for being part of the solution. We must forever stamp out racism and discrimination.”

Floyd’s death in Minneapolis is the latest example of “a broken society, fueled by a variety of factors but all connected by inherent bias and systemic racism. Society must take action on multiple levels and in new ways.“

How do we turn these words into action?

Why CEJ Works on and Promotes Insurance

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

In the U.S., fair and unfair discrimination is defined in two ways, typically found in rating and unfair trade practice statutes and regulations.

- **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 Does This Map of Cleveland Present?

- Concentration of Minority Population
- Eviction Rates
- COVID Infections and Deaths Rates
- Flood Risk
- Environment-related Illnesses
- Intensity of Policing
- Predatory Lending
- Federal Home Loan Eligibility 1930’s to 1960’s

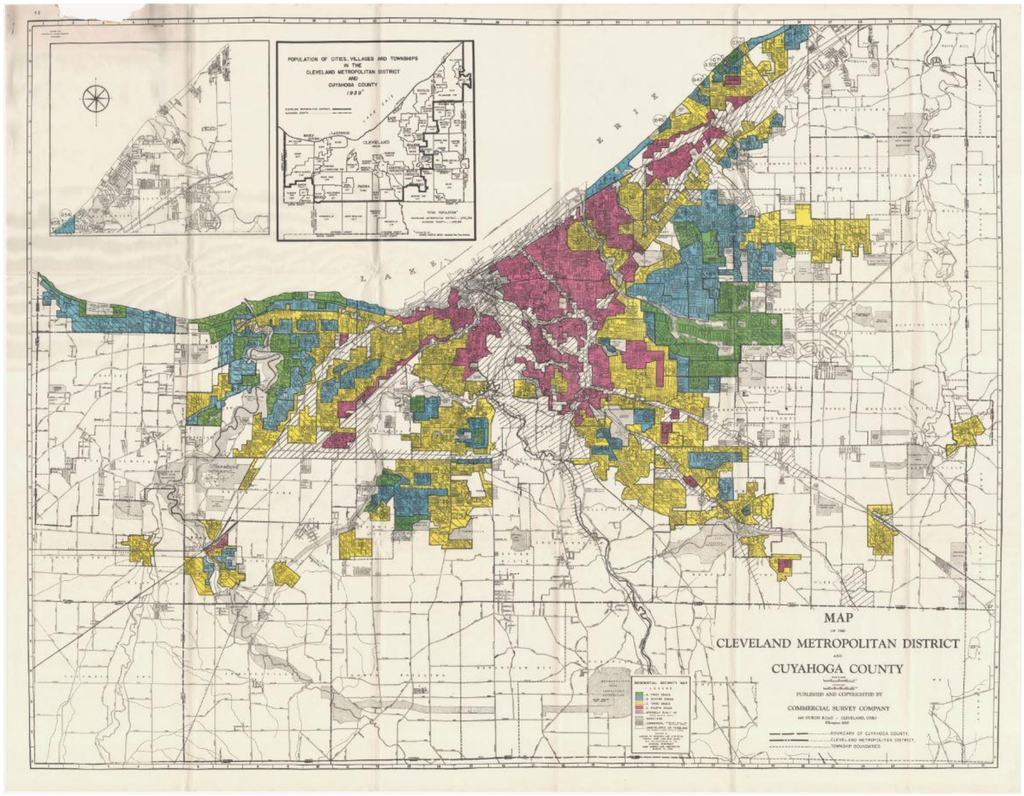

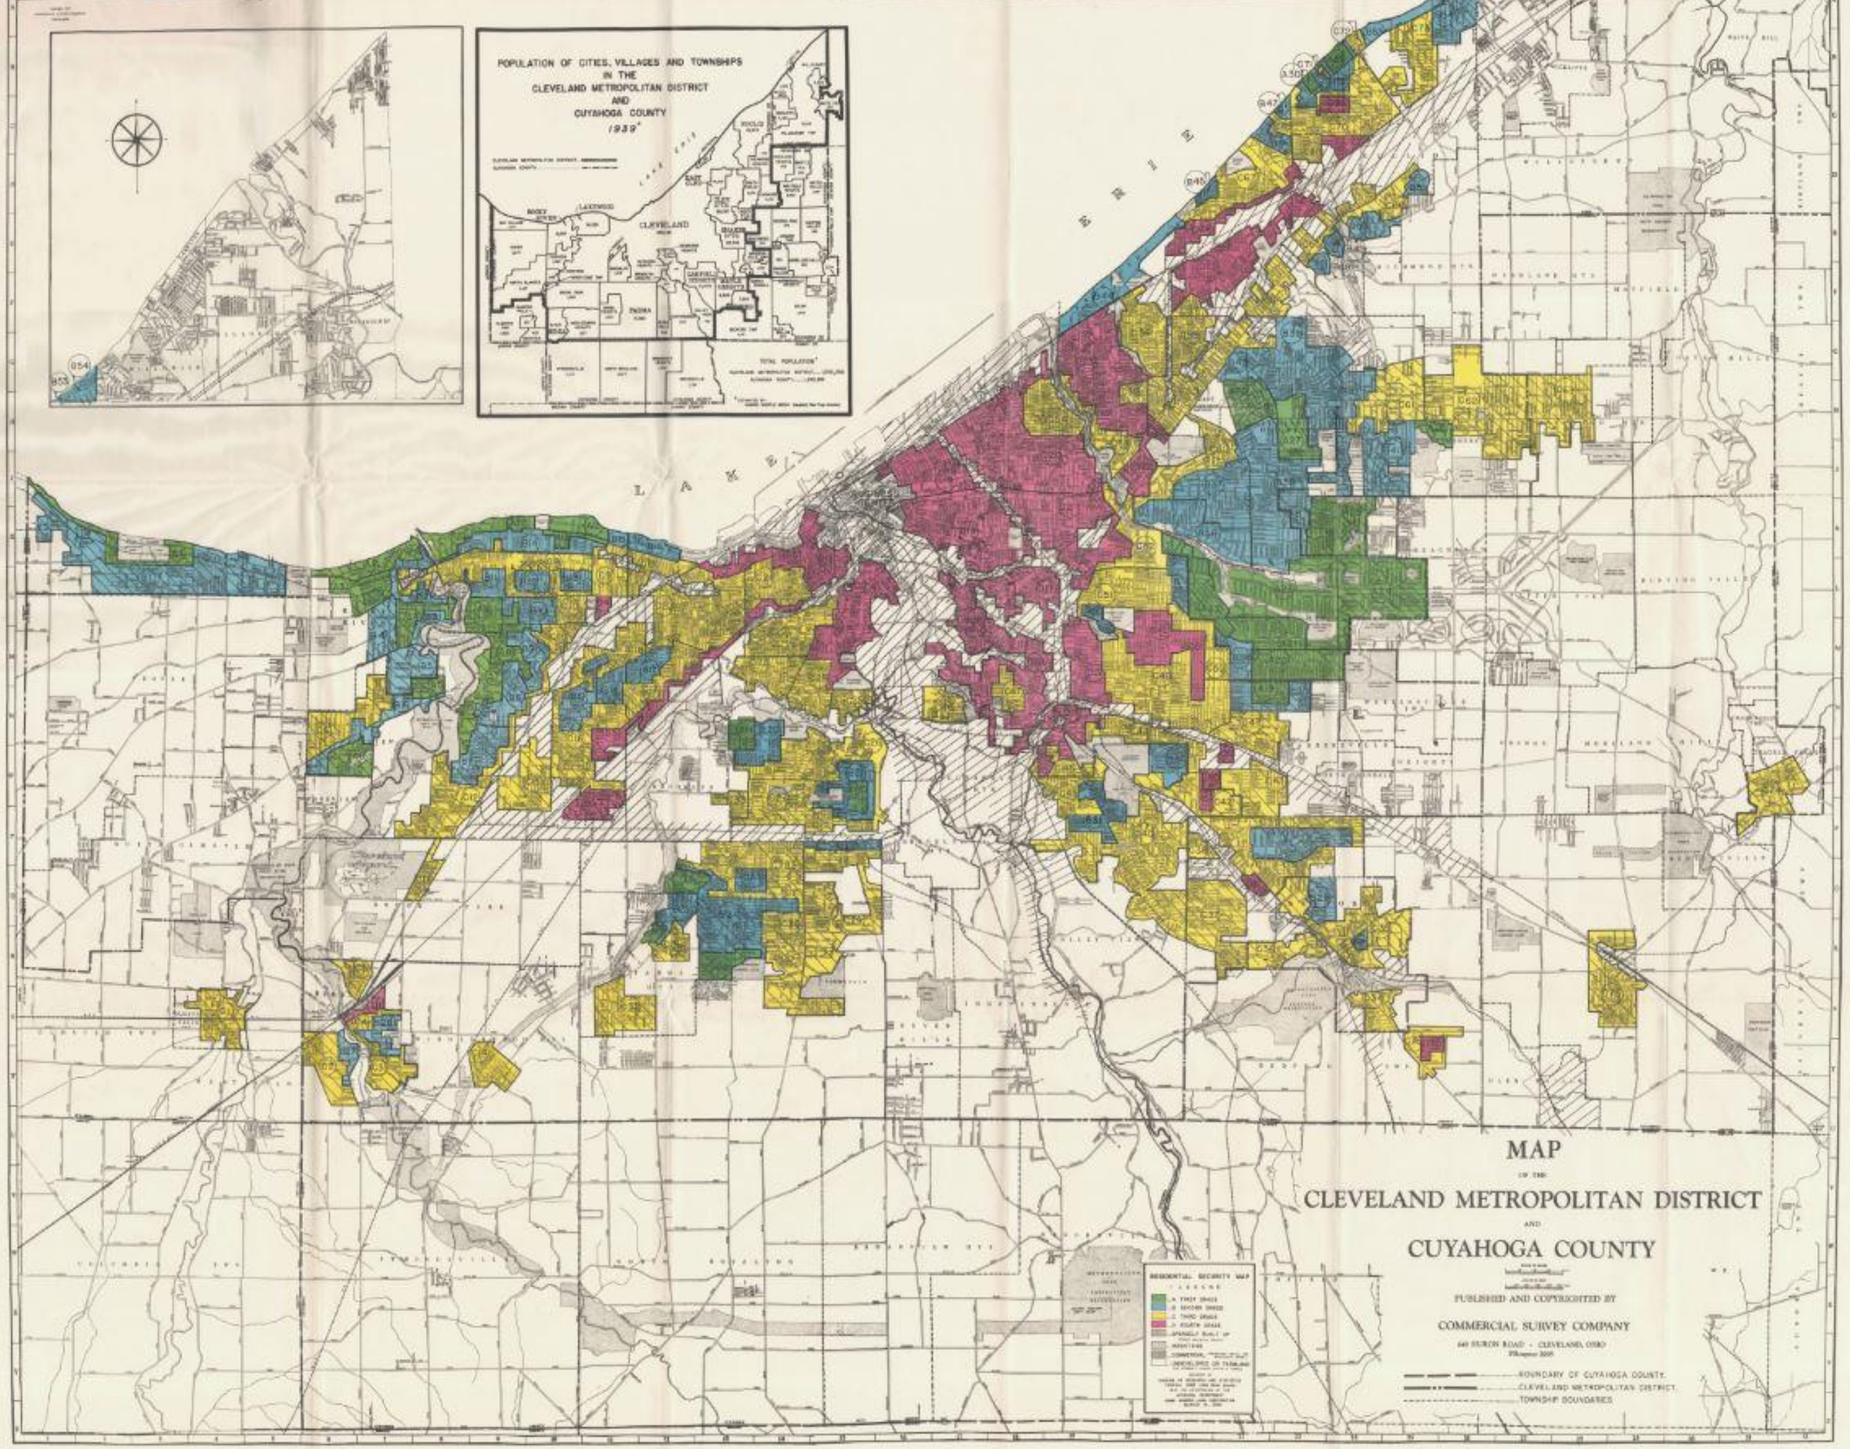
POPULATION OF CITIES, VILLAGES AND TOWNSHIPS
IN THE
CLEVELAND METROPOLITAN DISTRICT
AND
CUYAHOGA COUNTY
1939
LAKE ERIE
CLEVELAND
LAKEWOOD
MAP
OF THE
CLEVELAND METROPOLITAN DISTRICT
AND
CUYAHOGA COUNTY
PUBLISHED AND COPYRIGHTED BY
COMMERCIAL SURVEY COMPANY
640 HURON ROAD - CLEVELAND, OHIO
BOUNDARY OF CUYAHOGA COUNTY
CLEVELAND METROPOLITAN DISTRICT
TOWNSHIP BOUNDARIES

Why are Race and Other Protected Classes Carved Out of Fair Actuarial Discrimination?

Of course, this is a map of federal home loan eligibility from 1940 – The red areas represent parts of Cleveland that were excluded from federal housing loans because Black Americans were the predominant inhabitants of these areas. But, in fact, the map shows all the things I mentioned – all the legacy of historic racial discrimination.

Race and protected class characteristics are carved out regardless of actuarial fairness because there is a history of discrimination that, at best, has left a legacy of outcomes that are embedded in the data used for actuarial analysis and, at worst, continues today with racist practices – intentional or unintentional – unrelated to risk or cost of insurance.

Protected class unfair discrimination in insurance recognizes that historical discrimination has long-lasting effects that have disadvantaged these groups. The shorter life expectancy of Black Americans is not caused by their skin color, but by the historical and ongoing discrimination in housing, health care, policing and other parts of our lives.

What is Structural 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

What is Structural Racism?

Structural racism affects every institution in our society – including insurance.

In many or most instances, the racial bias is unrecognized and unintentional – **“intent” cannot be a determining factor.**

Structural Racism – Current Impacts

"Persistence of mortgage lending bias in the United States: 80 years after the Home Owners' Loan Corporation security maps"

https://pubmed.ncbi.nlm.nih.gov/35992214/

"We compare historic HOLC grades and contemporary levels of mortgage lending bias using spatially detailed HMDA data. We further examine the relationship between HOLC risk grades and contemporary racial and ethnic settlement patterns. **Results suggest that historical mortgage lending risk categorizations and settlement patterns are associated with contemporary mortgage lending bias and racial and ethnic settlement patterns."**

Why are public health scholars studying mortgage discrimination?

Structural Racism – Current Impacts

"The legacy of structural racism: Associations between historic redlining, current mortgage lending, and health"

https://pubmed.ncbi.nlm.nih.gov/33997243/

Structural racism, which is embedded in past and present operations of the U.S. housing market, is a fundamental cause of racial health inequities. Using historic redlining score and current lending discrimination, we created a 4-level hierarchical measure of lending trajectory.

Our findings illustrate ongoing legacies of government sponsored historic redlining. Structural racism, as manifested in historic and current forms of lending disinvestment, predicts poor health in Milwaukee's hypersegregated neighborhoods. We endorse equity focused policies that dismantle and repair the ways racism is entrenched in America's social fabric.

Structural Racism – Current Impacts

Racial Discrimination in Property Appraisals

https://therealdeal.com/2022/08/20/gut-punch-black-redlining-expert-faces-appraisal-bias/

https://www.nytimes.com/2022/08/18/realestate/housing-discrimination-maryland.html

Black professional couple aimed to refinance their mortgage last summer to take advantage of the low interest rates. They paid $450,000 for the home in 2017 and had since made $42,000 in renovations and upgrades.

Lanham, the owner of 20/20 Valuations, appraised the home for just $472,000, much less than the couple expected. The result: loanDepot denied a refinance loan.

Months later, the couple applied for another loan, this time with Swift Home Loans, but made a few changes. They removed photographs of their family and Connolly's white colleague from the university posed as the homeowner.

It was appraised for $750,000 – almost 60 percent more.

(97% of appraisers are white)

Structural Racism – Current Impacts

Environmental Racism

https://www.wilderness.org/articles/blog/people-color-have-been-left-out-nature-equitable-transit-can-help

74 percent of communities of color **in the contiguous United States live in nature-deprived areas,** compared with just 23 percent of white communities.

Studies show historically redlined areas are more likely to have fewer green spaces, and that **formerly redlined areas could be 12.6 degrees Fahrenheit hotter than non-redlined ones,** due to a lack of vital heat-reducing tree canopy.

People living in nature-deprived areas are also **exposed to more pollution** and are more likely to develop respiratory problems such as asthma.

Helping connect communities of color to nature is a critical climate and environmental justice issue. These populations face acute health problems from disproportionate pollution exposure. Recent studies show people of color are exposed to more pollution than white people from power plants, industry, traffic and other sources.

Structural Racism – Current Impacts

Environmental Racism

"Historical Redlining Is Associated with Present-Day Air Pollution Disparities in U.S. Cities"

- https://pubs.acs.org/doi/10.1021/acs.estlett.1c01012

Communities of color in the United States are systematically exposed to higher levels of air pollution. We explore here how redlining, a discriminatory mortgage appraisal practice from the 1930s by the federal Home Owners' Loan Corporation (HOLC), relates to present-day intraurban air pollution disparities in 202 U.S. cities. We find that pollution levels have a consistent and nearly monotonic association with HOLC grade . . . within each HOLC grade, racial and ethnic air pollution exposure disparities persist, indicating that redlining was only one of the many racially discriminatory policies that impacted communities. **Our findings illustrate how redlining, a nearly 80-year-old racially discriminatory policy, continues to shape systemic environmental exposure disparities in the United States.**

Modern air pollution disparities in historically redlined areas

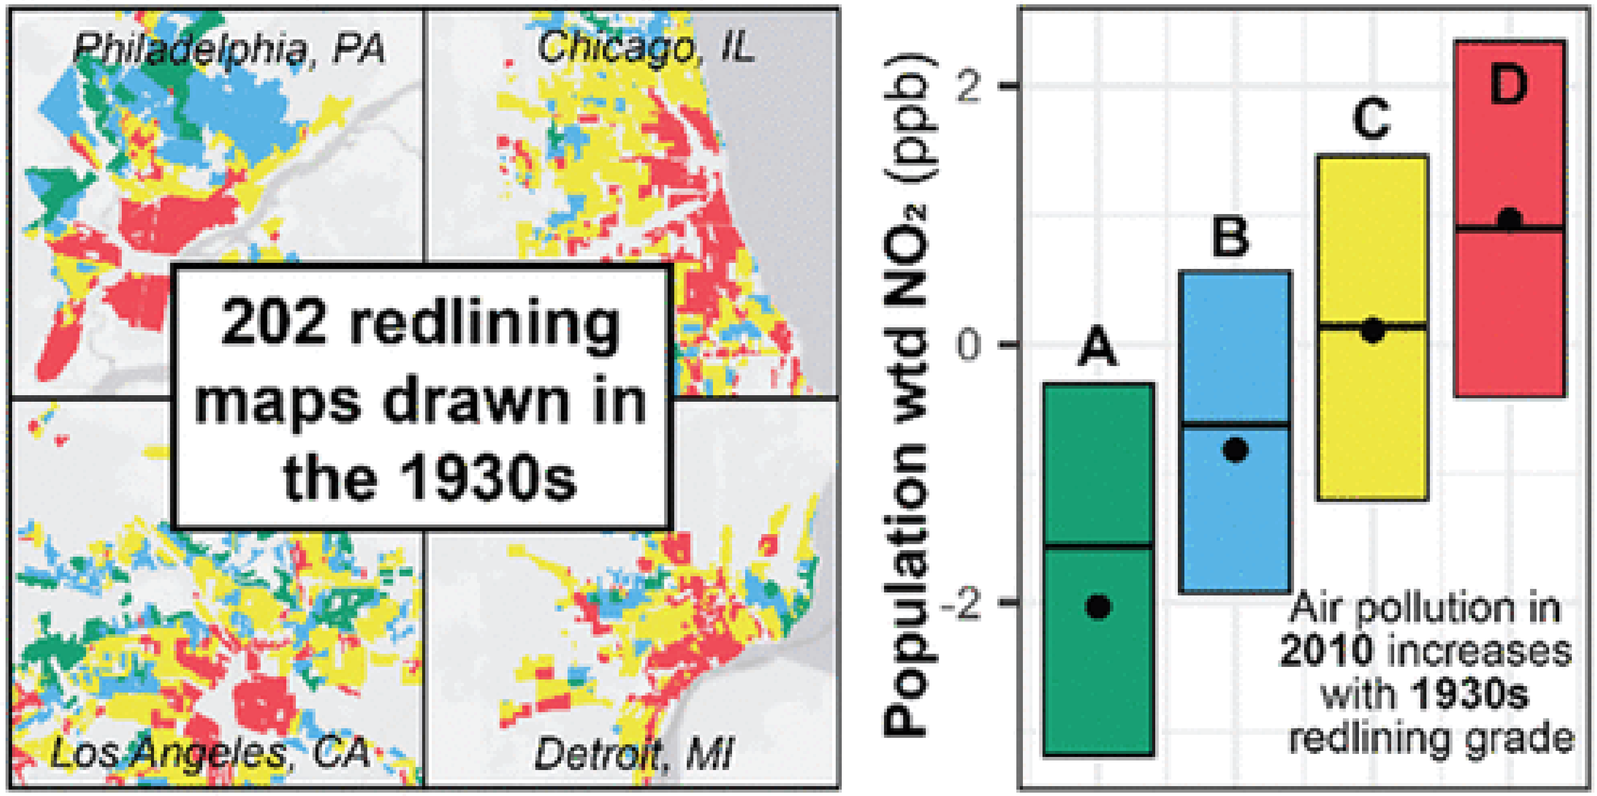

Structural Racism – Current Impacts

New Study Finds Historically Redlined Communities at a Higher Risk of Flooding”

https://nlihc.org/resource/new-study-finds-historically-redlined-communities-higher-risk-flooding

The study found that across 38 metro areas in the nation more than $107 billion worth of homes at a high risk for flooding were found in redlined communities. This amount is 25% more than in non-redlined communities. People of color are more likely to live in historically redlined communities. In fact, 58% of households in formerly redlined communities are nonwhite, compared to 40.4% of households in neighborhoods that were labeled “desirable” during the New Deal era. **The study further highlights the impact that past discriminatory housing policies continue to have on communities of color today and how housing justice is deeply connected to climate justice.**

Structural Racism – Current Impacts

Biased Policing – US DOJ Investigation into Policing in Ferguson, MO

"Ferguson's approach to law enforcement both reflects and reinforces racial bias, including stereotyping. ***The harms of Ferguson's police and court practices are borne disproportionately by African Americans, and there is evidence that this is due in part to intentional discrimination on the basis of race.***

Our investigation indicates that this disproportionate burden on African Americans cannot be explained by any difference in the rate at which people of different races violate the law. Rather, our investigation has revealed that these disparities occur, at least in part, because of unlawful bias against and stereotypes about African Americans."

Algorithms Learn the Bias in Data and Modelers

"Advocates of algorithmic techniques like data mining argue that they eliminate human biases from the decision-making process. ***But an algorithm is only as good as the data it works with.*** *Data mining can inherit the prejudices of prior decision-makers or reflect the widespread biases that persist in society at large.* ***Often, the "patterns" it discovers are simply preexisting societal patterns of inequality and exclusion. Unthinking reliance on data mining can deny members of vulnerable groups full participation in society."*** *(Barocs and Selbst)*

The fact that an insurer doesn't use race in an algorithm does not logically or factually result in no discrimination on the basis of race.

In fact, the only way to identify and eliminate the impacts of structural racism in insurance is to measure that impact by explicit consideration of race and other protected class factors**.**

Explicit Consideration of Race Needed to Test for Racial Bias

New York Times, August 10, 2015: Algorithms and Bias: Q. and A. With Cynthia Dwork

Q: Some people have argued that algorithms eliminate discrimination because they make decisions based on data, free of human bias. Others say algorithms reflect and perpetuate human biases. What do you think?

A: Algorithms do not automatically eliminate bias. . . .Historical biases in the . . .data will be learned by the algorithm, and past discrimination will lead to future discrimination.

Fairness means that similar people are treated similarly. ***A true understanding of who should be considered similar for a particular classification task requires knowledge of sensitive attributes, and removing those attributes from consideration can introduce unfairness and harm utility.***

Addressing Structural Racism in Insurance

Precise, Common Definitions of What We're Trying to Measure Critical

Disparate Intent: Intentional Use of Race

Proxy Discrimination: Disproportionate Racial Outcomes Tied to Use of Proxies for Race, Not to Outcomes

Disparate Impact: Disproportionate Racial Outcomes Tied to Historic Discrimination and Embedded in Insurance Outcomes

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 not only empirical analysis and metrics but public policy considerations for fair trade-offs between algorithmic efficiencies and reducing racially-biased outcomes.

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

Fortunately, there are statistical techniques – some that have been used for over 40 years – 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

Testing for racial bias shows where on the continuum from pure proxy to no racial bias the algorithm lies. Whether there is proxy discrimination or disparate impact would then depend on the quantitative definitions / thresholds established for proxy discrimination and disparate impact

Holistic Approach: 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 guidance for bias thresholds and equity trade-offs.
5. *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.*
6. *Modernizing definitions of and oversight of advisory organizations.*

Insurer Testing of Algorithms and Actual Consumer Outcomes

To their credit, the ACLI has engaged constructively regarding the implementation of Colorado SB 169. ACLI has 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.

Common Insurance Outcome Metrics Permit Common Testing Methods and Metrics

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?

Testing of Algorithms for All Phases of Insurance Life-Cycle

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."

Testing of Marketing Algorithms (con't)

- "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."

2005 was the stone age of big data analytics. 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, he 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"

The advances in algorithmic micro-targeting consumers in real time have advanced significantly since 2017.

Testing of Pricing Algorithms

Two Tests Needed:

- Testing for Actuarial Soundness Using Loss Costs as Outcome Variable
- Testing Using Premium as Outcome Variable

The first test addresses the actuarial fairness prong of unfair discrimination as well as for proxy discrimination and disparate impact in predictive factors.

The second test addresses the potential bias introduced when insurers move from indicated rates to selected rates and, for example, whether price optimization or consumer lifetime value scores introduce racial bias.

Testing of Pricing Algorithms

Testing of premium outcomes must include application data, not just data on policies issued. Inclusion of final quotes that didn't result in a policy are essential for at least two reasons. First, you can't meaningfully analyze bias in outcomes if you don't measure folks priced out of the market. The inclusion of data on all applications has been part of the Home Mortgage Disclosure Act data reporting from inception. Second, analyzing quote data will reveal the outcomes of insurers' marketing algorithms and, consequently, whether those marketing algorithms produce racially-biased outcomes

Holistic Testing of Algorithms, 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 correlations with race and reveal the factor's true contribution to explaining the insurance outcome and provide a statistical basis for addressing disparate impact.

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?

A Note on Risk-Based Pricing

The P/C trades and the III routinely seek to justify pricing freedom – the use of any data source or characteristic of the consumer, vehicle, property, built or natural environment if it is predictive of risk and more refined the risk prediction is always better and more fair. This formulation is problematic for a number of reasons, but for our purposes today, there are a couple of key points.

- The purpose of insurance is to create a risk pool through which individuals can transfer risk to that pool. Risk-based pricing is a means to manage that mechanism safely and fairly, but it is not the purpose of insurance.
- Testing for and addressing structural racism in insurance is 100% consistent with and improves risk- and cost-based practices.
- Unfettered risk-based pricing without attention to structural racism will reflect and perpetuate historic discrimination.

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 – a**nalysis of application data is essential to see if those algorithms systematically deny communities of color such access**

Univariate vs. Multivariate Analysis and the Problem of Multicollinearity

Basic stuff for the technicians, bear with me to bring the others along.

Let's say we price auto insurance using age, gender and credit score. If we analyze each factor's relationship to expected costs separately, the results for each factor are sure to be capturing some the effects of the other factors because there is correlation among the predictive factors – multicollinearity.

"Multicollinearity occurs when two or more independent variables are highly correlated and could therefore obfuscate each variable's true relationship with the target. "

CAS Research Paper, "Defining Discrimination in Insurance"

Testing for Proxy Discrimination and Disparate Impact: A Natural Extension of Current 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 Xs 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 Models

Suppose an insurer wants 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 <u>control variables.</u>

$$b_0 + b_1X_1 + b_2X_2 + b_3X_3 + \mathbf{b_4C_1} + e = y$$

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.

Using Race (or Other Protected Class Characteristic) as a Control Variable

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

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

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

This is one of many statistical techniques used is disparate impact analysis that all seek to remove the correlations between predictive variables and race.

How Do We Interpret the Testing Results?

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

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 Testing Results?

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

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 Testing Results?

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

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 Testing Results?

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

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 Applicants’ Race – How Can An Analyst Test for Racial Bias?

- Assign a racial characteristic to an individual based on racial characteristic of a small geographic area – Census data at the census block level.
- Utilize the Bayesian Improved Surname Geocoding Method (BISG), based on census geography and surname data or Bayesian Improved First Name Surname Geocoding (BIFSG).
- Reach out to data brokers and vendors for a new data service.

Just as insurers and vendors of algorithms work with individual transaction data, testing for racial bias in algorithms requires use of transaction data with consumer personal information

NAIC Efforts to Address Racial Bias

- Principles on Artificial Intelligence
- Special Committee on Race and Insurance
- Innovation, Cybersecurity and Technology Committee
  - Collaboration Forum on Algorithmic Bias
- Big Data Working Group
- Accelerated Underwriting Working Group
- Privacy Protections Working Group
- Casualty Actuarial Task Force
- Antifraud Task Force
- Producer Licensing Task Force
- Several States have also taken actions to address algorithmic and racial bias in insurance.

Inferring Race -- Resources

“Imperfect Inferences: A Practical Assessment”

https://dl.acm.org/doi/fullHtml/10.1145/3531146.3533140

“Using First Name Information to Improve Race and Ethnicity Classification”

https://www.researchgate.net/publication/322670070_Using_First_Name_Information_to_Improve_Race_and_Ethnicity_Classification

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/

NAIC Principles on Artificial Intelligence

Insurance-specific AI applications should be:

- Fair and Ethical
- Accountable
- Compliant
- Transparent
- Secure, Safe and Robust

Consistent with the risk-based foundation of insurance, AI actors should proactively engage in responsible stewardship of trustworthy AI in pursuit of beneficial outcomes for consumers and to avoid proxy discrimination against protected classes.

Disparate Impact & Proxy Testing

Disparate impact theory <> Proxy testing techniques <> Inferential analysis for protected data

Presentation by:
Morgan Williams, General Counsel &
Michael Akinwumi, Chief Tech Equity Officer
National Fair Housing Alliance

Civil Rights Liability: basic causes of action

- **Disparate treatment**: occurs when an entity explicitly or intentionally treats people differently based on protected characteristics, such as race, national origin, or sex.

- **Disparate impact**: in contrast, is focused on outcomes.
  1. Facially neutral policy or practice disproportionately adversely impacts members of protected classes,
  2. The policy or practice does not advance a legitimate business justification, or
  3. The policy is not the least discriminatory means to advance that interest.

© 2022 National Fair Housing Alliance. Please do not distribute.

Insurance Civil Rights Cases: NFHA, et al. vs Prudential Insurance

- Disparate Challenged practices included minimum underwriting requirements:
  1. Age of the home
  2. Market value of the home
  3. Difference between the replacement cost and the market value

Favorable Motion to Dismiss decision, Aug. 29, 2002.

Additional information: https://casetext.com/case/national-fair-housing-alli-v-prudential-ins-co

© 2022 National Fair Housing Alliance. Please do not distribute.

Insurance Civil Rights Cases: NFHA vs Travelers Insurance

- Challenged practices included
  - Policies that denied commercial habitational insurance to landlords who rented to voucher recipients

Favorable Motion to Dismiss decision, Aug. 21, 2017.

Additional Information: https://nationalfairhousing.org/travelers/

© 2022 National Fair Housing Alliance. Please do not distribute.

Civil Rights Testing Methods

1. Ensuring that models do not include protected characteristics or close proxies for protected characteristics, for example as variables or segmentations; and

2. Assessing whether facially-neutral models are likely to disproportionately lead to negative outcomes for a protected class, and if such negative impacts exist, ensuring the models serve legitimate business needs and evaluating whether changes to the models would result in less of a disparate impact while maintaining model performance.

© 2022 National Fair Housing Alliance. Please do not distribute.

Civil Rights Compliance Controls

Compliance management systems may include:

1. Board and management oversight;
2. Model governance;
3. Reviewing policies and procedures within which models operate, including underwriting policies, overlays, exclusions, overrides and the likes;
4. Assessing areas of discretion;
5. Providing fair lending training;
6. Ensuring teams have diverse backgrounds;
7. Effective monitoring; and
8. Consumer complaint resolution processes.

© 2022 National Fair Housing Alliance. Please do not distribute.

Proxy Testing: control variable

- Consider an ordinary regression model with a target variable y_i and a set of explanatory variables x_{ij} for each record $\mathbf{i}$ such that

$$y_i = \sum_{j=1}^{k} \beta_j x_{ij} + \varepsilon_i \quad (1)$$

- Consider a ratemaking application that uses a GLM (Generalized Linear Model) as a use case.
  1. systematic component is determined by all the $\mathbf{k}$ β_j post variable selection and multicollinearity checks
  2. **For all $\mathbf{i}$**, suppose each x_{ij} is a statistically significant explanatory variable for y_i, and that
  3. Each β_j is the contribution of an x_{ij} to the observed variation in y_i
- An insurer may want to mitigate effects like demographics or jurisprudence across state boundaries.
- This may be achieved with control variables x_{ij} , where j is from $k+1$ to n ,

$$y_i = \sum_{j=1}^{k} \beta_j x_{ij} + \sum_{j=k+1}^{n} \beta_j x_{ij} + \varepsilon_i \quad (2)$$

© 2022 National Fair Housing Alliance. Please do not distribute.

Proxy Testing: using control variables to mitigate discrimination

- Under the Fair Housing Act and Equal Credit Opportunity Act, some of the legally protected variables are race, color, national origin or disability.
- While each of these prohibited variables may not be directly used in, say, a GLM, not checking if any of the $x_{.j}$ is a potential proxy for the variables may lead to proxy discrimination or disparate impact.
- Suppose we want to control for Race effect in equation (1). This may be achieved in equation (3) with Race represented by $x_{.(k+1)}$

$$y_i = \sum_{j=1}^{k} \beta_j x_{ij} + \beta_{k+1} x_{i(k+1)} + \varepsilon_i \qquad (3)$$

- Consider equation (3)
  - Discrimination: suppose a clustering algorithm is trained on all the $\mathrm{k}\ x_{ij}$ to identify subgroups within records $\mathrm{i} = 1, \cdots, n$. Let's assume m subgroups were identified. If any of the $\mathrm{k}\ x_{ij}$ proxies for Race within one of the similarly situated m subgroups, this may lead to proxy discrimination or disparate treatment within that subgroup.

© 2022 National Fair Housing Alliance. Please do not distribute.

Proxy Testing: result interpretation and recommended action

There are four possible outcomes when proxy testing is conducted with the statistical approach described in

$$y_i = \sum_{j=1}^{k} \beta_j x_{ij} + \beta_{k+1} x_{i(k+1)} + \varepsilon_i$$

1. Neither proxy discrimination (PD) nor disparate impact (DI) is established:

   $x_{.(k+1)}$ is not statistically significant and there is marginal change in β_j for all $j = 1, \cdots, k$.

2. Proxy discrimination (PD) is established with no disparate impact (DI):

   $x_{.(k+1)}$ is statistically significant and statistical significance has been lost in at least one β_j where $j = 1, \cdots, k$.

PD	DI	
	yes	no
yes	May remove proxy/proxies and retest until model is proxy free NFHA	may remove proxy/proxies, and rebuild model
no	1. Are the variables business justified and statistically valid? i. If no, may be illegal. May not use model. ii. If yes, go to the next question. 2. Is there at least one less discriminatory alternative (LDA) with equal or better or acceptable performance as the tested model? i. If yes, it may be unlawful to not use the LDA-searched model	may use the model

© 2022 National Fair Housing Alliance. Please do not distribute.

Proxy Testing: result interpretation and recommended action

There are four possible outcomes when proxy testing is conducted with the statistical approach described in

$$y_i = \sum_{j=1}^{k} \beta_j x_{ij} + \beta_{k+1} x_{i(k+1)} + \varepsilon_i$$

3. No proxy discrimination (PD) but disparate impact (DI) is established:

   $x_{.(k+1)}$ is statistically significant with observed large disparities in y , and there is marginal change in β_j with no loss in statistical significance for all $j = 1, \cdots, k$.

4. Proxy discrimination (PD) and disparate impact (DI) are established:

   $x_{.(k+1)}$ is statistically significant with observed large disparities in y , and there is a material change in β_j with no loss in statistical significance for all $j = 1, \cdots, k$.

PD	DI	
	yes	no
yes	May remove proxy/proxies and retest until model is proxy free NFHA	may remove proxy/proxies, and rebuild model
no	1. Are the variables business justified and statistically valid? i. If no, may be illegal. May not use model. ii. If yes, go to the next question. 2. Is there at least one less discriminatory alternative (LDA) with equal or better or acceptable performance as the tested model? i. If yes, it may be unlawful to not use the LDA-searched model	may use the model

© 2022 National Fair Housing Alliance. Please do not distribute.

Inference for Protected Data: BISG and its variants

- Some of the legally protected variables such as race, color, national origin or disability required to conduct proxy tests may not have been collected. For example, applicants' race is not mandatory for an insurance underwriting.
- Statistical methods are often the go-to solutions when protected data are missing at insurance application time.
- Bayesian Improved Surname Geocoding (BISG) and its variants have become popular.

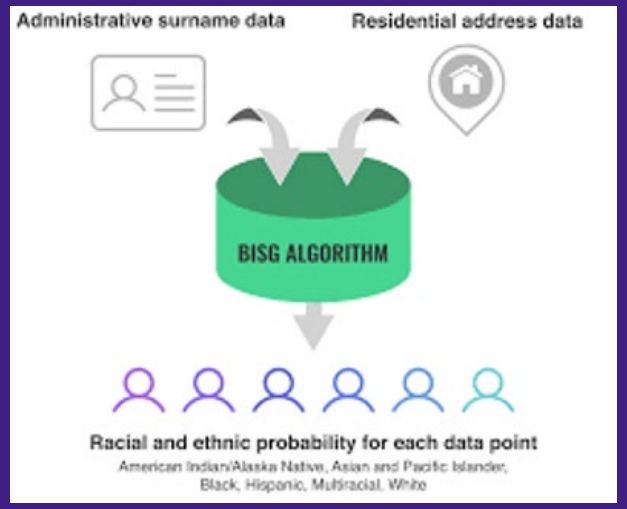

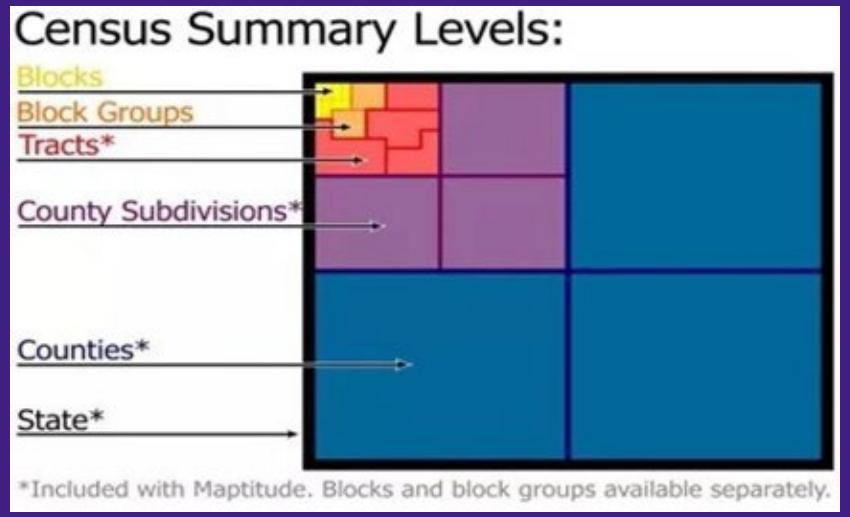

Bayesian indirect surname Geocoding (BISG). RAND Corporation. (n.d.). Retrieved March 15, 2022, from https://www.rand.org/health-care/tools-methods/bisg.html#accuracy.

What are blocks? Caliper Mapping Software and Transportation Software. (n.d.). Retrieved March 23, 2022, from https://www.caliper.com/maptitude/census2000data/summarylevels.htm

© 2022 National Fair Housing Alliance. Please do not distribute.

Inference for Protected Data: machine learning techniques

- Imageomics is a branch of science that attempts to infer biological traits and biometrics from a person's image
- There are Machine Learning techniques that use computer vision algorithms to infer gender and age from a person's image.
- In 2015 researchers from The Open University of Israel published research work that used CNN (Convolutional Neural Network) to infer a person's age and gender from their images.

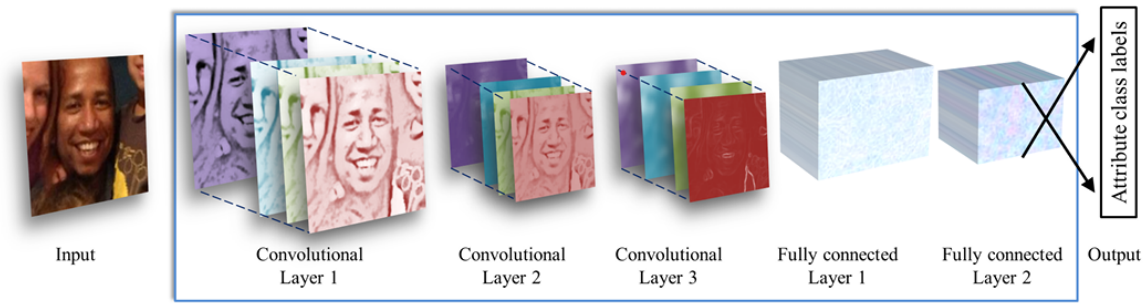

Figure 2. **Illustration of our CNN architecture.** The network contains three convolutional layers, each followed by a rectified linear operation and pooling layer. The first two layers also follow normalization using local response normalization [28]. The first Convolutional Layer contains 96 filters of 7×7 pixels, the second Convolutional Layer contains 256 filters of 5×5 pixels, The third and final Convolutional Layer contains 384 filters of 3×3 pixels. Finally, two fully-connected layers are added, each containing 512 neurons. See Figure 3 for a detailed schematic view and the text for more information.

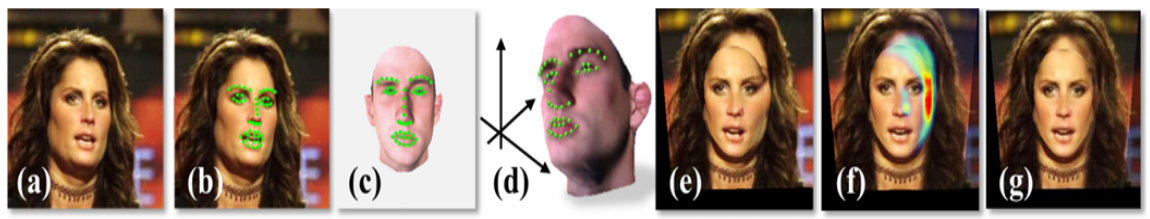

Figure 2: Frontalization process overview. (a) Query photo; (b) facial feature detections; (c) the same detector used to localize the same facial features in a reference face photo, produced by rendering a textured 3D computer graphics model (d); (e) from the 2D coordinates on the query and their corresponding 3D coordinates on the model we estimate a projection matrix which is then used to back-project query intensities to the reference coordinate system; (f) estimated visibility due to non-frontal poses, overlaid on the frontalized result. Warmer colors reflect less visible pixels. Facial appearance in these regions is produced by borrowing colors from corresponding symmetric parts of the face; (g) our final frontalized result.

Gil Levi and Tal Hassner at the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), Boston, June 2015, https://talhassner.github.io/home/projects/cnn_agegender/CVPR2015_CNN_AgeGenderEstimation.pdf Tal Hassner, Shai Harel, Eran Paz and Roee Enbar, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), Boston, June 2015, https://talhassner.github.io/home/projects/frontalize/CVPR2015_frontalize.pdf

© 2022 National Fair Housing Alliance. Please do not distribute.

Q & A

Conclusion

Discussion and Feedback

Receive Comments from Interested Parties