Disparate Impact: The Impact of the Social Justice Movement on Insurance Rating

CAS 2021 Ratemaking, Product and Modeling Seminar

March 16th, 2021

CAS Approach to Race and Insurance Pricing

Poll Question 1

How would you describe your race?

- Non-Hispanic White
- Hispanic or Latino/a/x
- Black or African American
- Asian or Asian American
- Other

Poll Question 2

Have you experienced discrimination in any context that you believe was based on your race or skin color?

- Yes
- No

Poll Question 3

How old are you?

- Under 25
- 25 to 34
- 35 to 49
- 50 to 64
- Over 65

Addressing Systemic Racism in Insurance

Presentation to Casualty Actuarial Society

Ratemaking and Product Management Seminar

March 16, 2021

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 is also a member of the U.S. Federal Reserve Board's Insurance Policy Advisory Committee.

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.

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.

What Information Does This Map of Omaha Nebraska Present?

a. Concentration of Minority Population

b. Home Insurance Rating Territories

c. Rates of COVID Infections and Deaths

d. Federal Home Loan Eligibility 1930’s to 1960’s

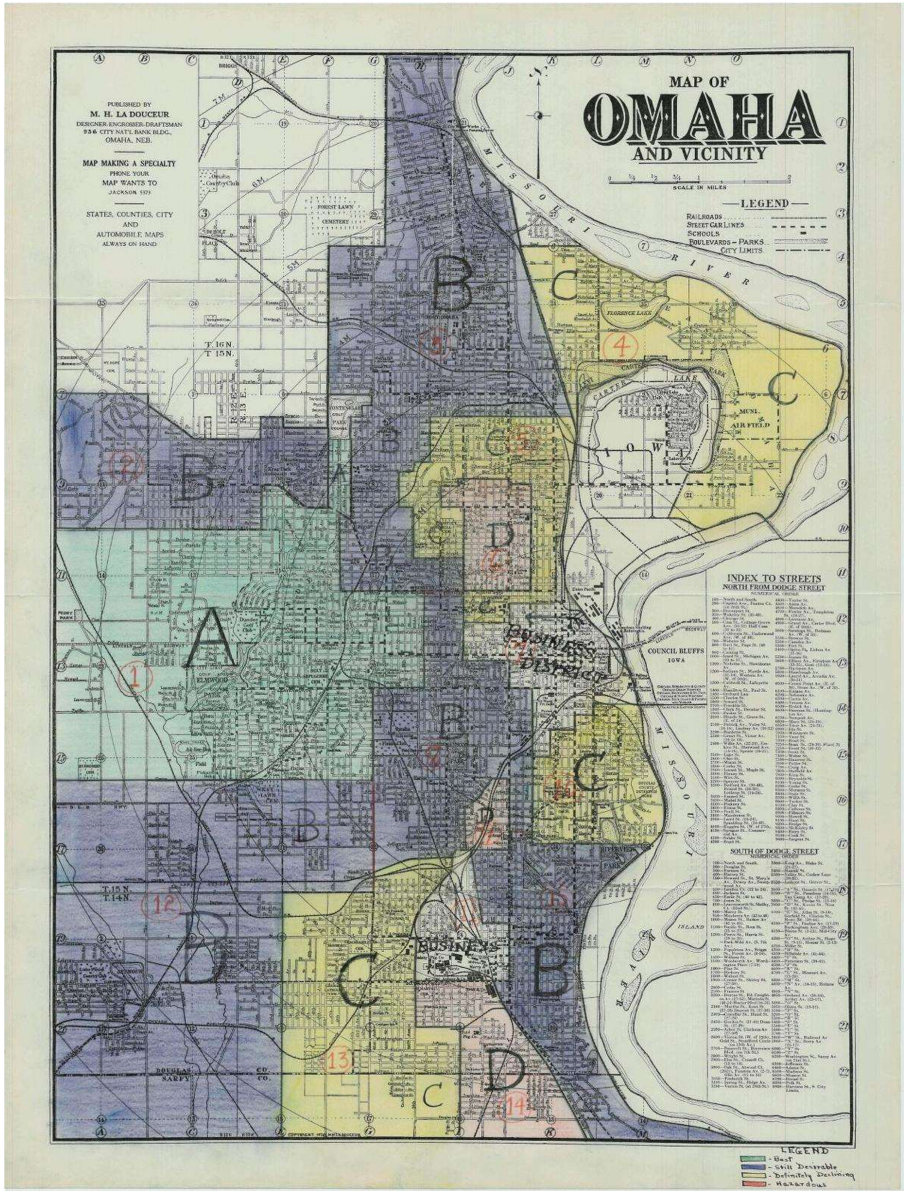

Birny Birnbaum
Center for Economic Justice

CAS RPM
March 16, 2021

What is Systemic Racism and Inherent Bias?

"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." Those are the words of Kirt Walker, Chief Executive Officer of Nationwide.

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. It also requires people of privilege—white people—to stand up for and stand with our communities like we never have before," Those are the words of Jack Salzwedel, the CEO of American Family.

Why Do State and Federal Laws Prohibition Discrimination on the Basis of Race?

Justice Kennedy for the Majority in the U.S. Supreme Court's 2015 *Inclusive Communities* Opinion upholding disparate impact as unfair discrimination under the Fair Housing Act.

Recognition of disparate-impact claims is also consistent with the central purpose of the FHA, which, like Title VII and the ADEA, was enacted to eradicate discriminatory practices within a sector of the Nation's economy.

Recognition of disparate-impact liability under the FHA plays an important role in uncovering discriminatory intent: it permits plaintiffs to counteract unconscious prejudices and disguised animus that escape easy classification as disparate treatment.

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 Systemic Racism Manifest Itself in Insurance – Whether for Marketing, Pricing or Claims Settlement?

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

Today's focus is on number 3 – practices that actuaries can identify and address.

Fair and Unfair Discrimination in Insurance

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

For life insurance, we look to the UTPAs and 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.

From the NAIC UTPA Section 4 G Unfair Discrimination

Actuarial Unfair Discrimination

(1) Making or permitting any unfair discrimination between individuals of the same class and equal expectation of life in the rates charged for any life insurance policy or annuity or in the dividends or other benefits payable thereon, or in any other of the terms and conditions of such policy.

Protected Class Unfair Discrimination

(6) Refusing to insure, refusing to continue to insure, or limiting the amount of coverage available to an individual because of the sex, marital status, race, religion or national origin of the individual;

Poll Question 4

Which best describes your experience?

a. My company has a policy to examine the development and impact of our algorithms for racial bias

b. My company has been receptive to suggestions to examine the development and impact of our algorithms for racial bias

c. My company has not been receptive to suggestions to examine the development and impact of our algorithms for racial bias

d. I would not suggest examining the development and impact of our algorithms for racial bias because of the reaction I think such a suggestion would create.

e. My company doesn't use, collect or consider racial characteristics so there is no need to test our practices for racial impact.

Correlation is Not the Standard for Fair Actuarial Discrimination

Statutes and actuarial standards don't refer to correlation, but demand a more robust relationship. Why? Here's an example of an almost perfect correlation – over 99%.

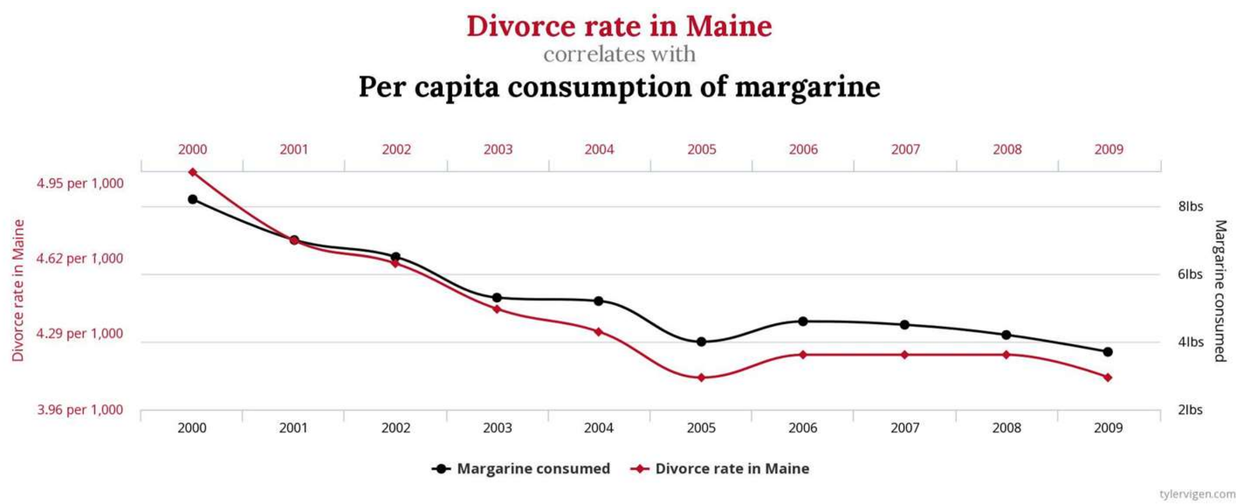

Why isn't a simple correlation relied upon or sufficient?

Because a predictive characteristic (or variable) may not be correlated in whole or in part to the outcome, but may also be correlated to other predictive variables.

Consider the difference between an outcome – say, mortality – and one predictive variable versus an outcome and multiple predictive variables.

Age to Mortality, Gender to Mortality, Tobacco Use to Mortality

Each of these represents a one-to-one – or univariate – relationship. But each predictive variable may be replicating part of another variable because of correlation between the predictive variables. Tobacco Use may be correlated with age or gender.

Eliminating Correlation among Predictive Variables:
Multi-variate Analysis

The issue of correlation among predictive variables has become more important in life insurance as insurers have started to use new data and predictive variables.

Over the last 30 years, insurers and actuaries have developed new techniques to address the problems with univariate analysis.

Insurers use a variety of techniques to eliminate correlations <u>among predictive variables</u> in order to isolate each individual predictive variable's unique contribution to explaining the outcome.

How Does Multi-Variate Analysis Work?

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 – mortality (or predicting the decision produced by traditional underwriting).

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 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, different frequencies of accidental accidents or differences in jurisprudence. An insurer would add one or more control variables.

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

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.

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:

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

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

What if X1 is a perfect proxy for Race?

Then once we add the control variable for Race, X1 no longer has any predictive value because all it was doing was predicting race, not the outcome y.

What if X1 is both predictive of mortality and correlated to Race? Then, the model still shows X1's (now different) predictive value, but shorn of its correlation to Race, leaving the unique contribution of X1 to explaining mortality.

Disparate Impact Analysis Improves Cost-Based Pricing

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.

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

What if X1, X2 and X3 are not perfect proxies for Race, but still have high correlation? Then, the disparate impact analysis – and our simple model – removes that correlation and the remaining values for b1, b2 and b3 are the unique contributions of each predictive variable to explaining the outcome. The result is more – not less – accurate cost-based or risk-based analysis.

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. From Barocas and Selbst:

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.***

Consider Criminal History Scores

"TransUnion recently evaluated the predictive power of court record violation data (including criminal and traffic violations)

"Also, as court records are created when the initial citation is issued, they provide insight into violations beyond those that ultimately end up on the MVR—such as violation dismissals, violation downgrades, and pre-adjudicated or open tickets."

What is the likelihood that TU Criminal History Scores have a disparate impact against African-Americans? Consider policing records in Ferguson, Missouri.

US DOJ Investigation of the Ferguson Police Department

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.***

Ferguson’s law enforcement practices overwhelmingly impact African Americans. Data collected by the Ferguson Police Department from 2012 to 2014 shows that African Americans account for 85% of vehicle stops, 90% of citations, and 93% of arrests made by FPD officers, despite comprising only 67% of Ferguson’s population.

US DOJ Investigation of the Ferguson Police Department (2)

FPD appears to bring certain offenses almost exclusively against African Americans. For example, from 2011 to 2013, African Americans accounted for 95% of Manner of Walking in Roadway charges, and 94% of all Failure to Comply charges.

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.

Why is it Reasonable and Necessary to Require Insurers to Test for and Minimize Disparate Impact?

Insurer practices and algorithms do not necessarily use expected claims as the outcome variable. Sometimes the desired outcome is based on non-cost factors and these non-cost factors has disproportionate impact on communities of color.

In 2005, then CEO of Allstate, Ed Liddy told investment analysts about how credit scoring was helping Allstate avoid the wrong customers:[1]

> 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. That drives growth on both the top and bottom line.
>
> This year, we've expanded from 7 basic price levels to 384 potential price levels in our auto business.

---

[1] Transcript of Presentation to Edward M. Liddy, Chairman and CEO, The Allstate Corporation Twenty-First Annual Strategic Decisions Conference, Sanford C. Bernstein & Co., June 2, 2005.

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.

These statements were made in the Stone Age of Big Data – 2005. Since then, insurers' use of new, bigger and more granular personal consumer data has exploded.

Practices That Raise Concerns About 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.

What are the Benefits and Costs of Requiring Insurers to Test For and Minimize Disparate Impact?

If racial and economic justice are a priority, if cost-based insurer practices are a priority, if closing the protection gap and making insurance more affordable and available in traditionally underserved communities is a priority, then the benefits of requiring insurers to test for and minimize disparate impact far, far outweigh the costs.

While there are examples of disparate impact claims brought against insurers under the federal Fair Housing Act that have resulted in improved risk-based pricing and improved insurance availability in communities of color – e.g., challenges against underwriting based on age and value of the home – industry has not been able to cite a single example of a successful disparate impact claim that has harmed risk-based pricing.

Why Do Efforts to Address Discrimination on the Basis of Race Require Explicit Consideration of Race?

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.***

Steve Bellovin, "Yes, 'algorithms' can be biased. Here's why. A computer scientist weighs in on the downsides of AI."[2]

This is what's important: machine-learning systems—"algorithms"—produce outputs that reflect the training data over time. If the inputs are biased (in the *mathematical* sense of the word), the outputs will be, too. Often, this will reflect what I will call "sociological biases" around things like race, gender, and class.

One thing is to exercise far more care in the selection of training data. Failure to do that was the likely root cause of Google Images labeling two African-Americans as gorillas. Sometimes, fixing the training data can help.

Of course, this assumes that developers are even aware of the bias problem. Thus, another thing to do is to test for biased outputs—and some sensitive areas, such as the criminal justice system, simply do not use these kinds of tools.

---

[2] https://arstechnica.com/tech-policy/2019/01/yes-algorithms-can-be-biased-heres-why/

There are several reasons to be wary of the "algorithmic" approach. One reason is that people put too much trust in computer output. Every beginning programmer is taught the acronym "GIGO:" garbage in, garbage out. To end users, though, it's often "garbage in, gospel out"—if the computer said it, it must be so. (This tendency is exacerbated by bad user interfaces that make overriding the computer's recommendation difficult or impossible.) We should thus demand less bias from computerized systems precisely to compensate for their perceived greater veracity.

The second reason for caution is that computers are capable of doing things—even bad things—at scale. There is at least the perceived risk that, say, computerized facial recognition will be used for mass surveillance. Imagine the consequences if a biased but automated system differentially misidentified African-Americans as wanted criminals. Humans are biased, too, but they can't make nearly as many errors per second.

Our test, then, should be one called disparate impact. "Algorithmic" systems should be evaluated for bias, and their deployment should be guided appropriately. Furthermore, the more serious the consequences, the higher the standard should be before use.

"The Real Reason Tech Struggles with Algorithmic Bias"[3]

These are mistakes made while trying to do the right thing. But they demonstrate why tasking untrained engineers and data scientists with correcting bias is, at the broader level, naïve, and at a leadership level insincere.

No matter how trained or skilled you may be, it is 100 percent human to rely on cognitive bias to make decisions. Daniel Khaneman's work challenging the assumptions of human rationality, among other theories of behavioral economics and heuristics, drives home the point that human beings cannot overcome all forms of bias. But slowing down and learning what those traps are—as well as how to recognize and challenge them—is critical. As humans continue to train models on everything from stopping hate speech online to labeling political advertising to more fair and equitable hiring and promotion practices, such work is crucial.

[3] Yael Eisenstat at https://www.wired.com/story/the-real-reason-tech-struggles-with-algorithmic-bias/

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. [4]

3. Reach out to data brokers and vendors for a new data service.

---

[4] 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:Ablueprintforahumanrights,socialandethicalimpactassessmentAlessandroMantelero

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

Disparate Impact – The Impact of the Social Justice Movement on Insurance Rating

March 16, 2021

Roosevelt C. Mosley, Jr., FCAS, CSPA, MAAA

Principal & Consulting Actuary

Poll Question 5

In my current job, with respect to insurance pricing, I have…

1. Started examining issues of racial bias
2. Not started, but have considered examining issues of racial bias
3. Not yet considered examining issues of racial bias

The Impact of the Social Justice Movement on Insurance Rating

- Guiding principle for ratemaking
- How is unfairly discriminatory defined?
- Implications for insurance ratemaking

Guiding Principle for Ratemaking

- Rates are to be not inadequate, not excessive, and not **<u>unfairly discriminatory</u>**
  - State rating laws
  - Statement of Principles Regarding Property and Casualty Insurance Ratemaking (recently rescinded)
- Principle 4 of the Statement of Principles Regarding Property and Casualty Insurance Ratemaking
  - A rate is reasonable and not excessive, inadequate, or unfairly discriminatory if it is an **<u>actuarially sound estimate of the expected value of all future costs associated with an individual risk transfer</u>**.

Unfairly Discriminatory?

- **Discrimination**: the process of evaluating insurable risks and determining premiums on the basis of likelihood of loss
- **Proxy discrimination**: occurs when a facially neutral trait acts as a stand-in (or a proxy) for a prohibited trait
- **Unfair discrimination**: formulation of rates on the basis of criteria that do not fairly measure the actual risk involved
- **Disparate treatment**: treatment of an individual that is less favorable than treatment of others for discriminatory reasons
- **Disparate impact**: an unnecessary discriminatory effect on a protected class caused by a practice or policy that appears to be nondiscriminatory

What Does This Mean For Insurance Pricing?

1. Do nothing
2. Exclude risk characteristics from rating plans
3. Control for protected characteristics in pricing
4. Adjust final pricing outcomes for protected classes

Exclude Risk Characteristics from Rating Plan

- Prohibit certain risk characteristics determined to be problematic from being used in developing premiums
- Considerations
  - Approach taken related to protected classes
  - Does it really address proxy discrimination?
  - Applied in some states (gender, credit score)
  - What is the ultimate criteria for exclusion?
  - Does this result in achievement of the ultimate goal?

What Does This Mean For Insurance Pricing?

- **Control for Protected Risk Characteristics**
  - Include protected characteristics in pricing models to control for proxy discrimination effect of other rating variables
  - Considerations
    - Systematically controls for proxy discrimination
    - Protected class data
    - Only directly applicable to certain model types
- **Evaluate Final Pricing Outcomes for Protected Classes**
  - Statistical measures can be calculated to determine if a rating plan has a disparate impact on protected classes. If it is determined, adjustments can be made to the rating plan to mitigate the impact
  - Considerations
    - Allows insurers to continue using complexity in risk segmentation
    - Shifts the focus from regulation of inputs and process to regulation of output
    - Requires collection of protected class data

Bias Mitigation Techniques

Bias in Data: Pre-process	Bias in Models: In-process	Bias in Predictions: Post-process
• Reweighting • Disparate Impact Remover • Optimized Preprocessing • Learning Fair Representations	• Adversarial De-biasing • Prejudice Remover • Meta Fair Classifier	• Reject Option Classification • Equalized Odds • Calibrated Equalized Odds

Poll Question 6

What Do You Think This Means for Insurance Pricing?

1. Do nothing
2. Exclude risk characteristics from rating plans
3. Control for protected characteristics in pricing
4. Adjust final pricing outcomes for protected classes