



Comments of the Center or Economic Justice

To NAIC Accelerated Underwriting Working Group

December 3, 2021

The Center for Economic Justice offers the following comments on the November 8, 2021 exposure draft of the working group's paper.

The paper continues to miss the key distinction between traditional underwriting and so-called accelerated underwriting – namely, the use of non-traditional, non-medical data. Insurers can accelerate the underwriting process in a number of ways that don't use non-traditional, non-medical data. Suppose that insurers were able to obtain traditional medical information in a faster, easier manner. Instead of asking consumers to provide a history of their prescription medicines or instead of obtaining and reviewing medical records from many providers, suppose an insurer could obtain that information electronically from a single source, like a prescription database. Although the insurer is still using the same traditional medical data, the insurer has accelerated the underwriting process.

If all insurers were doing was speeding up traditional underwriting methods, this group would not have been created. Just as property/casualty insurers speeded up auto and home underwriting by using all-claims databases and motor vehicle record databases instead of relying upon the consumer to provide that information, life insurers acceleration of access to and analysis of traditional medical information did not raise concerns among regulators.

It is not the use of predictive models or machine learning that distinguishes traditional underwriting from AUW – insurers have been applying such techniques to traditional underwriting data for years by more intensely analyzing traditional medical and other traditional data sources. The factor that most distinguishes AUW from traditional life insurance underwriting is the acquisition and use of non-traditional, non-medical data. This is evidenced by the fact that AUW models don't predict mortality – they can't because there is insufficient mortality data to develop a predictive model based on only a few years of data relating non-traditional, non-medical data to mortality. Rather, as the actuaries have stated, AUW is used to predict the same outcomes that would have occurred with traditional underwriting.

The problem with the current proposed definition of AUW is that, by lacking focus on the key differentiator of AUW from traditional underwriting, it obscures the new regulatory oversight steps needed to protect consumers from unfair discrimination and racial bias.

As we have urged in the past, the relevant definition for purposes of examining the adequacy of regulatory oversight and educating regulators and the public about AUW is:

Accelerated underwriting is life insurers' application of big data, artificial intelligence and machine learning to life insurance underwriting. What distinguishes AUW from traditional life insurance underwriting is the use of non-traditional, non-medical data using predictive models and machine learning.

The above definition focuses on the key differentiator between traditional underwriting and AUW and better sets the path for examining whether current regulatory structures require updating to protect consumers.

The error in the definition is reflected in the incorrect description of the differences between traditional and accelerated underwriting in the third paragraph on page 2.

Traditional life insurance underwriting involves assessing the applicant's physical health, then determining whether an applicant is eligible for coverage and the risk class to which that individual belongs. Accelerated underwriting relies on predictive models or machine learning algorithms to perform some of the tasks of an underwriter.

Traditional underwriting has always examined more than an applicant's health, including an applicant's financial situation (are they in bankruptcy?) and activities (are they a sky-diver?) as well as proxies for physical health (family history). Speeding up or more intensively analyzing these traditional data sources is not the reason why AUW is an issue of regulatory and consumer concern. It is the use of non-traditional, non-medical sources of data used with predictive models and machine learning that distinguishes AUW from traditional underwriting.

The fourth paragraph on page 2 continues to blur the needed understanding of AUW. The paper states that insurers use AUW to triage applicants and place applicants in different risk categories. Traditional underwriting has always done the same things. What distinguishes AUW from traditional underwriting is how the insurer does these two things – and for AUW that is the use of non-traditional, non-medical data sources in data-mined algorithms to accomplish these things.

At the bottom of page 2, the paper states that “increasing automation” of life insurance underwriting presents new regulatory challenges. Again, increased automation by itself is not the issue of concern. Automation can simply speed up manual processes using the same rules used by the formerly-manual process. And if all insurers were doing was speeding up traditional

underwriting by automating the acquisition and categorization of traditional medical information, there wouldn't be an AUW working group. Recall the origins of the NAIC's efforts on AUW at the Life Actuarial Task Force over five years ago – it was LATF's concern about the use of non-traditional, non-medical data used to predict traditional underwriting results as opposed to the use of traditional mortality tables.

The paper states that “the process” – referring to increasing automation – must be fair, transparent and secure, but offers no reference to the source or definition of these terms. As the text moves on to page three, the paper states a particular challenge is unfair discrimination, but offers no reference or definition of what is meant by unfair discrimination. As we have noted in several presentations, there are currently two types of unfair discrimination in insurance – discrimination not supported by actuarial analysis and discrimination based on protected class characteristics. The paper should explain why AUW raises new concerns about unfair discrimination.

We suggest that the paper discuss the history of life insurers' use of racial proxies for long periods of time as an example of the protected class unfair discrimination and life insurers' use of travel history as an example of unfair discrimination without actuarial basis (using Florida's actions to restrict such unfair discrimination).

We suggest that the first paragraphs on page 4 more clearly discuss the type of unfair discrimination at issue and how particular AUW data sources and applications raise concern for each of the two types of unfair discrimination. For example, this section of the paper discusses a concern about spurious correlations – where there may be a correlation between a particular data source and the insurer's outcome variable, but that correlation does not support the use of that data source as a reasonable or reliable predictor of that outcome variable.

Further, this section of the paper discusses testing, but is vague on the types of testing needed. The development of predictive models generally relies upon testing. The historical data is broken into two groups – one for development of the model and one for testing the model. The first group of data is data mined to develop the data elements and model specifications – the predictive model describes the relationship between the historical data and the outcome variable. The model is now run with the set-aside data to test whether model is reliable. We don't think this is the type of testing the paper is envisioning. The discussion of testing should reference testing for actuarial soundness on one hand and testing for protected class discrimination on the other hand.

In the recommendation section, the paper states AUW should be fair and transparent, but doesn't state to whom AUW should be transparent. AUW should be transparent to regulators, consumers and policymakers. The paper then states insurers should be accountable for operating in compliance with applicable laws. It is unclear why the word “should” is used or why this

statement is included. Insurers must operate in compliance with applicable laws. The question is whether existing laws are adequate to ensure that AUW is fair (which we expect to mean not unfairly discriminatory under current statutory standards) and transparent. Clearly, current regulatory requirements for life insurers need to be modernized for these consumer protections to be achieved. Yet, the paper makes no recommendation for or recognition of the needed improvements.

In the recommendations section on page 4, the paper sets out a list of “shoulds” for insurers and other parties involved in AUW. There is nothing in this list that distinguishes AUW from traditional underwriting – all of these “shoulds” apply to traditional underwriting data and methods.

This section fails to recommend specific testing for racial bias or algorithmic auditing to identify spurious correlations. This section fails to identify the needed requirement that insurers disclose to consumers the types and sources of data used – to actually implement the goal of transparency. The section fails to recommend requiring insurers to file AUW models with regulators – particularly credit-based AUW models. The section fails to recommend that life insurers using consumer credit information be held to the same standards as auto and home insurers. The section fails to discuss how the use of consumer credit, criminal history or consumer lifetime value information raises concern about algorithms reflecting and perpetuating historical racial bias. This section fails to recommend the development of regulatory guidance for what is needed to effectively implement the list of “shoulds” in a manner that complies with statutory standards.

In the section on traditional data, we suggest some authorities or sources be cited for the various assertions. For example, what is the source of the statement that “consumers understand how the elements impact their risk classification or premium charged?” What source does the paper rely upon to assert that consumers understand how MVRs and financial and tax information impact their premium charge or risk classification?

This section also claims that presentations to the working group represented significant time and costs associated with obtaining and reviewing traditional data. We suggest this is too generic a statement. While some sources of information remain costly and time-consuming to obtain – fluid and medical examinations – other sources of information have become readily and inexpensively available in digital formats – medical records, MVRs, MIB data, public records, prescriptions.

More importantly, the paper distinguishes traditional data from other types of data used in AUW, while largely dismissing any concerns with insurers' use of traditional data. We suggest that the paper's treatment of traditional data demonstrates and validates our points above about what distinguishes traditional underwriting from AUW – the use of non-traditional, non-medical data in predictive models and machine learning.

The next section of the paper discusses FCRA-compliant data. It is unclear why FCRA-compliant data is distinguished from traditional and non-traditional data. FCRA-compliant data are found in both the traditional and non-traditional data buckets. For purposes of defining AUW, whether data is FCRA compliant or not is not a distinguishing feature. For purposes of regulatory guidance and consumer protection, the FCRA provides a baseline of regulatory requirements for users of data and consumer protections. But, since FCRA-compliant data are found in both the traditional and non-traditional data buckets, it is not a third category of data. The FCRA serves as a guide for some of the regulatory changes and new consumer protections needed for AUW.

Further, the FCRA section is significantly incomplete. It is more accurate to describe data as FCRA-compliant or subject to the FCRA. The FCRA defines a consumer reporting agency and a consumer report and sets out a number of requirements for both consumer reporting agencies who collect and disseminate consumer reports and for companies using consumer reports provided by a consumer reporting agency. The list of consumer protections is far greater than those listed, including consent by the consumer for the use of the data, a notice of any adverse action, the ability to request a consumer report, the ability to correct erroneous data in a consumer report and the ability to request a reconsideration of the adverse action with corrected data. The FCRA also provides for oversight of the practices of consumer reporting agencies.

We suggest the working group review some documents regarding the FCRA.¹ For example the Consumer Financial Protection Bureau publishes a list of consumer reporting companies.² Some sources of traditional life insurance underwriting information are subject to the FCRA including data from the MIB, prescription drug histories and personal insurance claims information.

¹ https://files.consumerfinance.gov/f/201504_cfpb_summary_your-rights-under-fcra.pdf

² https://files.consumerfinance.gov/f/documents/cfpb_consumer-reporting-companies-list_2021-06.pdf

The paper then discusses non-traditional data sources. In the list of data sources, we suggest a broad grouping of biometric information, including facial, voice and other analytics based on personal biometric information. We suggest reference to the Illinois Biometric Information Privacy Act would be useful both to help describe biometric information and to identify needed new consumer protections. We also suggest more detailed descriptions of the sources of data and the uses and algorithms associated with those data. For example, how is specific biometric information used and for what purposes (e.g., to determine truth telling, biological age, body mass index)?

In the considerations section for non-traditional data, we see again a statement that such data may be used to predict mortality. We suggest a clear distinction between predicting the outcomes of the traditional underwriting process versus predicting mortality. It is unclear if insurers have sufficient historical data to associate non-traditional data sources with actual mortality.

The first consideration states that while non-traditional data may be used to predict mortality, there may not be a reasonable explanation for that correlation. This statement is problematic because it seems to assume that correlation is the same as actuarial soundness – it isn't. Further, it is unclear what a reasonable explanation means and how an insurer or regulator would interpret that term.

While we agree with the general thrust of the second bullet about racial bias, we suggest that the impacts of structural racism affect both traditional and non-traditional data sources. Further, we suggest the use of the term proxy discrimination as well as disparate impact. Proxy discrimination is the term used in the NAIC's principles for AI and is distinguishable from disparate impact. CEJ has presented the following definitions to the NAIC on several occasions:

Disparate Impact: Use of a non-prohibited factor that causes disproportionate outcomes on the basis of prohibited class membership and that such disproportionate outcomes cannot be eliminated or reduced without compromising the risk-based framework of insurance.

Proxy Discrimination: Use of a non-prohibited factor that, due in whole or in part to a significant correlation with a prohibited class characteristic, causes unnecessary, disproportionate outcomes on the basis of prohibited class membership.

Or

Proxy Discrimination: Use of an external consumer data and information source, algorithm, or predictive model whose predictive capability is derived in substantial part from its correlation with membership in one or more of such protected classes.

We attach a recent presentation to the NAIC to help the working group better understand these issues.

On the third bullet, we discussed above that some – even many – sources of non-traditional data used in AUW are FCRA-compliant data. This is not only important to better clarify the data categories used in the paper, but to demonstrate that many vendors of non-traditional AUW data sources and algorithms are the same consumer reporting agencies who provide data and algorithms auto and home insurers **and who currently file these algorithms with regulators**. A key consideration – and related recommendation – should be that life insurers (or third-party providers of AUW algorithms) file their models with regulators under the same types of regulatory requirements that exist for insurers filing credit-scoring models or catastrophe models for auto and home insurance. There is no rationale for treating auto and home insurers' use of credit and other non-traditional information differently from life insurers' use of the same data.

The recommendation in this section are significantly inadequate. The recommendations suggest that market conduct examinations are sufficient to ensure that AUW algorithms meet all the stated regulatory goals. We strongly disagree. First, market conduct examinations are infrequent and are typically triggered by some identified problem. Consequently, market conduct examinations cannot meaningfully address the activities of many insurers in a timely fashion. Nor are there existing metrics or data sources available to market analysts to trigger the types of concerns raised in the paper regarding racial bias or problems with data or algorithms. Second, there are no standards for market conduct examiners for most of the issues / considerations raised by the paper. Third, market conduct examinations are after-the-fact and not timely. Significant consumer harm – some irreparable -- will have occurred in the time it takes to start and complete a market conduct exam. Fourth, market conduct examinations are not the appropriate tool to establish the new guidance needed for insurers' use of big data and AI. You can't simply give market conduct examiners the NAIC principles for AI and expect enforcement or compliance or expect all insurers to discern regulatory guidance from the market conduct examination outcomes for one insurer.

The recommendations regarding form and rate are particularly puzzling since there is no rate regulation of life insurance and no current filing of life insurance rates. The only routine filing by life insurers is policy forms and applications. While a review of an application may indicate the use of a particular non-traditional data source, it's more likely that non-traditional data sources are not revealed in the application – so there would be nothing to trigger a form reviewer's question.

The first, third and fourth bullets under “form and rate reviewers” all assume some form of filing for a reviewer to analyze. There is a need for up-front filing and review of data sources and pricing models used by life insurers. Consumer protection demands that regulators stop the use of biased and unreliable data sources prior to use by insurers in the same way that regulators now stop the use of unfair, deceptive and prohibited policy form provisions prior to use by insurers.

Finally, we don't understand the last two paragraphs of this section telling life insurers they "should" engage in certain practices. How does the working group expect life insurers to respond to these "should" statements and what should regulators or consumers do if the insurers don't follow these exhortations?

The paper largely repeats the guidance for insurers set out in the NAIC principles for artificial intelligence. The purpose of the NAIC AI principles was to serve as the foundation for working groups to develop the application-specific regulatory guidance needed to operationalize those principles. ***We see no purpose served by only repeating those principles in a paper discussing a specific application of AI. The paper should be making recommendations for specific regulatory actions – new uses of existing regulatory authorities and tools and new regulatory authorities – needed to ensure that the AI principles are implemented for AUW.*** But, the paper offers no recommendations for how regulators and insurers should implement the AI Principles for life insurers' use of AUW.

Toward that end, the paper should be recommending specific statutory and regulatory changes, including:

1. Require life insurers to routinely file a list of the types, sources and uses of non-medical data for life insurance marketing, underwriting, claim settlement and anti-fraud. Regulatory attention to data and sources used for marketing is particularly important in the context of AUW because new data sources permit the micro-targeting of highly granular marketing to consumers, which effectively serves as pre-underwriting of life insurance. Regulators should pro-actively identify the types, sources and uses of data used by life insurers to timely stop the use of data that is biased, unfair or counter to public policy – instead of only learning about such data and data uses in a market conduct exam or through a media report. Further, regulators should not only collect this information, but publish summary reports to inform the public and policymakers about life insurers' data use.
2. Require life insurers to routinely file and regulators to routinely review algorithms used for marketing, underwriting, claims settlement and anti-fraud in the same manner that auto and home insurers are required to file credit-based insurance scoring models.
3. Require that all data sources used by insurers meet the consumer protection requirements of the FCRA, including consent, disclosure, challenge and correction.
4. Develop specific guidance and requirements for insurer testing of data sources and algorithms for actuarial soundness and protected class bias. It makes no sense to suggest that racial bias is a concern with AUW or other life insurer algorithms and then do nothing to prompt insurers to test for such bias and provide guidance for what sort of testing is reasonable and necessary. Why doesn't the paper recommend that all states – and the IIPRC – take the approach used by the New York Department of Financial Services in the cited Circular 1?

5. Recommend the development of guidance for life insurer collection and treatment of applicant data on race, ethnicity and other demographic characteristics to assist insurers and regulators in assessing proxy discrimination and disparate impact based on protected class characteristics. Again, if the potential for racial bias with AUW is a concern, then the relevant data must be collected to test and measure for such bias. The work of the health work stream of the Committee on race is relevant and instructional on this issue.
6. Develop / update guidance for third parties providing pricing algorithms to insurers. A third party vendors that collects information from insurers, combines that information with other data sources and then provides insurers with an algorithm for underwriting or pricing or claims settlement is engaged in collective decision-making with the insurers. Absent oversight of vendors providing these collective-pricing or collective-claims settlement algorithms, the third party algorithm provider may be engaging in prohibited antitrust and anti-competitive activities. This is vividly illustrated by comparing the regulatory oversight over mortality tables – review and approval by regulators of the raw material used by insurers for pricing and reserving life insurance – with the lack of oversight of non-traditional data sources and algorithms that are used for the same purpose.
7. Request that the Market Regulation D Committee direct the Market Conduct Annual Statement (MCAS) Blanks Working Group to complete its work on the AUW revisions to the Life Insurance MCAS line independently of the work of the AUW WG. The MCAS Blanks WG efforts on adding AUW reporting to the Life MCAS was stopped earlier this year to wait for a definition of AUW adopted by this working group for its educational paper. That directive was justified by an argument for coordination and consistency of terms among working groups. While “coordination and consistency” are generally reasonable considerations, this rationale was never logical or applicable in this context.

The MCAS effort is directed at data collection for specific market analysis purposes, so a precise definition is necessary to ensure the right data goes into the right data buckets. The AUW WG effort is directed at a different audience for a different purpose and, to date, has produced a vague and imprecise definition of AUW. It could never be used to generate reliable MCAS data reporting. We suggest that the AUW WG would benefit from review of the last version of the definition of AUW considered by the MCAS Blanks AUW subject matter expert group in which regulators and consumer stakeholders found agreement. AUW WG members will see a sharp focus on non-medical data obtained from other than the applicant (which would help inform this AUW WG’s definition) and the sharp difference in purposes of MCAS reporting and the charge of the AUW WG.

Text to accompany CEJ Presentation to Committee on Race, 12/1/2021

Slide 4

For those of you who don't know me, I'm Birny Birnbaum from the Center for Economic Justice.

The first few slides in the deck, which are available for download, provide background on me, my training as an economist at MIT, my service as an insurance regulator, my 30 years of work on racial justice in insurance. I'm speaking for both the Center for Economic Justice and the Consumer Federation of America and the nearly 300 state and national consumer organizations that are members of CFA.

Jump to Slide 5

To lay some groundwork, let's start by reviewing what fair and unfair discrimination in insurance means.

Unfair discrimination is generally defined in two ways. The first is actuarial – there must be an actuarial basis for different treatment of different groups of consumers. That is the “not unfairly discriminatory” portion of the statutory rate standards – not excessive, not inadequate and not unfairly discriminatory.

The second type of unfair discrimination is protected class discrimination – statutes that prohibit distinctions among groups defined by certain characteristics – race, religion, national origin. This type of discrimination is prohibited regardless of actuarial basis.

My question to you to start things off. Why is race a prohibited factor for underwriting or pricing even if there is an actuarial basis for such discrimination?

We know, at least for some lines of insurance, that race is predictive of insured loss. Black Americans have a lower life expectancy than other Americans – why are life insurers prohibited from using race as an underwriting or pricing factor? And, if race were predictive of auto insurance claims, why shouldn't insurers be able to use that or any factor predictive of claims? One reason could be that a person has no control over their race – they're born with it. But, there are plenty of pricing factors based on characteristics that consumers have little or no control over – like age or gender for auto insurance. So, again, why do state and federal laws declare racial discrimination as unfair discrimination in insurance?

Move to Slide 7

Slide 7 shows a map of Cleveland – 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

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.

Next Slide, 8

Let me suggest the reason that race and protected class characteristics are carved out regardless of actuarial fairness is that 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 – whether intentional or unintentional – that are unrelated to risk or cost of insurance. The 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.

That's why US federal civil rights and anti-discrimination laws in employment, credit and housing have always been understood to prohibit not just intentional discrimination, but practices – intentional or unintentional – that result in disparate outcomes based on race. Federal laws – and every court that has opined on the issue – have recognized both disparate treatment and disparate effect as unfair discrimination – that is intentional discrimination as well as facially-neutral practices that have the same effect as intentional discrimination.

Move to Slide 10

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

Systemic racism refers to policies, practices, or directives that result in advantages or disadvantages to individuals or communities **based on race**, including harm caused by infrastructures that determine access and quality of resources and services.

Slide 11

Let me identify 3 ways in which systemic racism can manifest in any aspect of the insurance life-cycle:

Intentional discrimination on the basis of race – disparate intent.

For today's presentation, I want to focus on two types of, hopefully, unintentional forms of racial bias.

Proxy Discrimination -- Disproportionate Outcomes On the Basis of Race Resulting from Proxies for Race; and

Disparate Impact -- Disproportionate Outcomes on the Basis of Race Because of Historic Discrimination Embedded in Insurance Outcomes

Next Slide 12

Proxy Discrimination – this is when a predictive factor – say, a rating variable – is actually predicting race and not the intended outcome. The result is unnecessary racial bias because the predictive factor is not, in fact, predicting the outcome. For example, consider the use of criminal history information in, say, Ferguson Missouri. Using criminal history as a predictive variable would simply be a proxy for the racist policing.

The other category is disparate impact – this occurs when the insurance outcomes are racially-biased because the racial bias is embedded in the insurance outcomes. Recall the map of Cleveland from earlier, an accurate assessment of flood risk will have a racial bias because of racial bias in housing.

It is important to distinguish between proxy discrimination and disparate impact. With proxy discrimination, insurers have or should have interest in stopping this unnecessary discrimination.

Disparate impact, however, requires a policy decision based on equity considerations – specifically – does prohibiting the use of a particular data source or consumer characteristic compromise the cost-based and risk-based foundation of insurance? We know that such equity-based policy decisions have been made – that’s why intentional use of race is prohibited.

Next Slide 13

While there is an important distinction between disparate impact and proxy discrimination, **there is a common methodology to test for both and** such testing is consistent with the predictive analytic methods that insurers already use.

In the Big Data / AI era, it is essential for insurers to test their algorithms and for regulators to test actual consumer market outcomes for proxy discrimination and disparate impact.

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.

Next Slide 14

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.

We also hope that you reject as absurd the p/c trades argument that they can't discriminate on the basis of race because they don't consider race. Anyone who works with predictive modeling and algorithms knows that algorithms will reflect and perpetuate any bias in historical outcomes embedded in the historical data.

Move to Slide 16

It is Reasonable and Necessary to Recognize Proxy Discrimination and 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.

Next Slide 17

At the urging of the P/C Trades, NCOIL adopted a definition of proxy discrimination that profoundly misunderstands how structural racism affects insurance. NCOIL's defines proxy discrimination only as "the **intentional** substitution of a neutral factor for a factor based on race, color, creed, national origin, or sexual orientation **for the purpose of discriminating** against a consumer **to prevent that consumer from obtaining insurance or obtaining a preferred or more advantageous rate due to that consumer's race, color, creed, national origin, or sexual orientation.**

At best, this action represents a profound misunderstanding of how systemic racism affects insurance. At worst, it is a conscious act of stopping insurance regulators and states from even attempting to address racial justice. The language memorializes insurer practices that indirectly discriminate on the basis of race, discourages insurers from examining such racial impact and restricts current regulatory efforts. It is based on a profoundly-flawed legal argument and NCOIL's mistaken belief that actuarial soundness requires only a simple correlation.

If there is to be any progress towards racial justice in insurance, the NCOIL definition of proxy discrimination must be rejected.



**Presentation to Property / Casualty Work Stream of
NAIC Special Committee on Race**

Proxy Discrimination and Disparate Impact in Insurance

December 1, 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 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.

Fair and Unfair Discrimination in Insurance

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

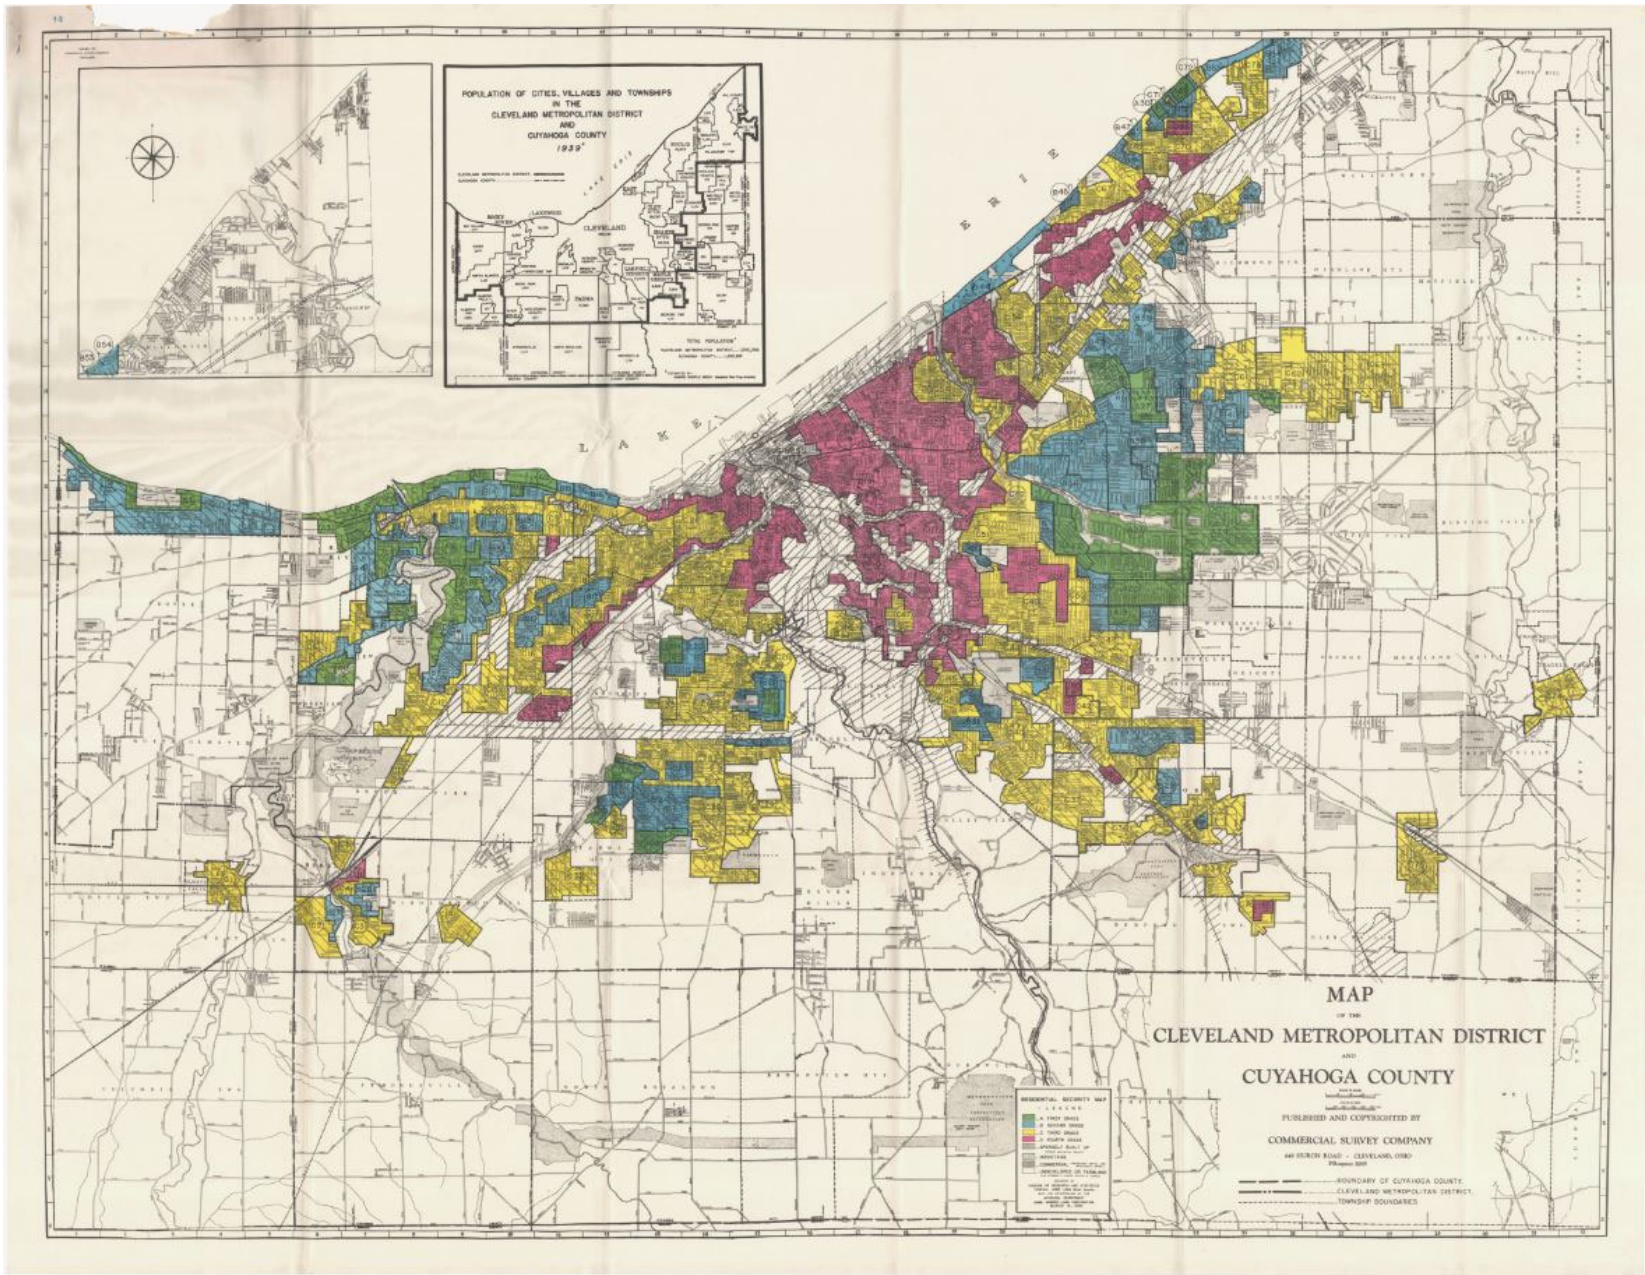
We find two types of unfair discrimination:

- 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.

Why do state and federal laws prohibit discrimination on the basis of certain characteristics even if there is an actuarial basis for such discrimination?

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



Why Do State and Federal Laws Prohibit 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.

Systemic Racism¹

Structural racism is the policies and practices that normalize and legalize racism in a way that creates differential access to goods, services, and opportunities based on race.

Systemic racism refers to policies, practices, or directives that result in advantages or disadvantages to individuals or communities based on race, including harm caused by infrastructures that determine access and quality of resources and services.

¹ <https://new.finalcall.com/2021/03/09/death-by-zip-code-housing-discrimination-neighborhood-contamination-and-black-life/>

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 – Disparate Impact
3. Disproportionate Outcomes Tied to Use of Proxies for Race, Not to Outcomes – Proxy Discrimination

Definitions

Disparate Impact: Use of a non-prohibited factor that causes disproportionate outcomes on the basis of prohibited class membership and that such disproportionate outcomes cannot be eliminated or reduced without compromising the risk-based framework of insurance.

Proxy Discrimination: Use of a non-prohibited factor that, due in whole or in part to a significant correlation with a prohibited class characteristic, causes unnecessary, disproportionate outcomes on the basis of prohibited class membership.

Or

Proxy Discrimination: Use of an external consumer data and information source, algorithm, or predictive model whose predictive capability is derived in substantial part from its correlation with membership in one or more of such protected classes.

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.

NCOIL's "Definition" of Proxy Discrimination Must Be Rejected

At the urging of the P/C Trades, NCOIL recently adopted the following:

For purposes of this Act, as well as for the purpose of any regulatory material adopted by this State, or incorporated by reference into the laws or regulations of this State, or regulatory guidance documents used by any official in or of this State, "Proxy Discrimination" means the **intentional** substitution of a neutral factor for a factor based on race, color, creed, national origin, or sexual orientation **for the purpose of discriminating** against a consumer **to prevent that consumer from obtaining insurance or obtaining a preferred or more advantageous rate due to that consumer's race, color, creed, national origin, or sexual orientation.**

At best, this action represents a profound misunderstanding of how systemic racism affects insurance. At worst, it is a conscious act of stopping insurance regulators and states from even attempting to address racial justice. The language memorializes insurer practices that indirectly discriminate on the basis of race, discourages insurers from examining such racial impact and restricts current regulatory efforts.

Algorithms Learn the Bias Reflected 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.**²*

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.

² Barocas and Selbst

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

Among the various parts of the insurance life-cycle – marketing, underwriting, pricing, claims settlement, antifraud – new data sources and complex algorithms for pricing currently get the most attention from regulators because in most states most insurers file personal lines rates. Data and algorithms used for marketing, in contrast, get little or no attention. ***Yet, it is the marketing function – and the new data sources and algorithms used in micro-targeting consumers – that has become the true gatekeeper for access to insurance.***

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"

And just recently, the telematics subsidiary of this insurer pitched its ability to identify the most valuable customers in real time:

Attract the most profitable drivers with telematics-based targeting

Traditionally, insurance marketing has relied on demographic and behavioral data to target potential customers. While useful at a high level, these proxies fall short when it comes to considering customer value and retention. Now, you can reach the most profitable customers from the outset using the nation's first telematics-based marketing platform.

Company intelligently layers driving score onto insurer campaign targeting criteria to purchase the ideal audience based on quartiles of driving risk. [The] Scored user receives a targeted offer via awareness and performance channels

Not to be outdone, another telematics data vendor announced a partnership with an auto manufacturer

Insurers can harness the power of connected Hyundai vehicles as a new marketing channel to support the profitable growth of their behavior- or mileage-based programs. Discount Alert allows insurers to deploy personalized marketing offers directly to drivers through Hyundai's online owner portal and contains robust tools to anonymously segment ideal risk targets—ensuring your offers are only sent to qualified leads.

All of this begs the questions, what about consumers and businesses who don't have the wealth to provide the value sought by insurers? How do these strategies line up with public policies against discrimination on the basis of race and promoting widespread availability of insurance?

The Murder of George Floyd Raised Awareness of Systemic Racism How Did Insurer CEOs React?

“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.

How Have the U.S. Insurer Trades – Particularly NAMIC and APCIA – Responded to the Insurer CEOs’ Calls?

- Opposed the inclusion of “Consistent with the risk-based foundation of insurance, AI actors should proactively . . . avoid proxy discrimination against protected classes” in the NAIC Principles for Artificial Intelligence.
- Have opposed the application of disparate impact liability under the federal Fair Housing Act to home insurance.
- Supported the gutting of the U.S. Housing and Urban Development’s disparate impact rule – despite pleas from several insurers to leave the rule alone in the aftermath of the murder of Black Americans at the hands of police.
- Pushed NCOIL to adopt a resolution opposing the CASTF White Paper because it suggested that regulators could ask insurers to show a rational relationship between new data sources and insurance outcomes.

How Have the Insurer Trades – Particularly NAMIC and APCIA – Responded to the Insurer CEOs’ Calls? (con’t)

- Opposed state bills to limit the impacts of credit-based insurance scores during a pandemic, citing insurers’ need for “risk-based pricing,” while supporting efforts to permit such deviations when insurers find it convenient – price optimization, consumer lifetime value.
- Sued regulators in NV and WA who sought temporary limits on the use of credit-based insurance scores disrupted by the pandemic and the CARES Act.
- Pushed NCOIL to adopt a definition of proxy discrimination that would block any efforts to identify and address disparate impact and proxy discrimination and shield insurers from any accountability for their practices.

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.

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.”³

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.

³ <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”⁴

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 Kahneman’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.

⁴ Yael Eisenstat at <https://www.wired.com/story/the-real-reason-tech-struggles-with-algorithmic-bias/>

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 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

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.⁵
3. Reach out to data brokers and vendors for a new data service.

⁵ 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>