American Academy of Actuaries
P/C Committee on Equity and Fairness (PCCEF)
Colorado Division of Insurance Stakeholder Meeting on Private Passenger Auto

August 24, 2023

About the Academy

AMERICAN ACADEMY
of ACTUARIES

- The American Academy of Actuaries is a 19,500-member professional association whose mission is to serve the public and the U.S. actuarial profession. For more than 50 years, the Academy has assisted public policymakers on all levels by providing leadership, objective expertise, and actuarial advice on risk and financial security issues.
- The Academy also sets qualification, practice, and professionalism standards for actuaries in the United States.

For more information, please visit:

www.actuary.org

Disclaimer

Please note: The American Academy of Actuaries' representatives' statements and opinions are their own and do not necessarily represent the official statements or opinions of the ABCD, ASB, any boards or committees of the American Academy of Actuaries, or any other actuarial organization, nor do they express the opinions of their employers.

Agenda

- Review the Academy's P/C Committee on Equity and Fairness goals and activities
- Share recent Academy work related to bias and discrimination
- Highlight key discussion points from our recent issue brief on approaches to identifying and mitigating bias in P/C insurance

P/C Committee on Equity and Fairness

Provide independent actuarial perspective on property and casualty equity and fairness topics while informing public policymakers on property and casualty practices that could potentially disadvantage people of color and/or other people in marginalized groups by:

- Providing key actuarial considerations regarding proposed laws and regulations;
- Offering regulators and legislators actuarial insights to consider as they evaluate proposals from industry groups, consumer groups, and other stakeholders;
- Considering and evaluating potential solutions to concerns associated with the use of external data, complex models, and outputs from actuarial models.

P/C Committee on Equity and Fairness

Engagement to-date:

- NAIC Special Committee on Race and Insurance
- NAIC Big Data and Artificial Intelligence Working Group
- NAIC Casualty Actuarial and Statistical Task Force
  - Recent presentation by Mike Woods at 8/12/23 CASTF meeting
- DC Department of Insurance, Securities and Banking (DISB) data call on unintentional bias in auto insurance
- NCOIL Special Committee on Race in Insurance Underwriting
- Federal Insurance Office
- Colorado Senate, Division of Insurance re SB 21-169

P/C Committee on Equity and Fairness

- Colorado SB 21-169 activities
    - March 7, 2023, joint letter to Division of Insurance on draft regulation for life insurance
    - February 4, 2022, joint letter to Division of Insurance following enactment
    - March 29, 2021, letter to Sen. Buckner

Recent Academy Issue Briefs

- *An Actuarial View of Data Bias: Definitions, Impacts, and Considerations*, July 2023
- *Approaches to Identify and/or Mitigate Bias in Property and Casualty Insurance*, February 2023
- *An Actuarial View of Correlation and Causation – From Interpretation to Practice to Implications*, July 2022
- *Sourcing Protected Class Information in P&C Insurance*, June 2022

Approaches to Identify, Mitigate Bias in Property and Casualty Insurance Issue Brief

- Discusses principles to be considered that may assist regulators in selecting suitable methodologies for identifying and/or mitigating bias.
- Structure of issue brief
  - Actuarial standards of practice
  - Definitions of unfair discrimination and disproportionate outcomes
  - Principles for approaches to identify and address unfair discrimination
  - Data collection, classification, and other considerations
  - Methods of identifying potential bias
  - Methods of preventing and addressing potential bias

Principles When Identifying and Addressing Unfair Discrimination

- Understandable to public
- Rates that continue to differentiate based on expected cost
- Adaptable to new data, innovation, and technology
- Consider intersectionality of protected classes
- Consistent application to all insurers
- Consider multivariate effects
- Assess impact to insurance marketplace
- Monitor results after initial approval
- Continually refresh data on protected classes

Potential Approaches for Identifying Bias

Disproportionate Impact Analysis. Study the impact that each rating variable has on each protected class's premiums. *To what extent does each rating attribute cause higher premiums for each class of insureds?*

Fairness Metrics. Compare model predictions to actual outcomes. *Is there bias (by protected class) in the prediction error in the loss model that supports the rating plan?*

Insurance Data Disclosure. Require insurers to release data on protected classes (such as loss ratios, frequency/severity, bind rates, rejection rates, etc.). *Allow the public to see whether there is bias in an insurer's practices.*

Potential Approaches for Identifying Bias

Loss Ratio Test. Compare loss ratios by variable of interest to demonstrate if they are materially different by protected class.

Proxy Test. Include protected class data in the rating model and see if the variable of concern continues to have predictive power.

Rational Explanation. Require carriers to describe a potentially causal relationship between the variable of concern and losses.

Identifying Bias: Sample Data

Protected Class Categories

	1	2	3	Total
(1) Premium	$ 600	$ 2,000	$ 3,000	$ 5,600
(2) Exposures	10	25	30	65
(3) Loss	$ 385	$ 1,517	$ 2,000	$ 3,902
(4) = (3)/(2) Loss Cost	$ 39	$ 61	$ 67	$ 60

Protected Class Categories

	1	2	3	Total
(5) = (1)/(2) Average Premium	$ 60	$ 80	$ 100	$ 86
(6) = (3)/(1) Loss Ratio	64%	76%	67%	70%

Assumptions:

- Coverages are the same across protected class categories (deductibles, limits, etc.)
- The examples assume there are no significant distributional differences between protected classes
- Expenses are proportional to loss and do not vary by protected class
- The protected class in the illustration is not being used in rating (e.g., not age or gender)

When premiums are different

Three customers purchase auto insurance from a company with the same coverage, limits, and deductibles. One pays $60, one pays $80, and one pays $100.

Fair or not fair?

$60 **$80** **$100**

Lower expected losses vs. premiums

The same three customers have three different expected loss ratios: One has an expected loss ratio of 64%, one has an expected loss ratio of 76%, and one has an expected loss ratio of 67%.

Fair or not fair?

64% **76%** **67%**

Identifying Bias in Outcomes: Disproportionate Impact Analysis (Average Premium Test)

Protected Class Categories

	1	2	3	Total
Average Premium	$ 60	$ 80	$ 100	$ 86

Recall we assume that there are no significant distributional differences (such as vehicle age and symbol) between protected classes.

- Comparing the average premiums, one sees that the average premiums are not the same for all class categories.
- Class 1 has a lower average premium than classes 2 and 3.
- That is, classes 2 and 3 are paying more premium per vehicle than class 1.

Identifying Bias in Outcomes: Loss Ratio Test

	Protected Class Categories			
	1	2	3	Total
Loss Ratio	64%	76%	67%	70%

- Comparing the loss ratios, one sees that the loss ratios are not the same for all class categories.
- Class 1, for example, has a lower loss ratio than classes 2 and 3.
- That is, class 1 is paying more premium per dollar of loss than the other two classes.

Your perspective depends on the measuring tool

In the first scenario, the first customer who has a $60 premium seems to have an **advantage** over the other two customers.

In the second scenario, the first customer with an expected loss ratio of 64% seems to have a **disadvantage** compared to the other two customers.

Both cases arise from the same set of data, *but* the evaluation of fairness depends on the **metric of interest** and the metric of interest varies by stakeholder.

Identification Methods for Models: Fairness Metrics

- Fairness metrics evaluate the bias in a model by comparing a model's predictions to actual outcomes
- One example of a fairness metric is independence. Independence evaluates whether the model predicts the same outcome for each class. Independence is similar to the average premium method described above.
- Another example of a fairness metric is accuracy parity. Accuracy parity evaluates whether the model error for each protected group is the same. Accuracy parity is similar to the loss ratio method described above.

Identification Methods for Models: Proxy Test

- A variable is a statistical proxy if it is not directly relevant but instead derives its predictive power from its correlation to another factor (such as protected class)
- One way to test whether a variable is a statistical proxy for a protected class is to include protected class data in a model and check whether the variable continues to have predictive power while including protected class in the model

Model Fairness vs. Market Rates

- Statistical methods include the proxy method and fairness metrics, which assume the use of a statistical model to generate relativities for rating variables
- These methods focus on the model outputs, but the entire modeling process, including data collection, should be examined as bias can enter at many points.
- Companies often adjust the rating relativities generated by the statistical model for many reasons, including competitive reasons or to minimize policyholder disruption
- Note that these adjustments to an otherwise bias-free model may result in bias

Pricing vs. Underwriting

There is overlap between pricing and underwriting

- Use of multiple companies within a group (e.g., standard vs. non-standard)
- Use of underwriting tiers within a single company (e.g., using credit-based insurance scores and prior liability limits to tier consumers)

Carriers may use underwriting in response to rating variable restrictions

- Companies can impose underwriting guidelines to limit or deny coverage if rating variable restrictions lead to unprofitable groups of risks

Potential Framework for Testing for Bias

First, evaluate coverage selection and distributional differences across allowed rating factors between protected classes.

- This is important, as it could be a big driver of differences.
- Different "solutions" to address lack of fairness could come from understanding these distributional differences.

Control for differences and review multiple metrics to develop a fuller evaluation of fairness between protected classes (i.e., average charged premiums, loss ratios, expense ratios).

Consider fairness throughout the multiple layers of the pricing approach (model outputs, selected factors and rates, underwriting overlays like tier or company).

- Statistical approaches to evaluating fairness can be applied to understand if a lack of fairness is coming from the modeling process, selection process, and/or underwriting process.

Potential Approaches for Preventing/Addressing Bias

Allow Only Pre-Approved Variables. States would provide a list of variables that companies are allowed to use in policy rating.

Prohibit Named Variables. Each state would provide a list of variables that cannot be used in policy rating.

Limit Rate Spread. Limit the spread of rating factors (e.g., no surcharge can exceed 30%) or limit the spread of premiums (e.g., the highest premium cannot be 3x greater than the lowest premium).

Potential Approaches for Preventing/Addressing Bias

Rate Factor Adjustment. Adjust rate factors (manually or algorithmically) until a test to identify bias has been passed.

Solidarity Tax and Rebate. Collect a tax from all policyholders and redistribute that tax as a rebate to those that have been identified as deserving a subsidy.

Statistical Model. Build an initial model using all rating variables and the protected class variables; then, algorithmically remove any proxy effects from the rating variables (and the protected class variables).

Conclusion

- Growing discussion around unintended bias and unfair discrimination
- There are many potential methods to identify and/or mitigate bias that have been discussed
  - There are likely to be even more methods in the future as discussions continue
- The American Academy of Actuaries is ready to assist regulators in their review of the technical components of these methods as well as in identifying strengths and weaknesses, particularly in relation to the principles noted in this presentation
- We hope these observations are helpful and we welcome further discussion

Questions/Comments

For more information, contact:
Rob Fischer, casualty policy analyst
fischer@actuary.org