American Academy of Actuaries

Issue Brief

Approaches to Identify and/or Mitigate Bias in Property and Casualty Insurance

FEBRUARY 2023

Key Points

- Concerns about potential unintended bias in insurance practices have escalated in recent years, especially for property and casualty lines, and there is increasing activity in legislation and regulation in this area.
- There are many approaches that might be used to identify and address bias.
- Several principles can be considered that might assist regulators in the selection of suitable methodologies.

In recent years, there has been considerable discussion regarding unintended bias and unfair discrimination in insurance rating. The American Academy of Actuaries (Academy) views this as a significant public policy issue. This is especially true for the property and casualty lines (P/C) and more specifically in the personal lines coverage market. As such, through its P/C Committee on Equity and Fairness (formerly Racial Equity Task Force, RETF), the Academy has actively participated in these dialogues. This has been accomplished through direct presentations, comment letters, and papers. This paper is intended to add to the discussion by providing a survey of methods aimed at helping to identify and/or mitigate unfair discrimination and unintended bias in rating for property and casualty lines.

Recognizing the importance of the role played by actuaries in designing and implementing risk classification plans, the paper starts with a brief review of key actuarial documents that support the committee's efforts in this regard. Following that, a few key definitions are discussed and then several principles and considerations that are important to regulators when determining how to address concerns about potential bias. The paper then presents a discussion of various methods for identifying potential bias and methods of preventing or addressing potential bias.

The appropriateness of any given method to identify and prevent any potential bias may be determined by the circumstances within which it is being considered. The use of risk classification plans with an emphasis on recognition of relative costs differences among the insured population is critical to a well-functioning insurance market. Avoiding unintended bias is also critical, and the committee supports efforts to eliminate unintended

American Academy of Actuaries
1850 M Street NW, Suite 300
Washington, DC 20036
202-223-8196 | www.actuary.org

Rich Gibson, MAAA, FCAS—Senior Pension Fellow

bias and unfair discrimination and recognizes the importance of finding a balance in the use of personal characteristics, external data, algorithms, and predictive models by the P/C insurance market. Please note that there may be other considerations that may be important for other practices or other purposes that are not addressed here.

Actuarial Standards and Guidance

Practicing actuaries in the U.S. are subject to professional guidance through the Code of Professional Conduct and the actuarial standards of practice (ASOPs) developed by the Actuarial Standards Board. Several ASOPs provide insight into actuarial risk classification. For example:

- ASOP No. 12, *Risk Classification*, provides guidance to actuaries when performing professional services with respect to designing, reviewing, or changing risk classification systems.
- ASOP No. 23, *Data Quality*, provides guidance to actuaries when performing actuarial services involving data.
- ASOP No. 56, *Modeling*, provides guidance to actuaries when performing actuarial services with respect to designing, developing, selecting, modifying, using, reviewing, or evaluating models.

Ideally, for the practicing actuary, laws and regulations and actuarial standards of practice are consistent. If a law or regulation conflicts with the guidance in an actuarial standard of practice, actuaries are required to comply with the requirements of the law and disclose any conflict.

Of particular relevance to P/C actuaries are the following sections from ASOP No. 12.

ASOP No. 12, section 3.2.1 states:

> A relationship between a risk characteristic and an expected outcome, such as cost, is demonstrated if it can be shown that the variation in actual or reasonably anticipated experience correlates to the risk characteristic.

Members of the Committee on Equity and Fairness, which authored this issue brief, include Lauren Cavanaugh, MAAA, FCAS—*Chairperson*; Dorothy Andrews, MAAA, ASA; Amy Angell, MAAA, FCAS; Rebecca Armon, MAAA, FCAS; Steven Armstrong, MAAA, FCAS; Rose Barrett, MAAA, FCAS; Mallika Bender, MAAA, FCAS; George Davis, MAAA, FCAS; Kwame Davis, MAAA, FCAS; Kevin Dyke, MAAA, FCAS; Richard Gibson, MAAA, FCAS; David Heppen, MAAA, FCAS; Susan Kent, MAAA, FCAS; Benyamin Kosofsky, MAAA, ACAS; Sandra Schrader, MAAA, FCAS; Benjamin Turner, MAAA, ACAS; and Michael Woods, MAAA, FCAS.

ASOP No. 12, 3.2.2 states:

> While the actuary should select risk characteristics that are related to expected outcomes, it is not necessary for the actuary to establish a cause and effect relationship between the risk characteristic and expected outcome in order to use a specific risk characteristic.

Definition of Unfair Discrimination and Disproportionate Outcomes

Rates are generally assumed to be actuarially sound if they account for the expected future costs of losses and expenses. Unfair discrimination in insurance has commonly been understood to be the lack of a relationship between price differences between groups of people and expected differences in losses and expenses. Laws and regulations in many states impact the information insurers can use in developing rates and prohibit insurers from using specific information, for example, race, directly in ratemaking models. Groups of people legally protected from discrimination are called protected classes. Insurers are also prohibited in many states from intentionally using proxy rating variables for protected classes. For example, a National Council of Insurance Legislators (NCOIL) model law defines proxy discrimination 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."[1]

As defined in a 2002 Academy P/C Risk Classification Subcommittee report to the National Association of Insurance Commissioners (NAIC),[2] disproportionate impact occurs when a rating tool results in higher or lower rates, on average, for a protected class, controlling for other distributional differences. Care should be taken to note that "disproportionate impact" is different from "*disparate* impact," which is a concept used in legal contexts. The Casualty Actuarial Society (CAS) research paper "Defining Discrimination in Insurance" explores these definitions, including the three-step process for determining disparate impact, in more detail.[3]

1 National Council of Insurance Legislators, "NCOIL Special Committee on Race in Insurance Underwriting Holds Virtual Interim Meeting: Adopted Definition of 'Proxy Discrimination,'" March 8, 2021.
2 American Academy of Actuaries, *Use of Credit History for Personal Lines of Insurance*, 2002.
3 Kudakwashe F. Chibanda, *Defining Discrimination in Insurance*, Casualty Actuarial Society, 2022.

Principles for Approaches to Identify and Address Unfair Discrimination

Insurance practices are becoming increasingly complex, which will pose challenges for those regulating those practices. When considering various approaches to identifying and addressing unfair discrimination, regulators might consider the following principles that are consistent with actuarial standards of practice, support a consistent approach among insurers while allowing for appropriate flexibility, and can lead to best practices.

1. **Readily Understandable to All Stakeholders**
   In order to promote consistent application across all lines of insurance, regulators might consider methods that may be understood consistently by insurers, regulators, and the public.

2. **Rates That Continue to Differentiate Based on Expected Cost**
   Regulators may continue to allow insurers to differentiate rates based on expected cost. Risk-based rates incentivize safe behaviors and loss mitigation, encourage competition among insurers, and have led to a large reduction in the number of consumers being forced to buy insurance through "assigned risk pools."

3. **Adaptable to New Data, Innovation and Technology**
   Insurers are continuously innovating, and methods for data collection and technology are evolving. Regulators might consider the practicality and efficiency of the regulations governing the process used by insurers and regulators to adapt to new data, innovation, and technology.

4. **Definitions and Intersectionality of Protected Classes**
   Individuals in protected classes could fall into more than one category of protected class. Thus, regulators might consider the interconnected nature among different protected classes and the impact of such intersectionality. In addition, consideration may be given to the fact that policies might cover multiple individuals under a single policy where association with a particular class is difficult (e.g., a homeowner's policy for a multi-racial household).

5. **Consistent Application for All Insurers**
   Given that regulations will place certain requirements on insurers, regulators might consider methods that can be applied consistently across insurers.

6. **Multivariate Effects**
   Regulators might consider that rating variables are used within the context of a complex risk classification system, and multivariate effects are important to identify and quantify.

7. **Impact to Insurance Marketplace**
   While consistent application of methods across insurers is advisable, as regulators are aware, this will result in costs and practical challenges for some insurers, and unintentional impacts on accessibility, availability or affordability for consumers should be avoided in order to preserve a healthy insurance marketplace.

8. **Monitoring After Initial Approval**
   After a particular practice is approved as not unfairly discriminatory, given changes in demographics and other inputs over time, it is important for regulators to consider whether and how often they will monitor the practice and/or require re-application by the carrier.

9. **Frequency of Refreshing Data on Protected Classes**
   Currently, insurers do not usually collect most protected class information and may not be allowed to collect this information. However, if that were to change, given that an individual's protected class identification can change over time (e.g., religion, gender identity, or disability status), consideration of how often an individual's class information should be collected is advisable. Any data used in a model should be appropriate for the intended purpose and sufficiently current.

Data Collection Considerations

To study the treatment of protected classes, data will be required to perform an analysis and arrive at an assessment. However, the insurance industry currently does not have accurate and readily available sources of most forms of protected class data. Given that insurance companies generally do not collect protected class data today, new approaches for sourcing this data may need to be developed. These data sources could be developed by (1) obtaining data directly from the insureds, (2) capturing existing data from third-party databases, or (3) imputing the data using statistical methods. If companies were required by regulators to obtain protected class data, regulators may need to review regulations and statutes to ensure this is permitted in their states.

Each of these approaches for sourcing protected class data has benefits and drawbacks. A combination of approaches may need to be used in the short and medium terms until accurate policyholder data can be securely sourced. A more thorough discussion can be found in the American Academy of Actuaries' issue brief *Sourcing Protected Class Information in P/C Insurance*.[4]

Classification Considerations

When considering the definition of protected classes and the collection of data needed to demonstrate that practices are not unfairly discriminatory, regulators may want to consider the following principles.

1. **Capable of Being Objectively Determined**
   In order to promote consistent approaches and applications, protected classes should be defined in a way that allows for the collection of data that is objectively determined. This poses a number of challenges, including:
   - How will they be categorized?
   - How to handle individuals who identify themselves in more than one category, e.g., race for those identifying with multiple races?

2. **Practical Limitations in Collecting Data**
   Consideration of practical limitations in data collection for a risk classification variable, including cost and efficiency, is advisable.

3. **Ability to Achieve Credible Results**
   Given that individuals within a given protected class are likely to exhibit a number of varying and different risk characteristics, and that insurers will have differing insured populations and mixes of business, it is important that the class definitions balance homogeneity with the ability to achieve credible results when demonstrating that a particular insurance practice is not unfairly discriminatory, possibly across intersections of protected classifications.

4. **Frequency of Reviewing Definitions, After Established**
   Given that categorization of classes can change over time, it is important that consideration be given as to how often regulators would review and update the class definitions.

Many of the above principles are considered within ASOP No. 12, such as objectivity, practicality, and credibility.

4 American Academy of Actuaries, *Sourcing Protected Class Information in P/C Insurance*, Racial Equity Task Force issue brief, July 2022.

Other Considerations

Regulators may also consider the following:

1. **Unintended Impacts to Consumers—**It is prudent to consider potential unintended marketplace impacts that could result from implemented regulations, including potential negative impacts regarding availability, accessibility, and affordability of coverage.

2. **Multiple Methods—**There are many methods that could be considered as appropriate means to demonstrate that insurance practices are not unfairly discriminatory. Multiple methods can provide additional insight to regulators, rather than relying on one method.

3. **Small Companies—**Given that smaller insurers could have additional challenges in complying with the regulations, due to credibility and practical limitations (among others), regulators may wish to adopt methods that consider these challenges.

4. **Data Protection and Cost of Implementation—**There will be costs to the insurers related to complying with regulations or laws, which could impact premiums, including:
    a. Gathering and protecting sensitive data
    b. Storing data
    c. Performing analysis to support non-discriminatory rates

5. **Data Granularity—**It is important to consider whether the data that is used to evaluate unfair discrimination must be specifically linked to each individual (e.g., by self-reporting) or whether the data can be imputed from other sources (e.g., by using demographic data at a geographic level).

6. **Field Test—**Regulators may want to consider implementing a field test prior to the final adoption of any proposed methods in order to more thoroughly understand the impact that the methods might have.

7. **Use of the Data, Algorithms, or Predictive Models—**It is important to understand how any data (internal or external to the insurer), algorithms, or predictive models are used. Work examining how to test for bias in algorithms and predictive models is relatively new and evolving. An Academy Data Science and Analytics Committee issue paper provides considerations in this regard.[5] These considerations could be

5 American Academy of Actuaries, *Big Data and Algorithms in Actuarial Modeling and Consumer Impacts*, November 2021.

helpful to identify algorithms that may be biased. Similarly, there are public tools available for testing an algorithm for bias. While the American Academy of Actuaries does not endorse any of these publicly available resources, it should be acknowledged that these public tools can be helpful to inform regulators.

Methods of Identifying Potential Bias

This section contains a discussion of different methods to identify potential bias. The methods are listed alphabetically. These tests are formulated under the assumption that data related to protected classes is available. Each method has its own advantages, disadvantages, and purpose. The "best" method will depend on the goals of an analysis and the questions that are trying to be answered. Also, the same method may be interpreted differently on the same data by two different users. Each method contains a definition for "pass" or "fail." Choosing a threshold that defines "pass" or "fail" is critical, and thresholds should be set with consideration that the volatility of insurance loss data may necessitate that a range of outcomes be acceptable. It may also be prudent to consider using multiple methods simultaneously because different methods can give different results and different methods attempt to answer different "questions." For example:

1. Disproportionate impact analyses ask, "How much does each rating attribute cause higher premiums for each class of insureds?"
2. Fairness metrics ask, "Is there a bias in the prediction error of the rating plan model?"
3. Insurance data disclosures do not ask any specific questions but allow the public to evaluate whether there is bias in insurers' data.
4. Loss ratio tests ask, "Are premiums appropriate in relation to cost for each protected class?"
5. Proxy tests ask, "Do any rating attributes derive their predictive power from their correlation to a protected class?"
6. Rational explanations tests ask, "Is there a relevant, understandable relationship between each rating attribute and insured losses?"

Below is a description of each test, including a discussion of each approach's advantages and disadvantages.

Disproportionate Impact Analysis—

Disproportionate impact occurs when a rating variable results in higher or lower rates, on average, for a protected class, controlling for other distributional differences. This method compares premiums, like the Average Premium Analysis method referred to in the Loss Ratio test section below, but controls for distributional differences among protected classes. Many rating variables likely have disproportionate impact because protected classes (and all other classes) could have different risk characteristics than the average policyholder. For example, in regards to automobile insurance, if any protected class has a younger average age than the average age in the portfolio, the use of age as a rating variable would have a disproportionate impact on that class. In a disproportionate impact analysis, the magnitude of disproportionate impact that would be deemed "acceptable" would need to be defined since some level of disproportionate impact is likely to occur.

One approach to carry out a disproportionate impact analysis is to use Nonparametric Matching.[6] This approach involves matching insureds that are in different protected classes but have similar risk characteristics except for an evaluation variable. A model is then built on the matched dataset with the evaluation variable included and another model is built without the evaluation variable. The average predictions from the two models are then compared for each protected class. If the average prediction from the two models is significantly different for the same protected class, then the evaluation variable is determined to have a disproportionate impact on that protected class.

Disproportionate impact analysis is adaptable to new data and technology, considers multivariate effects, and could be monitored as rating plans are updated. The limitations of this method include that it is not easily understandable to the public, does not consider intersectionality, and does not have a consistent application to all insurers due to its reliance on each individual insurer's data. Requiring the removal of variables creating disproportionate impact could create rates that do not differentiate on all expected costs.

Fairness Metrics—

Fairness metrics evaluate the bias in a model by comparing a model's predictions to actual outcomes. Details on this method can be found in the paper *Methods for Quantifying Discriminatory Effects on Protected Classes in Insurance* by Mosely and Wenman.[7] An example of a fairness metric is accuracy parity. Accuracy parity evaluates whether the

6 Daniel E. Ho, Kosuke Imai, Gary King, and Elizabeth A. Stuart, *Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference*, Political Analysis (2007) 15:199–236.

7 Roosevelt Mosley, FCAS, and Radost Wenman, FCAS, *Methods for Quantifying Discriminatory Effects on Protected Classes in Insurance*, Casualty Actuarial Society, 2022.

model error for each protected group is the same. Accuracy parity is similar to the loss ratio approach described below. In addition to accuracy parity, there are other fairness metrics, defined in the Mosely and Wenman paper, which could be applied to insurance such as Equalized Odds, Equal Opportunity, and Calibration. Multiple fairness metrics can be assessed simultaneously, but fairness metrics can conflict with each other, so it may be impossible for all chosen fairness metrics to be satisfied simultaneously. After choosing a fairness metric, or a set of fairness metrics, the next step is to determine how to mitigate any bias identified. Removing the bias involves adjusting the data, the model, or the model outcomes until the fairness metrics are satisfied. After the fairness metrics are satisfied, then the model is determined to be "fair."

An approach using fairness metrics has similar advantages and disadvantages as the loss ratio method, given the similarity between the two methods. However, the method is more statistically sophisticated and thus may be less understandable to stakeholders.

Insurance Data Disclosure—

This method involves requiring carriers to disclose information about how their algorithms impact members of a protected group.[8] An example of this method in practice is the federal Home Mortgage Disclosure Act, which requires certain financial institutions to provide mortgage data to the public. An example of data could be loss ratio and distributions of customers written. Thus, a loss ratio test as discussed below could be performed if insurance data disclosure was implemented, but different measures could also be analyzed.

This method allows for variation across insurers in how the same rating variables are used, grouped, and combined in a model. There is a clearer connection to the impact felt by consumers than other methods. Either the procurement of protected group data or an acceptably accurate prediction of belonging to a protected group would be required. A challenge would be the inclusion of variable interactions and interpretation of the data. Some companies may not be equipped with the tools or support to meaningfully analyze this data. In addition, the method may require setting a threshold for acceptable/unacceptable impacts on protected classes.

8 If a disclosure requirement were imposed, data privacy would need to be considered.

Loss Ratio Test —

A loss ratio test could ask insurers to demonstrate that loss ratios are not materially different by protected class. There are different loss ratio metrics that could be used including historical loss ratios or loss ratios for a prospective period. Additionally, loss ratios could include or exclude loss adjustment and other expenses. This type of test can be done in a crosstab format where one variable is a desired rating variable and the other variable is a protected class. The losses and the premiums of each policyholder would then be summed up in crosstab format.

As an illustration, suppose one wishes to test a credit-based insurance score versus a protected class. Losses and premiums could be summed in a crosstab format like this:

	Non-protected Class	Protected Class
No Credit	$LR_{\text{no, non}}$	$LR_{\text{no, protected}}$
Low Credit	$LR_{\text{low, non}}$	$LR_{\text{low, protected}}$
Medium Credit	$LR_{\text{medium, non}}$	$LR_{\text{medium, protected}}$
High Credit	$LR_{\text{high, non}}$	$LR_{\text{high, protected}}$

Using this approach, a risk classification system should show consistent loss ratios across all cells in the crosstab, demonstrating that the rating variable is effectively matching premiums to losses independent of class status. This type of test can be done on all filed rate factors, but for continuous rate factors the variables will need to be reasonably grouped.

This method would be understandable to stakeholders, including the public, would allow rates to reflect expected costs, would be easier for all companies to implement (if data is available by protected class), and could address most forms of disproportionately negative outcomes. Because insurers should be simultaneously adjusting all rating variables in their rate filing based on multivariate effects, this technique should effectively deal with multivariate effects. Furthermore, if one wishes to test intersectionality, the protected class defined in the test could be an intersectional group. However, if there are variables acting as proxies for protected class variables, the loss ratio method would not detect this. Additionally, there could be limited credibility depending on the amount of data available for each class/variable grouping.

A seemingly similar but philosophically different methodology to the loss ratio test is an average premium test. The average premium test compares the average premium by protected class. However, comparing average premiums ignores other relevant differences between classes. For example, one protected class may frequently choose higher limits of insurance, and this could cause a difference in average premium that is explainable. The loss ratio test considers these differences since the premiums in the denominator reflect relevant differences between classes.

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.

One should also consider whether there are interactions between multiple variables that may be acting as a proxy for a protected class. To test this, the predictive power of every variable in a rating plan, and not just the one variable in question, should be analyzed after including protected class data in the model.

One version of the proxy test was developed by the Federal Trade Commission (FTC) in their study of credit-based insurance scores. The FTC proxy test involves three questions (the third of which is similar to the discussion above):

1. *Does expected cost differ by protected class?*
   a. If yes, there exists a potential for a proxy effect.
2. *Does the rating variable predict expected cost within protected class groups?*
   a. If no, then the variable in question is likely a proxy.
3. *Does controlling for protected class in a predictive loss model impact the rating variable's effectiveness?*
   a. If yes, then the variable in question is likely a proxy.

This method is fairly understandable to all stakeholders, adaptable to new data and technologies, considers multivariate effects, and could be monitored as new rating variables are introduced. However, requiring the removal of proxy variables would not allow insurers to reflect all expected costs and therefore could have an impact on the insurance marketplace. This test would not have consistent application to all insurers since it relies on the data of each insurer, which could be different due to the volatility of insurance claims, and small insurers may not have enough data to produce credible test results.

Require a Rational Explanation Between Variables and Loss Experience—
This method requires carriers to describe a potentially causal relationship between a variable and losses. It is similar to guidance provided by the National Association of Insurance Commissioners to regulators in the *Regulatory Review of Predictive Models* white paper. The rational explanation should explain "why a rating variable is correlated to expected loss or expense, and why that correlation is consistent with the expected direction of the relationship."[9]

This method is easy to understand, and rates would continue to differentiate based on expected costs. This method is adaptable to new data, innovation, and technology. On the other hand, this method does not directly address the issue of whether there is unfair bias and may be applied inconsistently from state-to-state or regulator-to-regulator since a rational explanation is subjective. Also, this method may affect the insurance marketplace if many variables are disallowed.

Methods of Preventing and Addressing Potential Bias

This section contains a discussion of different methods to prevent or address bias if it has been found. The methods are listed alphabetically. Again, each method has its own advantages and disadvantages, and it may be best to consider using multiple methods to address potential sources of bias.

Allow Only Pre-Approved Variables—
Allowing only certain pre-approved variables in the premium calculation is a simple and easy to apply method of addressing unfair discrimination.

In addition to being easy to apply, this method is unambiguous in that the variables are either on the pre-approved list or not. However, this approach has disadvantages since it may not necessarily eliminate proxy discrimination or recognize variable interactions. In addition, the process to determine and approve acceptable variables could be difficult and time consuming. Finally, it may be difficult to keep up with evolving innovation and appropriately add other variables.

9 National Association of Insurance Commissioners, *Regulatory Review of Predictive Models*, Casualty Actuarial and Statistical (C) Task Force white paper, 2020.

Limit Rate Spread—

Limiting the spread of rating factors (e.g., no surcharge can exceed 30%) or limiting the spread of premiums (e.g., the highest possible premium cannot be greater than three times the lowest possible premium) are potential methods to address unfair discrimination.

This method is understandable to the public and has consistent application. However, this method limits insurers' ability to fully differentiate based on expected cost. It may prevent very high premium differentials by class, but also would prevent very low premium differentials as it leads to narrower premium ranges and could affect affordability in some cases. The insurance marketplace may be impacted if rate differentiation is limited and some unintentional availability issues could result.

Prohibit Named Variables—

This is a straightforward and easy to implement method based on disallowing the use of certain named variables. For example, California prohibits the use of credit-based insurance scores for underwriting or rating auto insurance policies or for setting rates for homeowners insurance. Several other states have similar restrictions on credit information.

While the ease of use and unambiguous application of this method is appealing, it does not necessarily recognize variable interactions. In addition, the prohibition of certain variables may result in rates not reflecting all expected costs and may not actually resolve disparate outcomes. The selection of variables to exclude could be subjective and based on public opinion rather than founded in actual data. Further, as companies introduce new rating variables, the list of disallowed variables could be constantly evolving and may not always adequately consider these new variables. Additionally, it could vary by jurisdiction and product lines.

Rate Factor Adjustment—

Rate factors can be manually or algorithmically adjusted until a test to identify bias has been passed. In this method, rating factors correlated with protected classes are the most likely to be adjusted to pass a discrimination test.

Manual rate factor adjustment would not be readily understandable to all stakeholders because it would not be clear how insurers would choose which rating factors to adjust. There may also not be consistent application among insurers. Rates may not be able to differentiate based on expected cost, and this could have an impact on the insurance marketplace and the availability of insurance.

Solidarity Tax and Rebate—

This method uses the determination of an appropriate premium adjustment for members of protected classes, a taxation of all policyholders, and a process for reimbursing those who qualify for the premium adjustment. This method has been recommended by Daniel Schreiber, the CEO of Lemonade. He refers to it as a Solidarity Tax and Rebate and suggests that carriers collect the tax and regulators determine which people would receive the rebates.[10]

This method is relatively easy for customers and carriers to understand. It would have little impact to the insurance marketplace outside of requiring carriers to assess a surcharge on each policy, which would be redistributed as reimbursements following some process. It allows for the resolution of social goals, while having rates, prior to the tax and rebate, that continue to reflect expected costs. It would place the burden on parties outside of insurance carriers to determine who should receive a reimbursement and how much that reimbursement should be. From the perspective of the parties determining the reimbursements, attention would need to be paid to the definitions and intersectionality of protected classes. One downside to this approach would be the difficulty of administration since reimbursements would need to be determined, processed, and tracked each year. Furthermore, certain individuals with low incomes who do not qualify for the rebates would experience higher premiums, thus causing potential affordability issues.

Statistical Model—

Another approach is to build a non-discriminatory model, as described in the article *Discrimination-Free Insurance Pricing* by Lindholm et al.[11] This method could be used proactively by insurers to eliminate the proxy effects of rating variables prior to filing their algorithms. Essentially, the insurer would first build a model including all rating variables and the protected class variables. Then, the effect of the discriminatory information would be removed in such a way that the protected class variables are removed as well as any proxy effects from the remaining variables.

Considering the principles for addressing unfair discrimination identified earlier, this method can measure different types of disproportionately negative outcomes, is adaptable to new data, innovation, and technology, and can handle multivariate effects. However, the method may not be as understandable to all stakeholders, particularly the public, and it may be more challenging for some insurers without sufficient data or sophisticated pricing models to execute.

10 Daniel Schreiber, "Lemonade CEO: Why Regulators Need to Engineer Equity in Insurance Prices," *Carrier Management*, April 6, 2021.
11 M. Lindholm, R. Richman, A. Tsanakas, and M.V. Wüthrich, "Discrimination-Free Insurance Pricing," *ASTIN Bulletin—the Journal of the International Actuarial Association*, January 2022, 52: 55–89.

Use Technology to Reduce Reliance on Existing Rating Variables—
New technologies, such as driver telematics, offer the opportunity for insurance companies to collect additional types of information and greater detail of information than can be collected up front through policyholder applications. This new and increased data could make other factors used in auto insurance less powerful, or completely redundant, potentially allowing for the deletion of some variables. For example, evaluating a driver's braking habits may provide additional information about an individual's risk of accidents. Tracking the locations that a person drove in a month may be more reflective of insurance risk than simply reflecting a vehicle's garaging location.

Compared to simpler methods, this method is more difficult to explain to consumers and may be more difficult for some carriers to implement, which could have marketplace implications. This method would allow rates to continue to differentiate based on expected costs. It embraces new data, innovation, and technology and does not identify unfairly discriminatory variables or directly correct for bias. It also assumes that by adding variables that increase segmentation and are more directly related to loss, variables less correlated with losses would be removed.

Conclusion

These are methods among many others likely to emerge over time to identify and address unfair discrimination, and the Casualty Practice Council of 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 above.