



City Research Online

City, University of London Institutional Repository

Citation: Lindholm, M., Tsanakas, A., Richman, R. & Wüthrich, M. V. (2022). A Discussion of Discrimination and Fairness in Insurance Pricing. .

This is the preprint version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/29201/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

A DISCUSSION OF DISCRIMINATION AND FAIRNESS IN INSURANCE PRICING

Mathias Lindholm Ronald Richman Andreas Tsanakas Mario V. Wüthrich

Version of September 2, 2022

Abstract

Indirect discrimination is an issue of major concern in algorithmic models. This is particularly the case in insurance pricing where protected policyholder characteristics are not allowed to be used for insurance pricing. Simply disregarding protected policyholder information is not an appropriate solution because this still allows for the possibility of inferring the protected characteristics from the non-protected ones. This leads to so-called proxy or indirect discrimination. Though proxy discrimination is qualitatively different from the group fairness concepts in machine learning, these group fairness concepts are proposed to ‘smooth out’ the impact of protected characteristics in the calculation of insurance prices. The purpose of this note is to share some thoughts about group fairness concepts in the light of insurance pricing and to discuss their implications. We present a statistical model that is free of proxy discrimination, thus, unproblematic from an insurance pricing point of view. However, we find that the canonical price in this statistical model does not satisfy any of the three most popular group fairness axioms. This seems puzzling and we welcome feedback on our example and on the usefulness of these group fairness axioms for non-discriminatory insurance pricing.

Keywords. Discrimination, indirect discrimination, proxy discrimination, fairness, protected information, discrimination-free, unawareness, group fairness, statistical parity, independence axiom, equalized odds, separation axiom, predictive parity, sufficiency axiom.

1 Introduction

For legal and for societal reasons, there are several policyholder attributes that are not allowed to be used for insurance pricing, e.g., the European Council [7, 6] does not allow for the use of gender information in insurance pricing, or ethnicity is a critical attribute that may be declared as protected information by society. Frees–Huang [8] and Xin–Huang [21] give extensive overviews on protected information in insurance and implications for pricing, while Avraham et al. [3] and Prince–Schwarz [16] provide legal viewpoints on this topic. The critical issue is that just ignoring protected information (being unaware of protected information) does not solve the problem, as protected information can be inferred from non-protected one if corresponding variables are correlated. This is especially true in high-dimensional algorithmic models. Such inference is called proxy discrimination or indirect discrimination, and it is often implicitly performed during the fitting procedure of complex models.

There are several attempts to prevent this inference, e.g., there is a counterfactual approach from causal statistics, see Kusner et al. [11] and Araiza Iturria et al. [2], there is a probabilistic

approach called discrimination-free insurance pricing by Lindholm et al. [12, 13], and there are group fairness approaches inspired by machine learning concepts, see Grari et al. [9]. The purpose of this note is to discuss the three most popular group fairness axioms in the light of insurance pricing. We define a statistical model that is free of discrimination because the non-protected information fully describes the distribution of the response variable, and the protected one does not carry any additional information about the response. We analyze the three group fairness axioms in this model. Surprisingly, the canonical non-discriminatory price (which is uncontroversial here) does not satisfy any of these three group fairness axioms. This seems puzzling and it requires a broader discussion within our community about discrimination and fairness in insurance pricing. Therefore, we share our thoughts and welcome feedback on our example and on the use of the group fairness axioms for non-discriminatory pricing.

Organization. We introduce our statistical model that is free of discrimination in Section 2, and we calculate the best-estimate and the unawareness prices in this model. In Section 3, we formally introduce the three most popular group fairness axioms; the independence axiom (statistical parity), the separation axiom (equalized odds) and the sufficiency axiom (predictive parity); see Barocas et al. [4], Xin–Huang [21] and Grari et al. [9]. We show that in our model the canonical non-discriminatory price does not satisfy any of these three axioms. Finally, in Section 4 we conclude and discuss further issues. All mathematical results are proved in Appendix B.

2 Running example

To set the ground, we fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with \mathbb{P} describing the real world probability measure. On this probability space we consider the random vector $(Y, \mathbf{X}, \mathbf{D})$. The response Y describes the insurance claim that we try to predict (and price). The vector \mathbf{X} describes the *non-protected characteristics (non-discriminatory covariates)*, and \mathbf{D} describes the *protected characteristics (discriminatory covariates)*. We assume that the partition of the covariates into protected and non-protected ones is given exogenously, e.g., by law or by societal norms and preferences. We use the distribution $\mathbb{P}(\mathbf{X}, \mathbf{D})$ to describe the insurance portfolio, i.e., the random selection of a policy from the insurance portfolio. Different insurance companies may have different such distributions, and the insurance portfolio distribution $\mathbb{P}(\mathbf{X}, \mathbf{D})$ typically differs from the overall population distribution in a given society.

Best-estimate price. For insurance pricing, one designs a regression model that describes the conditional distribution of Y , given the explanatory variables (\mathbf{X}, \mathbf{D}) . The *best-estimate price* of Y , given full information (\mathbf{X}, \mathbf{D}) , is given by

$$\mu(\mathbf{X}, \mathbf{D}) = \mathbb{E}[Y | \mathbf{X}, \mathbf{D}]. \quad (2.1)$$

This price is called best-estimate because it has minimal prediction variance,¹ i.e., it is the most accurate predictor for Y , given (\mathbf{X}, \mathbf{D}) , in the L^2 -sense.

 In general, the best-estimate price *directly discriminates* because it uses the protected information \mathbf{D} as an input.

¹For simplicity, we assume that all considered random variables are square-integrable w.r.t. \mathbb{P} .

Fairness through unawareness. The most simple fairness concept in machine learning is the *fairness through unawareness* concept that just drops the protected information \mathbf{D} . This results in the *unawareness price* of Y , given \mathbf{X} , defined by

$$\mu(\mathbf{X}) = \mathbb{E}[Y | \mathbf{X}]. \quad (2.2)$$

The unawareness price does not *directly discriminate* because it does not use \mathbf{D} as an input, i.e., it is blind w.r.t. \mathbf{D} . However, it may *indirectly discriminate* because the knowledge of \mathbf{X} allows inference of \mathbf{D} through the tower property of conditional expectation

$$\mu(\mathbf{X}) = \int \mu(\mathbf{X}, \mathbf{d}) d\mathbb{P}(\mathbf{D} = \mathbf{d} | \mathbf{X}). \quad (2.3)$$

This formula shows that if there is statistical dependence between \mathbf{D} and \mathbf{X} w.r.t. \mathbb{P} , we implicitly use this dependence for inference, see last term in (2.3). This is what is called proxy discrimination and indirect discrimination, and it should be avoided.

☞ Fairness through unawareness indirectly discriminates, and it does not solve the problem of non-discriminatory insurance pricing.

We now present our example. In this example we have a response variable Y whose distribution function is fully described by the non-protected information \mathbf{X} , and the protected information \mathbf{D} does not carry any additional information about the response Y .² Therefore, we consider this statistical model to be free of discrimination. This model is simple enough to be able to calculate all terms of interest, and, even if it is unrealistic in practice, it allows us to draw conclusions about the group fairness axioms and discrimination in insurance pricing.

Example 2.1 (running example) Assume we have three-dimensional covariates $(\mathbf{X}, \mathbf{D}) = (X_1, X_2, D)$ having a multi-variate Gaussian portfolio distribution

$$(X_1, X_2, D)^\top \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \Sigma = \begin{pmatrix} 1 & 0 & \rho_1 \\ 0 & 1 & \rho_2 \\ \rho_1 & \rho_2 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0.1 \\ 0 & 1 & 0.9 \\ 0.1 & 0.9 & 1 \end{pmatrix} \right). \quad (2.4)$$

Thus, the non-protected covariates X_1 and X_2 are independent, but they are both dependent with the protected covariate D . We assume that this dependence is a purely statistical one implied by the portfolio structure of the considered insurance company, and not a causal one. That is, there might be a second insurance company with a portfolio distribution $\mathbb{P}(\mathbf{X}, \mathbf{D})$ described by a Gaussian distribution (2.4), but having, e.g., negative correlations between \mathbf{X} and \mathbf{D} .

For the response Y we assume that conditionally, given (\mathbf{X}, \mathbf{D}) ,

$$Y |_{(\mathbf{X}, \mathbf{D})} \sim \mathcal{N}(X_1, 1 + X_2^2). \quad (2.5)$$

That is, the response does *not* depend on the protected information \mathbf{D} , but only on the non-protected information \mathbf{X} . This means that \mathbf{X} is sufficient to describe the distribution of Y , and \mathbf{D} does not carry any additional information about Y , see also Footnote 2.

²Araiza Iturria et al. [2] call such a variable \mathbf{D} not a true risk factor, because \mathbf{X} is sufficient to describe the distribution of Y , see also European Commission [6].

The best-estimate and the unawareness prices coincide in this example, and they are given by

$$\mu(\mathbf{X}, \mathbf{D}) = \mu(\mathbf{X}) = X_1. \quad (2.6)$$

In this example, there is no indirect discrimination through a conditional expectation (2.3) because the best-estimate price coincides with the unawareness price, and the explicit structure of the portfolio distribution $\mathbb{P}(\mathbf{X}, \mathbf{D})$ is *not* needed here. Therefore, in this example, we call these prices *non-discriminatory*.³ From an actuarial viewpoint this price is non-discriminatory because the protected information \mathbf{D} has no influence on the response Y , given \mathbf{X} . This is true for *any* dependence structure between \mathbf{D} and \mathbf{X} , and no matter whether it is a causal or a purely statistical one. This is an additional justification for not having indirect discrimination. ■

Remark 2.2 Under the assumptions of Example 2.1, the non-discriminatory price $\mu(\mathbf{X})$ given in (2.6) is equal to the discrimination-free insurance price of Lindholm et al. [12].

3 Group fairness axioms

In this section, we introduce the three most popular group fairness axioms from machine learning; we also refer to Barocas et al. [4], Xin–Huang [21] and Grari et al. [9]. We will show that the non-discriminatory price of Example 2.1 given in (2.6) violates all three of these group fairness axioms.

We denote by $\hat{\mu}(\mathbf{X})$ any $\sigma(\mathbf{X})$ -measurable predictor of Y , which can be the unawareness price (2.2) or any other pricing functional that solely depends on the non-protected information \mathbf{X} .

(i) Independence axiom / statistical parity. Statistical parity is also called demographic parity, and following Definition 1 of Agarwal et al. [1] it is given as follows: We have *statistical parity* if

$$\hat{\mu}(\mathbf{X}) \text{ and } \mathbf{D} \text{ are independent under } \mathbb{P}.$$

This independence implies for the distribution of the insurance prices $\hat{\mu}(\mathbf{X})$, a.s.,

$$\mathbb{P}[\hat{\mu}(\mathbf{X}) \leq m | \mathbf{D}] = \mathbb{P}[\hat{\mu}(\mathbf{X}) \leq m] \quad \text{for all } m \in \mathbb{R}. \quad (3.1)$$

We have the following proposition.

Proposition 3.1 *The non-discriminatory price $\mu(\mathbf{X}) = X_1$ of Example 2.1 does not satisfy the independence axiom (statistical parity).*

 The non-discriminatory price of Example 2.1 violates the independence axiom (statistical parity). This violation is implied by the dependence of the protected and the non-protected information, which is described by the parameters (ρ_1, ρ_2) in Example 2.1. Remark that every

³More generally, the definition of the non-discriminatory price is uncontroversial in insurance in examples where the non-protected information \mathbf{X} is sufficient to describe the response Y , and the additional knowledge of the protected information \mathbf{D} is not needed because it does not reveal any additional information about the response Y ; we come back to this in formula (3.2) and in Remark 3.5, below.

insurance company may have a different portfolio distribution $\mathbb{P}(\mathbf{X}, \mathbf{D})$, e.g., Company 1 may have the portfolio distribution (2.4) with a positive correlation $\rho_1 = 0.1$ and Company 2 may have this portfolio distribution but with independence $\rho_1 = 0$. The independence axiom implies that Company 1 cannot use information X_1 for (statistical parity-fair) insurance pricing, whereas Company 2 would be allowed to include information X_1 in its tariff. Such a regulation would be very problematic as it generates unwanted distortions at the insurance market, in a situation where we have a pricing problem that is apparently free of discrimination (i.e., in Example 2.1). The only price that Company 1 could charge under the independence axiom is $\mu = \mathbb{E}[Y]$, which corresponds to the null model not considering any covariates. However, this is not in line with, e.g., Article 2.3.1(17) of the guidelines of the European Commission [6] which explicitly allows for the use of covariate X_1 . We conclude that the independence axiom (statistical parity) may be too restrictive in insurance pricing because the use of specific covariates will differ across insurance companies, and the resulting (statistical parity-fair) insurance pricing functionals may be impractical.

Remark 3.2 Assume there is a decomposition $\mathbf{X} = (\mathbf{U}, \mathbf{V})$ such that the protected information \mathbf{D} is independent from \mathbf{U} but not from \mathbf{V} . Moreover, as in Example 2.1, assume that, a.s.,

$$\mathbb{P}[Y \leq y | \mathbf{X}, \mathbf{D}] = \mathbb{P}[Y \leq y | \mathbf{X}] \quad \text{for all } y \in \mathbb{R}. \quad (3.2)$$

In this case, the best-estimate price coincides with the unawareness price, and we generally set for the non-discriminatory price $\mu(\mathbf{X}) = \mu(\mathbf{X}) = \mathbb{E}[Y | \mathbf{X}]$, see Footnote 3 on page 4. This non-discriminatory price does not satisfy the independence axiom in general. On the other hand, the insurance price based only on \mathbf{U} is equal to $\mu(\mathbf{U}) = \mathbb{E}[Y | \mathbf{U}]$, and it satisfies the independence axiom because \mathbf{U} and \mathbf{D} are independent by assumption. If we insist on using the independence axiom, here, we discard a possibly large part of the non-protected information \mathbf{X} , here \mathbf{V} . This seems to be too restrictive, because assumption (3.2) provides a statistical model that does not require any knowledge about \mathbf{D} (and of the portfolio distribution $\mathbb{P}(\mathbf{X}, \mathbf{D})$), and it is also not in line with the guidelines of the European Commission [6].

(ii) Separation axiom / equalized odds. Equalized odds is sometimes also called disparate mistreatment. It has been introduced by Hardt et al. [10], and it is defined as follows: We have *equalized odds* if

$$\hat{\mu}(\mathbf{X}) \text{ and } \mathbf{D} \text{ are conditionally independent under } \mathbb{P}, \text{ given the response } Y.$$

This conditional independence implies for the distribution of the prices $\hat{\mu}(\mathbf{X})$, a.s.,

$$\mathbb{P}[\hat{\mu}(\mathbf{X}) \leq m | Y, \mathbf{D}] = \mathbb{P}[\hat{\mu}(\mathbf{X}) \leq m | Y] \quad \text{for all } m \in \mathbb{R}. \quad (3.3)$$

Note that, in general, independence between \mathbf{X} and \mathbf{D} is not sufficient to receive equalized odds.

Proposition 3.3 *The non-discriminatory price $\mu(\mathbf{X}) = X_1$ of Example 2.1 does not satisfy the separation axiom (equalized odds).*

 The non-discriminatory price of Example 2.1 violates the separation axiom (equalized odds). This violation is again implied by the dependence of the protected and the non-protected information, which is described by the parameters (ρ_1, ρ_2) in Example 2.1. Therefore, basically, the same remarks apply as for the failure of the independence axiom.

(iii) **Sufficiency axiom / predictive parity.** For predictive parity we exchange the role of the response Y and the predictor $\hat{\mu}(\mathbf{X})$ compared to equalized odds. We have *predictive parity* if

Y and \mathbf{D} are conditionally independent under \mathbb{P} , given the prediction $\hat{\mu}(\mathbf{X})$.

This conditional independence implies for the distribution of the response Y , a.s.,

$$\mathbb{P}[Y \leq y | \hat{\mu}(\mathbf{X}), \mathbf{D}] = \mathbb{P}[Y \leq y | \hat{\mu}(\mathbf{X})] \quad \text{for all } y \in \mathbb{R}. \quad (3.4)$$

The notion of predictive parity is inspired by the definition of a sufficient statistics in statistical estimation theory. We denote the support of the protected information \mathbf{D} by \mathfrak{D} . We can then interpret $\mathcal{P} = \{\mathbb{P}_{\mathbf{d}}[Y \in \cdot] := \mathbb{P}[Y \in \cdot | \mathbf{D} = \mathbf{d}]; \mathbf{d} \in \mathfrak{D}\}$ as a family of distributions of Y being parametrized by $\mathbf{d} \in \mathfrak{D}$. In statistics we call $\hat{\mu}(\mathbf{X})$ sufficient for \mathcal{P} if (3.4) holds. Basically, this means that $\hat{\mu}(\mathbf{X})$ carries all the necessary information to predict Y , and the explicit knowledge of $\mathbf{D} = \mathbf{d}$ is not necessary.

Proposition 3.4 *The non-discriminatory price $\mu(\mathbf{X}) = X_1$ of Example 2.1 does not satisfy the sufficiency axiom (predictive parity).*

 The non-discriminatory price of Example 2.1 violates the sufficiency axiom (predictive parity). This violation is again implied by the dependence of the protected and the non-protected information, and the same remarks apply as for the failure of the independence and the separation axioms.

Remark 3.5 We briefly give the intuition why the sufficiency axiom fails in Example 2.1. The sufficiency axiom (3.4) is formulated w.r.t. the knowledge of the predictor $\hat{\mu}(\mathbf{X})$. This is different from a sufficiency condition w.r.t. the non-protected information \mathbf{X} where, a.s.,

$$\mathbb{P}[Y \leq y | \mathbf{X}, \mathbf{D}] = \mathbb{P}[Y \leq y | \mathbf{X}] \quad \text{for all } y \in \mathbb{R}. \quad (3.5)$$

The condition (3.5) has already been considered in (3.2), and we also refer to Footnote 2 on page 3. In general, the information set generated by \mathbf{X} is bigger than the information set generated by $\hat{\mu}(\mathbf{X})$. Therefore, the two sufficiency conditions typically differ. More specifically, the price functional $\mathbf{X} \mapsto \hat{\mu}(\mathbf{X})$ is a projection of the (multi-dimensional) covariate information \mathbf{X} to the sub-space generated by the one-dimensional prices $\hat{\mu}(\mathbf{X})$. In general, such a projection leads to a loss of information which is described by the complement of this projection (in the case of our Example 2.1 this is X_2). If the protected information \mathbf{D} is dependent with this complement, it will partly compensate for this loss of information which results in the failure of the sufficiency axiom (in case this complement is needed to fully characterize the distribution of Y).

4 Discussion and open questions

We have defined a statistical model in Example 2.1 where the non-protected information \mathbf{X} is sufficient to describe the distribution of the response variable Y , and the additional knowledge of the protected information \mathbf{D} does not reveal any additional information about this response variable Y ; Araiza Iturria et al. [2] call such a variable \mathbf{D} not a true risk factor. In this

statistical modeling set-up, there is a canonical non-discriminatory price because the knowledge of \mathbf{D} is completely irrelevant for best-estimate pricing, in fact, proxy discrimination (2.3) does not happen in this set-up. In Propositions 3.1, 3.3 and 3.4 we show that this non-discriminatory price does not satisfy any of the three group fairness axioms of 'statistical parity', 'equalized odds' and 'predictive parity'. In other words, these three group fairness axioms seem to exclude a huge class of reasonable pricing functionals. Moreover, this exclusion is insurance company-dependent, i.e., different companies may be forced to exclude different non-protected information (due to different portfolio distributions). This seems disputable, and it requires a broader discussion within our community under which circumstances we should postulate one of these three group fairness axioms; this is further highlighted in the examples of Appendix A.

As a consequence, the three group fairness axioms considered here do not seem to provide a quick fix of indirect discrimination, or they may but at a too restrictive price. In some sense, these fairness axioms aim at mitigating discrimination in a post-hoc way which does not seem to be appropriate in every situation, e.g., under the assumptions of Example 2.1.

We state a number of further points that are worth studying in more detail.

- In view of the previous conclusions, one could think of other fairness criteria that one may want to postulate. Alternatively, one may ask under which circumstances the group fairness axioms should be fulfilled, and how they can complement the notion of proxy discrimination in insurance, as the two concepts seem to be qualitatively different.
- Group fairness axioms are typically formulated in the binary classification case, and one may explore whether the one-to-one extension to continuous responses is sensible. Generally speaking, the binary classification case (Bernoulli model) is a one-parameter model where the expected value (parameter) fully describes the distribution of the response variable. Similarly, there are continuous models which are fully parametrized by their expected values, and in these cases the group fairness criteria will essentially behave the same as in the binary classification case. From an actuarial perspective, Tweedie's models (belonging to the exponential dispersion family) are often used with the mean function depending on covariates and a constant dispersion parameter. These models are examples of one-parameter models being fully described by their mean regression functional.

Our example is different because the response Y has an expected value of $\mu(\mathbf{X}) = X_1$ and a variance of $1 + X_2^2$. Therefore, the response distribution is not fully characterized by its mean, and applying the mean functional leads to the loss of information of X_2 . This may motivate a search for weaker forms of the independence, separation and sufficiency axioms, which do not rely on the full distribution of the response variable.

- Recently, notions of conditional group fairness have gained more popularity. We could further explore how these behave on simple insurance pricing examples. Intuitively, these conditional versions will not essentially behave differently from our outline above, but they could still be useful as they study fairness criteria on smaller units of the covariate space.
- There are many different terms in the field of discrimination such as disparate effect, disparate impact, disproportionate impact, etc.; see, e.g., Chibanda [5]. It would be desirable to translate these terms into more mathematical definitions.

- Adverse selection and unwanted economic consequences of non-discriminatory pricing should be explored. Non-discriminatory prices typically fail to fulfill the auto-calibration property which is crucial for having homogeneous risk classes, see Wüthrich [19].
- All considerations above have been based on the assumption that we know the true model. Clearly, in statistical modeling there is model uncertainty which may impact different protected classes differently because, e.g., they are represented differently in historical data (statistical and historical biases). There are several examples of this type in the machine learning literature; see, e.g., Barocas et al. [4], Mehrabi et al. [14] and Pessach–Shmueli [15].
- Non-protected covariates may be context-sensitive. E.g., the European Commission [6], footnote 1 to Article 2.2(14) – life and health underwriting – mentions the waist-to-hip ratio as a non-protected (useful) covariate for health prediction. Note the following:
 - The waist-to-hip ratio is gender dependent.
 - The waist-to-hip ratio is age dependent.
 - The waist-to-hip ratio is race dependent.
 - The waist-to-hip ratio for females depends on the number of born children.
 - Etc.

Thus, a waist-to-hip ratio X (if non-protected) can only be correctly interpreted under the knowledge of protected information. This will require that we pre-process non-protected characteristics to common units. A similar situation may, e.g., occur with age, as the chronological age may be predictive for car driving experience, whereas in life insurance we should rather choose a biological age (which uses gender information). This raises the general issue of how to consider non-protected information that may be context-sensitive and may only be fully explanatory under the inclusion of protected information. Generally, this will result in the question of how to pre-process non-protected information.

- We have been speaking about (non-)discrimination of insurance prices. These insurance prices are actuarial or statistical prices (technical premium), i.e., they directly result as an output from a statistical procedure. These prices are then modified to commercial prices, e.g., administrative costs are added, etc. An interesting issue is raised by Thomas [17, 18], namely, by converting actuarial prices into commercial prices one often distorts these prices with elasticity considerations, i.e., insurance companies charge higher prices to customers that are (implicitly) willing to pay more. This happens, e.g., with new business and contract renewals that are often priced differently though the corresponding customers may have exactly the same risk profile. In the light of discrimination and fairness one should also clearly question such practice of elasticity pricing as this leads to discrimination that cannot be explained by risk profiles (no matter whether we consider protected or non-protected information).
- Related to the previous item on converting the technical premium into a commercial price. The technical premium (actuarial or statistical price) focuses on the expected response for given covariates. It might be that this statistical price is accepted as being free of

discrimination. However, by converting a statistical price into a commercial price one also needs to add a risk margin (related to the prediction uncertainty in the statistical price), and it can well be that this risk margin (loading) adds discrimination to the (acceptable) statistical price.

References

- [1] Agarwal, A., Dudik, M., Wu, Z.S. (2019). Fair regression: quantitative definitions and reduction-based algorithms. *arXiv:1905.12843*
- [2] Araiza Iturria, C.A., Hardy, M., Marriott, P. (2022). A discrimination-free premium under a causal framework. *SSRN Manuscript ID 4079068*.
- [3] Avraham, R., Logue, K. D. and Schwarcz, D.B. (2014). Understanding insurance anti-discrimination laws. *Southern California Law Review* **87(2)**, 195-274.
- [4] Barocas, S., Hardt, M., Narayanan, A. (2019). *Fairness and Machine Learning: Limitations and Opportunities*. <https://fairmlbook.org/>
- [5] Chibanda, K.F. (2021). Defining discrimination in insurance. *CAS Research Papers: A Special Series on Race and Insurance Pricing*. <https://www.casact.org/publications-research/research/research-paper-series-race-and-insurance-pricing>
- [6] European Commission (2012). Guidelines on the application of COUNCIL DIRECTIVE 2004/113/EC to insurance, in the light of the judgment of the Court of Justice of the European Union in Case C-236/09 (Test-Achats). *Official Journal of the European Union* **C11**, 1-11.
- [7] European Council (2004). COUNCIL DIRECTIVE 2004/113/EC - implementing the principle of equal treatment between men and women in the access to and supply of goods and services. *Official Journal of the European Union* **L 373**, 37-43.
- [8] Frees, E.W.J., Huang, F. (2022). The discriminating (pricing) actuary. *North American Actuarial Journal*, in press.
- [9] Grari, V., Charpentier, A., Lamprier, S., Detyniecki, M. (2022). A fair pricing model via adversarial learning. *arXiv:2202.12008v2*.
- [10] Hardt, M., Price, E., Srebro, N. (2016). Equality of opportunity in supervised learning. In: *Advances in Neural Information Processing Systems*, 3315-3323.
- [11] Kusner, M.J., Loftus, J., Russell, C. and Silva, R. (2017). Counterfactual fairness. In: *Advances in Neural Information Processing Systems*, 4066-4076.
- [12] Lindholm, M., Richman, R., Tsanakas, A., Wüthrich, M.V. (2022). Discrimination-free insurance pricing. *ASTIN Bulletin* **52(2)**, 55-89.
- [13] Lindholm, M., Richman, R., Tsanakas, A., Wüthrich, M.V. (2022). A multi-output network approach for calculating discrimination-free insurance prices. *SSRN Manuscript ID 4155585*.
- [14] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A. (2019). A survey on bias and fairness in machine learning. *arXiv:1908.09635v3*.
- [15] Pessach, D., Erez Shmueli, E. (2022). A review on fairness in machine learning. *ACM Computing Survey* **55(3)**, article 51.
- [16] Prince, A.E.R., Schwarcz, D. (2020). Proxy discrimination in the age of artificial intelligence and big data. *Iowa Law Review* **105(3)**, 1257-1318.
- [17] Thomas, R.G. (2012). Non-risk price discrimination in insurance: market outcomes and public policy. *Geneva Papers on Risk and Insurance - Issues and Practice* **37**, 27-46.

- [18] Thomas, R.G. (2022). Discussion on “The discriminating (pricing) actuary”, by E.W.J. Frees and F. Huang. *North American Actuarial Journal*, in press.
- [19] Wüthrich, M.V. (2022). Model selection with Gini indices under auto-calibration. *arXiv:2207.14372*.
- [20] Wüthrich, M.V., Merz, M. (2015). Stochastic claims reserving manual: advances in dynamic modeling. *SSRN Manuscript ID 264905*.
- [21] Xin, X., Huang, F. (2021). Anti-discrimination insurance pricing: regulations, fairness criteria, and models. *SSRN Manuscript ID 3850420*.

A Appendix: further observations from our example

We give further observations on Example 2.1, but we state them as conjectures to keep this discussion paper short (though some of them would not be difficult to prove).

We come back to the Gaussian distribution (2.4), but we do not explicitly specify the values of the covariance parameters

$$(X_1, X_2, D)^\top \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \Sigma = \begin{pmatrix} 1 & 0 & \rho_1 \\ 0 & 1 & \rho_2 \\ \rho_1 & \rho_2 & 1 \end{pmatrix} \right).$$

As in (2.5), for the response Y we assume that conditionally, given $(\mathbf{X}, \mathbf{D}) = (X_1, X_2, D)$,

$$Y|_{(\mathbf{X}, \mathbf{D})} \sim \mathcal{N}(X_1, 1 + X_2^2).$$

We conjecture the following results about the three group fairness axioms.

	group fairness axiom		
	independence	separation	sufficiency
$\rho_1 > 0$ and $\rho_2 > 0$	NO	NO	NO
$\rho_1 > 0$ and $\rho_2 = 0$	NO	NO	YES
$\rho_1 = 0$ and $\rho_2 > 0$	YES	NO	NO
$\rho_1 = 0$ and $\rho_2 = 0$	YES	YES	YES

The conjectures on lines 2 and 3 are also supported by the fact that under suitable assumptions different group fairness criteria cannot jointly hold true; see, e.g., Chapter 3 in Barocas [4].

B Appendix: mathematical proofs

We prove the above statements of Propositions 3.1, 3.3 and 3.4 in this appendix.

Proof of Proposition 3.1. Under the assumptions of Example 2.1 we have

$$\text{Cov}(\mu(\mathbf{X}), \mathbf{D}) = \text{Cov}(X_1, D) = \rho_1 = 0.1 > 0.$$

Thus, we have a positive correlation, and (3.1) cannot hold. This completes the proof. \square

Proof of Proposition 3.3. In view of the separation axiom (3.3) we need to prove, a.s.,

$$\mathbb{P}[\mu(\mathbf{X}) \leq m | Y, \mathbf{D}] = \mathbb{P}[\mu(\mathbf{X}) \leq m | Y] \quad \text{for all } m \in \mathbb{R}.$$

Under the assumptions of Example 2.1 we have $\mu(\mathbf{X}) = X_1$ and $\mathbf{D} = D$, and if the separation axiom holds we need to have for all measurable functions $h : \mathbb{R} \rightarrow \mathbb{R}$

$$\mathbb{E}[h(X_1) | Y, D] = \mathbb{E}[h(X_1) | Y], \tag{B.1}$$

a.s., where these exist. We will give a counterexample to (B.1) which proves that the separation axiom cannot hold in Example 2.1. Because of continuity it suffices to prove that

$$\mathbb{E}[X_1^2 | Y = 0, D = 0] < \mathbb{E}[X_1^2 | Y = 0] < \mathbb{E}[X_1^2] = 1. \tag{B.2}$$

We give the rationale behind conjecture (B.2). Observe that the distribution of X_1 is symmetric around the origin under the Gaussian assumption (2.4). Therefore, the right-hand side of (B.2) is precisely the variance of X_1 . If we condition on $Y = 0$, the distribution of X_1 is still symmetric around the origin, therefore, the middle term in (B.2) is the conditional variance of X_1 , given $Y = 0$. Since $Y = 0$, we expect X_1 (being the mean of Y) to be

closer to zero, which implies that the conditional variance of X_1 , given Y , should be smaller than 1. Additionally conditioning on $D = 0$, still keeps X_1 symmetric around zero, and the positive correlation $\rho_1 = 0.1$ should further decrease the conditional variance, which motivates the first inequality in (B.2).

We start by analyzing the left-hand side of (B.2). We express the distribution of (Y, X_1, X_2) conditionally, given $D = 0$. Using a standard result from multi-variate Gaussian distributions, see, e.g., Corollary 4.4 of Wüthrich–Merz [20], we have conditional portfolio distribution

$$\mathbf{X}|_{D=0} = (X_1, X_2)^\top|_{D=0} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_D = \begin{pmatrix} 1 - \rho_1^2 & -\rho_1\rho_2 \\ -\rho_1\rho_2 & 1 - \rho_2^2 \end{pmatrix}\right). \quad (\text{B.3})$$

For the response Y conditionally given $(\mathbf{X}, D = 0)$, we still have

$$Y|_{(\mathbf{X}, D=0)} \sim \mathcal{N}(X_1, 1 + X_2^2). \quad (\text{B.4})$$

Remark that (B.3) would coincide with the distribution of \mathbf{X} , if we would have independence $\rho_1 = \rho_2 = 0$, see (2.4). Therefore, all considerations that we do under (B.3) can be translated to the middle term of (B.2) by just setting these correlation parameters equal to zero.

Thus, we aim at calculating for $h(x) = x$ and $h(x) = x^2$ the conditional expectation

$$\mathbb{E}[h(X_1)|Y = 0, D = 0] = \mathbb{E}_0[h(X_1)|Y = 0],$$

where we abbreviate the conditional probability measure, given $D = 0$, by \mathbb{P}_0 . This motivates us to consider

$$\mathbb{E}[h(X_1)|Y, D = 0] = \mathbb{E}_0[h(X_1)|Y] = \mathbb{E}_0[\mathbb{E}_0[h(X_1)|Y, X_2]|Y]. \quad (\text{B.5})$$

The joint density of $(Y, X_1, X_2)|_{D=0} \sim f_{Y, X_1, X_2}^{(0)}$ is given by

$$\begin{aligned} f_{Y, X_1, X_2}^{(0)}(y, x_1, x_2) &= \frac{1}{\sqrt{2\pi(1+x_2^2)}} \exp\left\{-\frac{1}{2} \frac{(y-x_1)^2}{1+x_2^2}\right\} \\ &\times \frac{1}{2\pi|\Sigma_D|^{1/2}} \exp\left\{-\frac{1}{2} \frac{x_1^2(1-\rho_2^2) + x_2^2(1-\rho_1^2) + 2x_1x_2\rho_1\rho_2}{1-\rho_1^2-\rho_2^2}\right\}. \end{aligned}$$

This gives for the conditional density of X_1 , given $(Y, X_2, D = 0)$, we drop normalizing constants in the proportionality sign \propto ,

$$\begin{aligned} f_{X_1|Y, X_2}^{(0)}(x_1|Y, X_2) &\propto \exp\left\{-\frac{1}{2} \frac{(Y-x_1)^2}{1+X_2^2}\right\} \exp\left\{-\frac{1}{2} \frac{x_1^2(1-\rho_2^2) + 2x_1X_2\rho_1\rho_2}{1-\rho_1^2-\rho_2^2}\right\} \\ &\propto \exp\left\{-\frac{1}{2} \left(\frac{x_1^2 - 2x_1Y}{1+X_2^2} + \frac{x_1^2(1-\rho_2^2) + 2x_1X_2\rho_1\rho_2}{1-\rho_1^2-\rho_2^2}\right)\right\} \\ &\propto \exp\left\{-\frac{1}{2} \left(\frac{x_1^2((2+X_2^2)(1-\rho_2^2) - \rho_1^2) - 2x_1(Y(1-\rho_1^2-\rho_2^2) - \rho_1\rho_2X_2(1+X_2^2))}{(1+X_2^2)(1-\rho_1^2-\rho_2^2)}\right)\right\}. \end{aligned}$$

This is a Gaussian density, and we have

$$X_1|_{(Y=0, X_2, D=0)} \sim \mathcal{N}\left(\frac{-\rho_1\rho_2X_2(1+X_2^2)}{(2+X_2^2)(1-\rho_2^2) - \rho_1^2}, \frac{(1+X_2^2)(1-\rho_1^2-\rho_2^2)}{(2+X_2^2)(1-\rho_2^2) - \rho_1^2}\right). \quad (\text{B.6})$$

We calculate (B.5) for the function $h(x) = x$, conditional on having a response $Y = 0$ and $D = 0$,

$$\mathbb{E}_0[X_1|Y = 0] = \mathbb{E}_0\left[\frac{-\rho_1\rho_2X_2(1+X_2^2)}{(2+X_2^2)(1-\rho_2^2) - \rho_1^2} \middle| Y = 0\right]. \quad (\text{B.7})$$

Intuitively, this should be equal to 0 which we are going to verify. The joint density of $(X_1, X_2)|_{Y=0, D=0}$ is given by

$$\begin{aligned} f_{X_1, X_2|Y=0}^{(0)}(x_1, x_2) &\propto \frac{1}{\sqrt{1+x_2^2}} \exp\left\{-\frac{1}{2} \left(\frac{x_1^2}{1+x_2^2} + \frac{x_1^2(1-\rho_2^2) + x_2^2(1-\rho_1^2) + 2x_1x_2\rho_1\rho_2}{1-\rho_1^2-\rho_2^2}\right)\right\} \\ &= \frac{1}{\sqrt{1+x_2^2}} \exp\left\{-\frac{1}{2} \left(\frac{x_1^2((2+x_2^2)(1-\rho_2^2) - \rho_1^2) + 2x_1x_2\rho_1\rho_2(1+x_2^2)}{(1+x_2^2)(1-\rho_1^2-\rho_2^2)} + \frac{x_2^2(1-\rho_1^2)}{1-\rho_1^2-\rho_2^2}\right)\right\}. \end{aligned}$$

This allows us to calculate the density of $X_2|Y=0, D=0$ by integrating the previous density over x_1

$$\begin{aligned}
f_{X_2|Y=0}^{(0)}(x_2) &= \int f_{X_1, X_2|Y=0}^{(0)}(x_1, x_2) dx_1 \\
&\propto \frac{1}{\sqrt{1+x_2^2}} \frac{\sqrt{(1+x_2^2)(1-\rho_1^2-\rho_2^2)}}{\sqrt{(2+x_2^2)(1-\rho_2^2)-\rho_1^2}} \exp\left\{-\frac{1}{2} \frac{x_2^2(1-\rho_1^2)}{1-\rho_1^2-\rho_2^2}\right\} \\
&\quad \times \int \frac{\sqrt{(2+x_2^2)(1-\rho_2^2)-\rho_1^2}}{\sqrt{(1+x_2^2)(1-\rho_1^2-\rho_2^2)}} \exp\left\{-\frac{1}{2} \frac{x_1^2 + 2x_1x_2\rho_1\rho_2(1+x_2^2)/((2+x_2^2)(1-\rho_2^2)-\rho_1^2)}{(1+x_2^2)(1-\rho_1^2-\rho_2^2)/((2+x_2^2)(1-\rho_2^2)-\rho_1^2)}\right\} dx_1 \\
&\propto \frac{1}{\sqrt{1+x_2^2}} \frac{\sqrt{(1+x_2^2)(1-\rho_1^2-\rho_2^2)}}{\sqrt{(2+x_2^2)(1-\rho_2^2)-\rho_1^2}} \\
&\quad \times \exp\left\{-\frac{1}{2} \frac{x_2^2(1-\rho_1^2)((2+x_2^2)(1-\rho_2^2)-\rho_1^2) - \rho_1^2\rho_2^2x_2^2(1+x_2^2)}{(1-\rho_1^2-\rho_2^2)((2+x_2^2)(1-\rho_2^2)-\rho_1^2)}\right\} \\
&\propto \frac{1}{\sqrt{1+x_2^2}} \frac{\sqrt{(1+x_2^2)(1-\rho_1^2-\rho_2^2)}}{\sqrt{(2+x_2^2)(1-\rho_2^2)-\rho_1^2}} \exp\left\{-\frac{1}{2} \frac{x_2^2[(2+x_2^2)-\rho_1^2]}{(2+x_2^2)(1-\rho_2^2)-\rho_1^2}\right\} \\
&\propto \frac{1}{\sqrt{1+x_2^2}} \frac{\sqrt{(1+x_2^2)(1-\rho_1^2-\rho_2^2)}}{\sqrt{(2+x_2^2)(1-\rho_2^2)-\rho_1^2}} \exp\left\{-\frac{1}{2} \frac{x_2^2(2+x_2^2)\rho_2^2}{(2+x_2^2)(1-\rho_2^2)-\rho_1^2}\right\} \exp\left\{-\frac{1}{2}x_2^2\right\}. \tag{B.8}
\end{aligned}$$

We observe that the last expression is symmetric w.r.t. x_2 around the origin. This implies that the conditionally expected value in (B.7) is zero. Therefore, the left-hand side of (B.2) is equal to the conditional variance, and we have from (B.6)

$$\begin{aligned}
\mathbb{E}[X_1^2|Y=0, D=0] &= \text{Var}(X_1|Y=0, D=0) = \mathbb{E}_0\left[\frac{(1+X_2^2)(1-\rho_1^2-\rho_2^2)}{(2+X_2^2)(1-\rho_2^2)-\rho_1^2} \Big| Y=0\right] \\
&= \frac{\mathbb{E}\left[\left(\frac{1}{1+X^2}\right)^{1/2} \left(\frac{(1+X^2)(1-\rho_1^2-\rho_2^2)}{(2+X^2)(1-\rho_2^2)-\rho_1^2}\right)^{3/2} \exp\left\{-\frac{1}{2} \frac{X^2(2+X^2)\rho_2^2}{(2+X^2)(1-\rho_2^2)-\rho_1^2}\right\}\right]}{\mathbb{E}\left[\left(\frac{1}{1+X^2}\right)^{1/2} \left(\frac{(1+X^2)(1-\rho_1^2-\rho_2^2)}{(2+X^2)(1-\rho_2^2)-\rho_1^2}\right)^{1/2} \exp\left\{-\frac{1}{2} \frac{X^2(2+X^2)\rho_2^2}{(2+X^2)(1-\rho_2^2)-\rho_1^2}\right\}\right]}, \tag{B.9}
\end{aligned}$$

with $X \sim \mathcal{N}(0, 1)$ and where we use (B.8) to calculate the corresponding integrals. That is, the denominator of (B.9) corresponds to the missing normalizing constant in (B.8), and it can be obtained by interpreting the corresponding integral as an expectation under a Gaussian distribution due to the last term in (B.8). A similar interpretation applies to the numerator in (B.9). From (B.9) we conclude that we can calculate the left-hand side of (B.2) by Monte Carlo simulation.

The middle term in (B.2) is calculated completely analogously, and it can be received from (B.9) by setting $\rho_1 = \rho_2 = 0$. This results in

$$\mathbb{E}[X_1^2|Y=0] = \text{Var}(X_1|Y=0) = \frac{\mathbb{E}\left[\left(1+X^2\right) \left(\frac{1}{2+X^2}\right)^{3/2}\right]}{\mathbb{E}\left[\left(\frac{1}{2+X^2}\right)^{1/2}\right]},$$

with $X \sim \mathcal{N}(0, 1)$. Evaluating these terms with Monte Carlo simulation using 10M Gaussian samples provides us with

$$\mathbb{E}[X_1^2|Y=0, D=0] = 0.520 < 0.604 = \mathbb{E}[X_1^2|Y=0].$$

This proves (B.2), which completes the proof of Proposition 3.3. \square

Proof of Proposition 3.4. Sufficiency (3.4) of $\mu(\mathbf{X})$ implies that

$$\text{Var}(Y|\mu(\mathbf{X}), \mathbf{D}) = \text{Var}(Y|\mu(\mathbf{X})). \tag{B.10}$$

Therefore, it suffices to give a counterexample to (B.10) under the model assumptions of Example 2.1 to prove Proposition 3.4. We first calculate the right-hand side of (B.10). It is given by

$$\begin{aligned}
\text{Var}(Y|\mu(\mathbf{X})) &= \text{Var}(Y|X_1) \\
&= \text{Var}(\mathbb{E}[Y|\mathbf{X}]|X_1) + \mathbb{E}[\text{Var}(Y|\mathbf{X})|X_1] \\
&= \text{Var}(X_1|X_1) + \mathbb{E}[1+X_2^2|X_1] \\
&= 0 + 1 + \text{Var}(X_2) = 2,
\end{aligned}$$

where we use that X_1 and X_2 are independent, see (2.4). Next we calculate the left-hand side of (B.10). It is given by

$$\begin{aligned}
\text{Var}(Y|\mu(\mathbf{X}), \mathbf{D}) &= \text{Var}(Y|X_1, \mathbf{D}) \\
&= \text{Var}(\mathbb{E}[Y|\mathbf{X}, \mathbf{D}]|X_1, \mathbf{D}) + \mathbb{E}[\text{Var}(Y|\mathbf{X}, \mathbf{D})|X_1, \mathbf{D}] \\
&= \text{Var}(X_1|X_1, \mathbf{D}) + \mathbb{E}[1 + X_2^2|X_1, \mathbf{D}] \\
&= 0 + 1 + \mathbb{E}[X_2^2|X_1, \mathbf{D}].
\end{aligned} \tag{B.11}$$

The distribution of X_2 , given (X_1, D) , can be calculated under the Gaussian assumption (2.4), and it is given by

$$X_{2|(X_1, D)} \sim \mathcal{N}\left(\frac{\rho_2}{1 - \rho_1^2}(D - \rho_1 X_1), \frac{1 - \rho_1^2 - \rho_2^2}{1 - \rho_1^2}\right),$$

this can be obtained, e.g., from Corollary 4.4 of Wüthrich–Merz [20]. This implies that in general the last term in (B.11) is different from 1 because it depends on (X_1, D) . Henceforth, identity (B.10) cannot hold, which completes the proof. \square