



City Research Online

City St George's, University of London

Citation: Lindholm, M., Richman, R., Tsanakas, A. & Wüthrich, M. V. (2026). Sensitivity-based measures of discrimination in insurance pricing. *European Journal of Operational Research*, 333(2), pp. 601-615. doi: 10.1016/j.ejor.2026.01.021

This is the published version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/36642/>

Link to published version: <https://doi.org/10.1016/j.ejor.2026.01.021>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, 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. For full details of reuse please refer to [City Research Online policy](#).



Innovative Applications of O.R.

Sensitivity-based measures of discrimination in insurance pricing

Mathias Lindholm ^{a,*}, Ronald Richman ^b, Andreas Tsanakas ^c, Mario V. Wüthrich ^d^a Department of Mathematics, Stockholm University, SE-106 91, Sweden^b insureAI, Floor 2 30 Melrose Boulevard, 2196, Melrose Arch Gauteng, South Africa^c Bayes Business School, City St George's, University of London, 106 Bunhill Row, EC1Y 8TZ, London, United Kingdom^d RiskLab, Department of Mathematics, ETH Zurich, Switzerland

ARTICLE INFO

Keywords:

Sensitivity analysis
Proxy discrimination
Demographic parity
Insurance pricing
Algorithmic fairness

ABSTRACT

Different notions of fairness and discrimination have been extensively discussed in the machine learning, operations research, and insurance pricing literatures. As not all fairness criteria can be concurrently satisfied, metrics are needed that allow assessing the materiality of discriminatory effects and the trade-offs between various criteria. Methods from sensitivity analysis have been deployed for the measurement of demographic unfairness, that is, the statistical dependence of risk predictions on protected attributes. We produce a sensitivity-based measure for the distinct phenomenon of proxy discrimination, referring to the implicit inference of protected attributes from other covariates. For this, we first define a set of admissible prices that avoid proxy discrimination. Then, the measure is defined as the normalised L^2 -distance of a price from the closest element in that set. We use variance-based sensitivity analysis and Shapley values to attribute the proxy discrimination measure to individual (or subsets of) covariates and investigate how properties of the data generating process are reflected in those metrics. Furthermore, we build on the global (i.e., portfolio-wide) measures of demographic unfairness and proxy discrimination to propose local (i.e., instance- or policyholder-specific) measures, which allow a fine-grained understanding of discriminatory effects. Finally, we apply the methods developed in the paper to a real-world insurance dataset, where ethnicity is a protected variable. We observe substantial proxy-discriminatory effects for one ethnic group and identify the key variables driving this.

1. Introduction

Questions of fairness and discrimination have become central to the machine learning literature (Barocas & Selbst, 2016; Mehrabi et al., 2021) and its applications. Increasingly, the wider operations research literature is building fairness considerations into problems as diverse as organ allocation (Bertsimas et al., 2013), efficiency measurement (Radovanović et al., 2022), hiring (Komiyama & Noda, 2024), and stress testing (Glasserman & Li, 2025). In particular, in applications such as credit scoring (Hurlin et al., 2024; Kozodoi et al., 2022) or insurance pricing (Frees & Huang, 2023; Lindholm et al., 2022), the risk of financial exclusion for particular demographic groups has a high societal salience. Furthermore, questions of fairness are intrinsically linked to questions of model interpretability, since the latter can be a pre-requisite for effectively identifying and addressing the former (Bock et al., 2024; Kraus et al., 2024). As a result, fairness and explainability considerations are increasingly build within the structure of predictive models (Carrizosa et al., 2026, 2025).

A variety of criteria have been formulated, which model predictions should satisfy in order to be considered fair (indicatively, Barocas & Selbst, 2016; Charpentier, 2024; Mehrabi et al., 2021). On the one hand, *group fairness* criteria interrogate the statistical relationship between responses, predictions, and protected attributes such as gender or ethnicity; these criteria are typically formulated via (conditional) independence statements. On the other hand, *individual fairness* criteria focus on whether individuals with similar risk profiles are treated similarly, e.g., are quoted comparable insurance prices. Here different fairness criteria arise from different notions of similarity and information restrictions. Hence, while group fairness criteria consider the outcome of a prediction algorithm, individual fairness notions revolve more around the way that these predictions are generated. Furthermore, as part of the rich literature on fairness and discrimination, an understanding has developed that such criteria are not necessarily consistent with each other and can even be mutually exclusive (Kleinberg et al., 2016; Lindholm et al., 2024b). Consequently, the need emerges to construct metrics for different forms of discrimination and unfairness: one does not just need

* Corresponding author.

E-mail addresses: lindholm@math.su.se (M. Lindholm), ron@insureai.co (R. Richman), A.Tsanakas.1@citystgeorges.ac.uk (A. Tsanakas), mario.wuethrich@math.ethz.ch (M.V. Wüthrich).<https://doi.org/10.1016/j.ejor.2026.01.021>

Received 23 December 2024; Accepted 13 January 2026

Available online 15 January 2026

0377-2217/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

to know whether such phenomena take place as part of a prediction/decision process, but also whether the effects are material enough to justify concern and eventual (e.g., regulatory) action. Furthermore, the potential incompatibility of different fairness notions means that one cannot require for all of them to hold at the same time. This creates the need to monitor many aspects of possible unfairness and measure the respective trade-offs, including those arising from any adjustments to model predictions.

Here we discuss measures for two distinct phenomena, *demographic unfairness* and *proxy discrimination*, with a clear emphasis on the latter. While our application context is insurance pricing, the measures are more broadly applicable to settings such as, e.g., credit scoring. Demographic unfairness relates to violations of demographic parity, that is, the requirement that prices are statistically independent of policyholders' protected attributes. While the deployment of this particular group fairness notion in insurance has been criticised (Lindholm et al., 2024b), we consider it for two reasons: first, because it is easy to explain and politically salient, therefore a potential source of reputational risk for insurers (e.g., Cook et al., 2022); and second because it helps with introducing the construction of the measures we are using. The notion of proxy discrimination builds on the understanding that some policyholder attributes, like gender or ethnicity, should not be used to calculate the price for individual policyholders, as this would constitute *direct discrimination*. Based on that premise, it is additionally desirable to avoid the effective proxying of protected attributes by other variables (e.g., car engine size, postal code) that are correlated to them. Several conceptualisations of proxy discrimination exist (Prince & Schwarcz, 2019; Tschantz, 2022), with the causal structure of covariates often taking centre stage (Côté et al., 2025; Iturria et al., 2024). Here we take a view of proxy discrimination as a form of omitted variable bias, which is not contingent on assumptions of causality, but focuses on the indirect inference of protected attributes from other covariates, in the sense of Lindholm et al. (2024a, 2022, 2024b).

Sensitivity analysis has been long deployed (e.g., Borgonovo & Plischke, 2016; Lamboni & Kucherenko, 2021; Mara & Tarantola, 2012; Pessenti et al., 2025; Saltelli et al., 2010, 2008) to provide insight into complex computational models, evaluate the relative importance of model inputs and identify model vulnerabilities; for applications specifically to insurance risk portfolios and insurance regulation, see respectively Rabbitt and Borgonovo (2020), Vallarino et al. (2024) and Borgonovo et al. (2024). Within that literature, Bénese et al. (2024), utilise contemporary sensitivity analysis methods to develop measures for a wide range of group fairness criteria, that include demographic unfairness. These fairness criteria relate to the joint distribution of model inputs and prices. Hence, they cannot capture proxy discrimination, which relates to the way that prices are derived, rather than their statistical properties as model outputs. To our knowledge, applications of sensitivity analysis to the problem of measuring proxy discrimination are currently lacking in the literature (a brief relevant discussion is found in Hiabu et al., 2023). In our work, we address this challenge, by focusing on the development of metrics for proxy discrimination.

In Section 2 we formally introduce the ideas of demographic unfairness and proxy discrimination, using the paradigm of insurance pricing in the exposition. Specifically, in Sections 2.3 and 2.4 we define, respectively, measures of demographic unfairness and proxy discrimination and discuss their properties. We take as an inspiration and comparison point, a simple measure of demographic unfairness, already discussed by Bénese et al. (2024), based on Sobol' indices (Saltelli et al., 2010, 2008; Sobol', 2001). The measure of proxy discrimination is new and addresses the challenging task of quantifying a phenomenon related to the way prices are calculated from observable quantities. The measure is based on the distance between any given price and the closest element in a set of prices that avoid proxy discrimination. This set of admissible prices arises as a convex combination of the discrimination-free prices in Lindholm et al. (2022) and constant prices that do not depend on any policyholder characteristics. This idea is operationalised through a

constrained regression of the price on best-estimate prices, calculated using different scenarios regarding the value of a protected attribute. The measure of proxy discrimination takes values between 0 and 1 for an insurance portfolio and can be evaluated for any system of prices, without reference to how these were calculated. As a result it lends itself to empirical evaluation and price auditing.

While a global measure of proxy discrimination for a portfolio is useful, it is also necessary to understand which covariates are the sources for such discrimination. In Section 2.5, we discuss how to attribute the measure of proxy discrimination to covariates. Given that the measure is a result of an L^2 -projection, this can be achieved by a simple adaptation of the variance-based Sobol' and total sensitivity indices (Saltelli et al., 2008) and further refined via Shapley values (Owen, 2014).

In Section 3.1, we discuss properties of the underlying data generating process and explain how the sensitivity measures respond to such properties. Subsequently, in Sections 3.2 and 3.3 we introduce, respectively, local measures of demographic unfairness and proxy discrimination. These measures are evaluated for individual policies and allow a more fine-grained understanding and visualisation of the way in which discriminatory effects may arise in a portfolio. The local measures are given by the price of a policy minus a benchmark price that is free of the particular type of unfairness considered. For the case of demographic unfairness, the benchmark price comes from an Output Optimal Transport transformation of the portfolio's prices (Charpentier, 2024; Lindholm et al., 2024b). For the case of proxy discrimination, the benchmark is given by the closest element in the set of prices that avoid proxy discrimination.

In Section 4, we apply the measures developed to a real-world motor insurance dataset, where ethnicity is the protected characteristic. Attribution of the proxy discrimination metric to covariates helps with identifying the key drivers of discriminatory effects. On the computational side, we propose a dependence-respecting calculation of Shapley values combining KernelSHAP (Lundberg & Lee, 2017) with conditional expectation networks (Richman & Wüthrich, 2023). The use of local measures and the consideration of global metrics on sub-portfolios give a more granular view, with specific groups (e.g. younger drivers of one particular ethnicity) seen to be more heavily affected by proxy discrimination. The application shows how the metrics introduced in this paper can be applied in practice and provide insights into the potentially discriminatory effects of algorithmic decisions.

Brief conclusions are stated in Section 5. Some additional technical details are given as appendices in the Supplementary materials.

2. Measures of proxy discrimination and demographic unfairness

2.1. Setup and notation

We work on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with \mathbb{P} being the real-world probability measure. On that space we consider the random vector (Y, X, D) . The random variable Y represents a quantity to be predicted based on covariates (X, D) . Of these, X is a vector of *non-protected covariates* (non-discriminatory characteristics), while D reflects a *protected attribute* (discriminatory or sensitive characteristic). From now on, we will refer specifically to an insurance pricing context, where Y is a loss or a loss frequency, while (X, D) capture policyholder characteristics. The variability of (X, D) under \mathbb{P} represents heterogeneity in the population of policyholders, while the variability of Y conditional on $\{X = x, D = d\}$ reflects loss uncertainty for a policyholder with known features (x, d) . We will assume throughout that D is discrete, taking values in the finite set \mathfrak{D} . Elements of \mathfrak{D} can be thought of as representing different demographic groups, e.g., by ethnicity and/or gender.

Throughout, for a generic random variable Z , we represent its distribution function by $\mathbb{P}(z)$; the conditional distribution of Z given $W = w$ is denoted accordingly by $\mathbb{P}(z|w)$. In the case of absolutely continuous random variables Z , probability density functions are given by $d\mathbb{P}(z)/dz$. In the case of discrete Z we have probability weights $\mathbb{P}(z) = \mathbb{P}(Z = z) > 0$.

2.2. Pricing functions and discriminatory effects

We define the *best-estimate price* as

$$\mu(x, d) := \mathbb{E}[Y \mid X = x, D = d], \tag{1}$$

such that $\mu(x, d)$ is the optimal (in L^2 -norm) prediction of the loss Y for a policyholder with features (x, d) . Best-estimate prices have discriminatory effects because of their direct dependence on protected characteristics D . The most straightforward way of correcting for such *direct discrimination*, is to calculate insurance prices without including the information D as a covariate for prediction of Y . The resulting conditional expectation based only on non-protected characteristics X is termed the *unawareness price* and is defined by

$$\mu(x) := \mathbb{E}[Y \mid X = x]. \tag{2}$$

Nonetheless, the unawareness price may still have discriminatory effects, arising from the potential dependence between the random vectors X and D . Such dependence need not have a causal source, but just be a feature of a particular portfolio structure of insurance policies, e.g., when female and male policyholders have different age distributions.

Out of many different notions of unfairness or discrimination we focus on two complementary perspectives on the impact of the dependence of protected attributes and non-protected covariates. First, given such dependence it will generally hold that $\mu(X)$ is also dependent on D . This means that insurance prices may vary across demographic groups, with the implication that that for some $d \neq d', d, d' \in \mathfrak{D}$, we have that $\mathbb{E}[\mu(X) \mid D = d] \neq \mathbb{E}[\mu(X) \mid D = d']$. For short, we call such insurance prices *demographically unfair*. Second, removing D from the set of covariates does not imply that these are not used *indirectly* in pricing. A concern is that, if D can be partially predicted from X , it is possible that insurance prices, derived with the aim of maximising predictive accuracy, implicitly use X to infer D . In fact, by observing that we can re-write unawareness prices as

$$\mu(x) = \sum_{d \in \mathfrak{D}} \mu(x, d) \mathbb{P}(d \mid x), \tag{3}$$

it becomes clear that unawareness prices do indeed rely on implicit inference of D from X , via the conditional probability $\mathbb{P}(d \mid x)$ used in averaging over best-estimate prices (if there is dependence between D and X). We say that prices utilising such inference of protected characteristics are subject to *proxy discrimination*.

2.3. Measuring demographic unfairness

We now put the previous ideas on a more formal footing, which applies to a general pricing functional $X \mapsto \pi(X)$. We start with the well-known idea of demographic parity.

Definition 1. The pricing functional $X \mapsto \pi(X)$ satisfies *demographic parity* with respect to $\mathbb{P}(X, D)$, if the random variable $\pi(X)$ is independent of D under \mathbb{P} . If π violates demographic parity, we say that it is *demographically unfair*.

Note that independence of X and D is a sufficient but not a necessary condition for demographic parity. Furthermore, establishing that a pricing functional π is demographically unfair does not necessarily mean that the resulting impact is substantial. To understand the materiality of unfairness, one can visualise the changes in (empirical estimates of) the conditional density of $\pi(X) \mid D = d$ across demographic groups $d \in \mathfrak{D}$. Sometimes though, a global numerical metric is useful. We now present such a metric, following Bénése et al. (2024) who apply ideas from sensitivity analysis to evaluate various concepts of algorithmic (un)fairness.

Definition 2. The *demographic unfairness metric* UF is defined as

$$\text{UF}(\pi) = \frac{\text{Var}(\mathbb{E}[\pi(X) \mid D])}{\text{Var}(\pi(X))}, \tag{4}$$

with the convention that if $\text{Var}(\pi(X)) = 0$, then $\text{UF}(\pi) = 0$.

The rationale for the construction (4) is well established in the different context of Global Sensitivity Analysis and the metric is known as a *Sobol' Index* (e.g. Saltelli et al., 2008; Sobol', 2001). Noting that $\text{Var}(\pi(X)) = \text{Var}(\mathbb{E}[\pi(X) \mid D]) + \mathbb{E}[\text{Var}(\pi(X) \mid D)]$, the numerator in (4) represents the amount of variation in $\pi(X)$ attributable to the protected variable D . For a discrete D and given prices $\pi(X)$, the conditional expectation $\mathbb{E}[\pi(X) \mid D]$ can be constructed empirically by simply averaging prices over the subsets of the data defined by $D = d, d \in \mathfrak{D}$.

The following properties are standard (e.g., Saltelli et al., 2008).

Proposition 1. *The unfairness metric UF satisfies the following properties.*

- i) $0 \leq \text{UF}(\pi) \leq 1$. Furthermore, for all $a, b \in \mathbb{R}$ it holds that $\text{UF}(a + b\pi) = \text{UF}(\pi)$.
- ii) If π satisfies demographic parity with respect to $\mathbb{P}(X, D)$, then $\text{UF}(\pi) = 0$.
- iii) If $\sigma(\pi(X)) \subseteq \sigma(D)$, i.e., $\pi(X)$ is D -measurable, then $\text{UF}(\pi) = 1$.

Example 1. Let $D \in \{0, 1\}$, $X \sim U(0, 1)$ and $\mathbb{P}(D = 1 \mid X) = \mathbb{E}[D \mid X] = X$. This implies $\mathbb{P}(D = 1) = \frac{1}{2}$. Assume that the best-estimate price is

$$\mu(X, D) = \frac{1}{2} + X + D,$$

which includes a fixed cost. In this model there is the potential for demographic unfairness, since (X, D) are dependent. The unawareness price equals

$$\mu(X) = \mathbb{E}\left[\frac{1}{2} + X + D \mid X\right] = \frac{1}{2} + 2X.$$

Straightforward calculations lead to

$$\begin{aligned} \mathbb{E}[\mu(X) \mid D = 0] &= \frac{1}{2} + 2\mathbb{E}[X \mid D = 0] = \frac{1}{2} + \frac{2}{3}, \\ \mathbb{E}[\mu(X) \mid D = 1] &= \frac{1}{2} + 2\mathbb{E}[X \mid D = 1] = \frac{1}{2} + \frac{4}{3}. \end{aligned}$$

It then follows that $\text{Var}(\mathbb{E}[X \mid D]) = 1/36$. Hence, if the unawareness price is used, the average premium for $D = 0$ is different compared to $D = 1$. Consequently demographic unfairness arises. We can quantify this effect via the UF metric:

$$\text{UF} = \frac{\text{Var}(\mathbb{E}[1/2 + 2X \mid D])}{\text{Var}(1/2 + 2X)} = \frac{\text{Var}(\mathbb{E}[X \mid D])}{\text{Var}(X)} = \frac{1}{3} > 0.$$

□

We conclude with some notes on the UF measure and potential alternatives. Definition 2 of the unfairness metric, already found in Bénése et al. (2024), is a classic first-order variance-based sensitivity measure in the sense of, e.g., Saltelli et al. (2008). Typically, along such a measure the *total sensitivity* would be stated. In an insurance pricing context, this measure is only meaningful if prices potentially include direct discrimination, i.e. where the prices are also functions of the protected characteristic D . So far we have assumed that this is not the case due to, e.g., legal or regulatory constraints. Nonetheless, the measure (4) can also be evaluated with respect to a broader classes of prices, by substituting for $\pi(X)$ some random variable Π that is not necessarily X -measurable. If direct discrimination is the focal issue, one can follow the suggestion of Bénése et al. (2024) and use the total sensitivity measure $1 - \text{Var}(\mathbb{E}[\Pi \mid X]) / \text{Var}(\Pi)$. Then, if $\Pi \equiv \pi(X)$ is X -measurable – as is the case throughout this paper – that metric becomes equal to zero.

Nonetheless, a suitable total sensitivity measure can be defined if one considers the representation $X = g(D, \mathbf{R})$, where \mathbf{R} is a random vector independent of D . Such constructions are always possible (e.g. via Rosenblatt transforms) and have been widely used in sensitivity analysis, when considering dependent model inputs (Lamboni & Kucherenko, 2021; Mara & Tarantola, 2012; Pesenti et al., 2021, 2025) and, in the context of fairness measures, by Bénése et al. (2024). We do not wish to replicate the discussion in those papers here and proceed succinctly. Given the representation $X = g(D, \mathbf{R})$, $D \perp \mathbf{R}$, one can define:

$$\widetilde{\text{UF}}(\pi) := 1 - \frac{\text{Var}(\mathbb{E}[\pi(X) \mid \mathbf{R}])}{\text{Var}(\pi(X))}.$$

The calculation of $\widetilde{\text{UF}}$ is illustrated in our running example.

Example 2. Returning to [Example 1](#), we can construct $\mathbf{R} = R$ as a (1-dimensional) uniform random variable, defined implicitly by the a.s. relation $X = g(D, R) := F_{X|D}^{-1}(R|D)$. Calculations not documented here lead to the required quantities:

$$X = \sqrt{R} \mathbb{1}_{\{D=1\}} + \left(1 - \sqrt{1-R}\right) \mathbb{1}_{\{D=0\}},$$

$$\mathbb{E}[\mu(X) | R] = \frac{3}{2} + \sqrt{R} - \sqrt{1-R},$$

$$\text{Var}(\mathbb{E}[\mu(X) | R]) = 1 - \frac{\pi}{4}.$$

Consequently the unfairness measure $\widetilde{\text{UF}}$ equals:

$$\widetilde{\text{UF}}(\mu) = 1 - \frac{1 - \pi/4}{1/3} \approx 0.3562,$$

which is very close to the value $\text{UF}(\mu) = 1/3$ obtained in [Example 1](#).

Finally, a common criticism of variance-based sensitivity measures of the form (4) is that a value of zero does not necessarily imply independence of D from $\pi(X)$. Indeed, the scenario $\text{Var}(\mathbb{E}[\pi(X) | D]) = 0$ would imply that $\mathbb{E}[\pi(X) | D] = 0$ a.s., that is, $\mathbb{E}[\pi(X) | D = d] = \mathbb{E}[\pi(X) | D = d']$ for all $d \neq d'$ such that $\mathbb{P}(D = d), \mathbb{P}(D = d') > 0$. This would correspond to a typical monitoring requirement in insurance pricing, assessing whether the average premium differs across demographic groups – hence the situation $\text{UF}(\pi) = 0$ is directly interpretable in our context.

Still, $\text{UF}(\pi) = 0$ is a weaker condition than achieving demographic fairness. This can be addressed by using alternative sensitivity measures, e.g., within the wide family comprehensively studied by [Borgonovo et al. \(2025\)](#), [Borgonovo and Plischke \(2016\)](#), [Plischke and Borgonovo \(2020\)](#), which admit variance-based sensitivity as a special case. In our context, define by G the cdf of $\pi(X)$ and by G_d its conditional cdf given $D = d$. For a suitably defined distance function $\zeta(\cdot, \cdot)$ on pairs of distributions (or measures), one can then build sensitivity measures by first formulating the function $\zeta : d \mapsto \zeta(G, G_d)$ and then evaluating the sensitivity by $\xi = \mathbb{E}[\eta(D)]$. The distance function ζ can be chosen in a way such that the resulting measure satisfies $\xi = 0 \implies D \perp\!\!\!\perp \pi(X)$. For example, one such choice is based on the Kolmogorov-Smirnov statistic:

$$\zeta(G, G_d) = \sup_y |G(y) - G_d(y)|, \quad \xi = \mathbb{E} \left[\sup_y |G(y) - G_D(y)| \right].$$

2.4. Defining and quantifying proxy discrimination

We now turn our attention to the measurement of proxy discrimination. Here the situation is different, compared to unfairness. Recall that the best-estimate prices $\mu(X, D)$ are not used, as they would give rise to direct discrimination and some other pricing functional $\pi(X)$ must be used instead. Effectively this corresponds to merging rating classes with the same non-protected profile $\{X = \mathbf{x}\}$ but different protected characteristics $\{D = d\}$. Hence, one can calculate the price in each of those new classes as a weighted average over d of the corresponding costs $\mu(\mathbf{x}, d)$. Thus, pricing relates to reallocating the claims costs $\mu(\mathbf{x}, d)$, following the removal of the explanatory effect of D . To avoid proxy discrimination – unlike the situation of unawareness prices (3) – the selected weights used should not depend on \mathbf{x} .

To address this issue, [Lindholm et al. \(2022\)](#) suggest the pricing formula

$$h^*(X) = \sum_{d \in \mathfrak{D}} \mu(X, d) \mathbb{P}^*(d), \tag{5}$$

for some distribution $\mathbb{P}^*(d)$ on \mathfrak{D} . Here, we build on [Lindholm et al. \(2022\)](#), to construct an expanded set of admissible pricing functionals that we consider to be free from proxy discrimination. We accept that a constant price $\pi(X) \equiv \pi$ that does not depend on the covariates X cannot proxy-discriminate, as it does not discriminate between policyholders in any sense. Consequently, we would not like to exclude convex combinations of the form $(1 - a)\pi + ah^*(X)$, $a \in [0, 1]$, from the admissible set of prices. Hence, we both allow for an additive constant offset and

also that the weights $v_d := \alpha \mathbb{P}^*(d)$ that are applied on the best-estimate prices $\mu(X, d)$ may sum to less than 1.

Proxy discrimination is now defined with the above arguments in mind. Define the set $\mathcal{V} := \{v \in [0, 1]^{\mathfrak{D}} : \sum_{d \in \mathfrak{D}} v_d \leq 1\}$.

Definition 3. The pricing functional $X \mapsto \pi(X)$ avoids proxy discrimination with respect to $\mu(X, D)$, if for \mathbb{P} -almost every X we can write

$$\pi(X) = c + \sum_{d \in \mathfrak{D}} \mu(X, d)v_d, \tag{6}$$

for some $c \in \mathbb{R}$ and $v \in \mathcal{V}$ that do not depend on X . If π does not have that structure, we say that it is *proxy-discriminatory*.

In general, prices of the form (6), will be dependent with the protected characteristic D and hence will not be demographically fair. Proxy discrimination, via its focus on indirect inference of D focuses on *how* prices are calculated rather than on the statistical *outcome* of the pricing exercise. For a detailed discussion of the difference between the two concepts, see [Lindholm et al. \(2024b\)](#).

Following [Definition 3](#), we can now define a measure of proxy discrimination.

Definition 4. The proxy discrimination metric PD is defined as

$$\text{PD}(\pi) = \frac{\min_{c \in \mathbb{R}, v \in \mathcal{V}} \mathbb{E} \left[\left(\pi(X) - c - \sum_{d \in \mathfrak{D}} \mu(X, d)v_d \right)^2 \right]}{\text{Var}(\pi(X))}, \tag{7}$$

with the convention that if $\text{Var}(\pi(X)) = 0$, then $\text{PD}(\pi) = 0$.

The metric PD quantifies the extent to which a pricing functional cannot be expressed as a weighted average of best-estimate cost terms $\mu(\mathbf{x}, d)$, allowing also for a fixed cost term. Note that the presence of the intercept c , even with the constraints on v , ensures that any solution (c^*, v^*) of the regression problem (7), satisfies $\mathbb{E}[\pi(X) - c^* - \sum_{d \in \mathfrak{D}} \mu(X, d)v_d^*] = 0$. We note that even though c^*, v^* need not be unique, the quantity $c^* + \sum_{d \in \mathfrak{D}} \mu(X, d)v_d^*$ is. Hence, we can explicitly solve for c and express (7) as

$$\text{PD}(\pi) = \frac{\text{Var}(\pi(X) - \sum_{d \in \mathfrak{D}} \mu(X, d)v_d^*)}{\text{Var}(\pi(X))}, \tag{8}$$

with $\text{PD}(\pi)$ reflecting the residual variance for the constrained regression of $\pi(X)$ on $\mu(X, d)$, $d \in \mathfrak{D}$.

The quantity $\pi^*(X) := c^* + \sum_{d \in \mathfrak{D}} \mu(X, d)v_d^*$ is the closest element to $\pi(X)$ in the set of prices that are free from proxy discrimination. We do not specifically suggest $\pi^*(X)$ as a suitable price correction for $\pi(X)$, given that the optimal (c^*, v^*) generally depends on the joint distribution of (X, D) . This would violate the requirements on avoiding proxy discrimination, as formulated by [Lindholm et al. \(2024b\)](#).

The evaluation of (7) requires estimation of the function $\mu(\mathbf{x}, d)$. This is typically carried out on a sample of size n , reflecting experience on the past performance of insurance policies, by regressing claims Y_k on the covariates (X_k, D_k) , $k = 1, \dots, n$. The techniques employed range from Generalised Linear Models to generalisations based on, e.g., Neural Networks; for a wide-ranging review see [Wüthrich and Merz \(2023\)](#). A practical issue can be the availability of information on the protected characteristic D . [Lindholm et al. \(2024a\)](#) show how $\mu(\mathbf{x}, d)$, as well as the classifier $\mathbb{P}(D = d | X = \mathbf{x})$ can be concurrently and accurately estimated by a multi-task neural network, even when information on D is only present for a subset of policies.

Again, simple properties of the metric can be stated.

Proposition 2. The proxy discrimination metric PD satisfies the following properties.

- i) $0 \leq \text{PD}(\pi) \leq 1$. Furthermore, for all $a \in \mathbb{R}, b \in \mathbb{R}_+$ it holds that $\text{PD}(a + b\pi) = \text{PD}(\pi)$.
- ii) $\text{PD}(\pi) = 0$ if and only if π avoids proxy discrimination with respect to $\mu(X, D)$.
- iii) If $\pi(X)$ is uncorrelated with $\mu(X, d)$ for all $d \in \mathfrak{D}$, then $\text{PD}(\mu) = 1$.

Parts i)-ii) of Proposition 2 show that PD is an interpretable metric for proxy discrimination, while part iii) describes a situation where proxy discrimination is maximal: the insurance prices are not at all explained by claim costs and, thus, any discrimination they achieve between policyholders must be undesirable.

In the definition of the proxy discrimination metric (7) we constrained the weights v to be less than one. The reason for this is that higher weights on terms $\mu(X, d)$ may also produce proxying effects, as the next example shows.

Example 3. We continue with the simple model of Example 1. In that model we expect the unawareness price to be subject to proxy discrimination, as (X, D) are dependent and the best-estimate prices are sensitive in D . Let us evaluate the numerator of (7), for $\pi(X) = \mu(X) = \frac{1}{2} + 2X$. We have that

$$\mu(X) - c - \sum_{d \in \{0,1\}} \mu(X, d)v_d = (2 - v_0 - v_1)X - c - \frac{1}{2}(v_0 + 3v_1 - 1).$$

Since c can be chosen to remove the bias for any choices of v_0, v_1 , we have that

$$\mathbb{E} \left[\left(\mu(X) - c^* - \sum_{d \in \{0,1\}} \mu(X, d)v_d^* \right)^2 \right] = (2 - v_0^* - v_1^*)^2 \text{Var}(X).$$

From this it is clear that the minimum is achieved by $v_0^* + v_1^* = 1$. Hence, the proxy discrimination metric (7) becomes

$$\text{PD}(\mu) = \frac{\text{Var}(X)}{\text{Var}(1/2 + 2X)} = \frac{1}{4}.$$

We now consider the alternative price $\pi(X) = 3X$, noting that it agrees on average with the unawareness price, i.e., $\mathbb{E}[\mu(X)] = \mathbb{E}[\pi(X)] = 3/2$. This price penalises further policyholders with X close to 1, which we know are more likely to satisfy $D = 1$. Hence, we expect the price to proxy-discriminate even more than $\mu(X)$; moreover this will be in a gratuitous way, as the increased level of proxy discrimination does not benefit prediction accuracy. Let us now calculate $\text{PD}(\pi)$. Using similar arguments as above, it follows that $\mathbb{E} \left[\left(\pi(X) - c^* - \sum_{d \in \{0,1\}} \mu(X, d)v_d^* \right)^2 \right] = \text{Var}(2X)$. Then,

$$\text{PD}(\pi) = \frac{\text{Var}(2X)}{\text{Var}(3X)} = \frac{4}{9} > \frac{1}{4} = \text{PD}(\mu),$$

such that the increase in the degree of proxy discrimination is reflected in our metric.

Finally, any price that is free of proxy discrimination according to Definition 3 will, for some $c \in \mathbb{R}$, $v \in \mathcal{V}$, take the form

$$\mu^*(X) = c + v_0\mu(X, 0) + v_1\mu(X, 1) = c + \frac{1}{2}(v_0 + 3v_1) + (v_0 + v_1)X.$$

This allows for different choices of prices that avoid proxy discrimination. For example

$$\begin{aligned} \mu_1^*(X) &:= 1 + X \quad (\text{for } v_0 + v_1 = 1) \quad \text{or} \\ \mu_2^*(X) &:= \frac{5}{4} + \frac{1}{2}X \quad (\text{for } v_0 + v_1 = 1/2), \end{aligned}$$

where $\mathbb{E}[\mu_1^*(X)] = \mathbb{E}[\mu_2^*(X)] = \mathbb{E}[\mu(X)]$. For the price $\mu_2^*(X)$ by choosing $v_0 + v_1 < 1$ we have a decreased sensitivity to claim costs and we compensate by a higher flat premium part c . \square

The measure PD of Definition 4 shares an interpretation as a variance ratio with UF, but has a different mathematical structure. One implication is that if $\text{PD}(\pi) = 0$, then we can be assured that the prices $\pi(X)$ are free from proxy discrimination – hence the above criticism of variance-based sensitivity measures no longer applies. In general, the numerator of PD reflects the variability of regression residuals with respect to a quadratic loss function.

We conclude with some comments on the measurement of proxy discrimination. In the construction of the PD measure of (7), we have assumed that the mechanism by which the prices $\pi(X)$ are constructed

is unknown. Definitions 3 and 4 are specifically designed to work for any set of prices $\pi(X)$, potentially arising from a black box model, and we consider this a crucial aspect of operationalising the measurement of proxy discrimination, with reference to a given cost function $\mu(x, d)$. For that reason, inclusion of the constant c is useful, to avoid biased regression coefficients in (7). However, if we restricted analysis to the canonical case of $\pi(X) = \mu(X) = \mathbb{E}[Y | X]$, then we could strengthen the Definitions 3 and 4, by the constraints $c = 0$ and $\sum_d v_d = 1$, which would be consistent with Lindholm et al. (2022).

An alternative lens on proxy discrimination draws on the fact that the protected covariate D , given its dependence on X , can be written as $D = h(X, U)$, where U is a random variable independent of X . This is analogous to the discussion at the end of Section 2.3 and draws on relevant sensitivity literature focusing on dependent inputs (Lamboni & Kucherenko, 2021; Mara & Tarantola, 2012). In the context of proxy discrimination, the function h precisely quantifies the extent to which D is indirectly driven by X ; see Lindholm et al. (2022). However, utilising this insight is not straightforward. As the price functional π is applied to X , it does not admit D as an input. Hence, in general, the extent to which the impact of X on prices operates via the proxying of D is not observable, unless the precise structure of the pricing functional is available. As discussed above, in the practice of price auditing this may not be the case.

Nonetheless, the ideas of dependency models can shed some light when the pricing model is explicitly known and we focus here on the canonical case of the unawareness price $\pi(X) = \mu(X)$. Consider the regression residual in the numerator of (7) and expression (3). Let U be uniform and write $h(X, U) = F_{D|X}^{-1}(U|X)$. Then, for a specific realisation (policyholder) x :

$$\begin{aligned} \mu(x) - c - \sum_{d \in \mathcal{D}} \mu(x, d)v_d &= \sum_{d \in \mathcal{D}} \mu(x, d)\mathbb{P}(d | x) - c - \sum_{d \in \mathcal{D}} \mu(x, d)v_d \\ &= \mathbb{E}[\mu(x, h(x, U))] - c - \sum_{d \in \mathcal{D}} \mu(x, d)v_d. \end{aligned}$$

The term $\mathbb{E}[\mu(x, h(x, U))]$ in the above expression varies in x in two ways: via the *direct effect*, $x \mapsto \mu(x, \cdot)$, and the *indirect effect* through D , $x \mapsto \mu(\cdot, h(x, U))$ – this latter source of variation generates proxy discrimination. When calculating the PD measure (7), we are approximating the prices $\mu(X)$ with the term $c + \sum_{d \in \mathcal{D}} \mu(X, d)v_d$, which is free from that indirect effect and, therefore, from proxy discrimination. Hence, the optimised residual in the numerator of the PD measure captures the ‘unexplained’ variability due to the indirect effect on predictions of X through D .

2.5. Attribution of proxy discrimination to individual covariates

Given the measurement of proxy discrimination by (7), a next question of interest is which (subsets of) covariates – elements of X – are mostly responsible. Here we draw again from literature on Global Sensitivity Analysis (e.g. Saltelli et al., 2008). Let the dimension of X be q and $S \subseteq \{1, \dots, q\} =: \mathcal{Q}$ a set of indices, such that X_S is the corresponding sub-vector of X . Analogously, denote $S^c = \mathcal{Q} \setminus S$, and X_{S^c} , such that $X = (X_S, X_{S^c})$.

Noting the form (8), we can attribute the variance in the numerator to the subset S of covariates, by conditioning on sub-vectors. Following Definition 4, denote the regression residual by

$$\Lambda(\pi, X) := \pi(X) - c^* - \sum_{d \in \mathcal{D}} \mu(X, d)v_d^*. \tag{9}$$

We now define two metrics that reflect the contribution of (a subset of) covariates to proxy discrimination.

Definition 5. For the proxy discrimination metric PD of (7) and $\Lambda(\pi, X)$ as in (9), we define the contribution of the sub-vector X_S , $S \subseteq \mathcal{Q}$ to proxy discrimination by the two metrics,

$$\text{PD}_S(\pi) = \frac{\text{Var}(\mathbb{E}[\Lambda(\pi, X) | X_S])}{\text{Var}(\pi(X))}, \tag{10}$$

$$\widetilde{\text{PD}}_S(\pi) = \frac{\text{Var}(\Lambda(\pi, \mathbf{X})) - \text{Var}(\mathbb{E}[\Lambda(\pi, \mathbf{X}) \mid \mathbf{X}_{S^c}])}{\text{Var}(\pi(\mathbf{X}))}. \tag{11}$$

When $S = \{i\}$, we write $\text{PD}_i(\pi)$, $\widetilde{\text{PD}}_i(\pi)$.

The metric PD_S is thus understood as the sensitivity of the residual $\Lambda(\pi, \mathbf{X})$ to the subset of covariates \mathbf{X}_S , reflecting the amount of variability in $\Lambda(\pi, \mathbf{X})$ driven by \mathbf{X}_S . The metric $\widetilde{\text{PD}}_S$ reflects the expected reduction in the variance of $\Lambda(\pi, \mathbf{X})$ achieved by averaging out \mathbf{X}_{S^c} . When $S = \{i\}$, then PD_i is identified by a (rescaled) Sobol’ Index (or first-order sensitivity), while $\widetilde{\text{PD}}_i$ is known as a *total sensitivity* (Saltelli et al., 2008). A difference to standard sensitivity measures is that here we are normalising with $\text{Var}(\pi(\mathbf{X}))$ – rather than $\text{Var}(\Lambda(\pi, \mathbf{X}))$ – to maintain the direct connection with the global PD metric (7) (note that this suppresses the scale of PD_S , $\widetilde{\text{PD}}_S$).

We note that the choice of Sobol’ indices in Definition 5 is not necessary; indeed any importance measure could be applied to the model output $\Lambda(\mathbf{X}, \pi)$ from the ξ_i measure (Plischke & Borgonovo, 2020) discussed at the end of Section 2.3 to, e.g., the Accumulated Local Effects studied in the machine learning literature (Apley & Zhu, 2020). However, we pursue the route of Definition 5, because, as the PD measure is expressed as a variance ratio, this makes Sobol’ indices a natural tool for studying its drivers.

Estimation of the metrics (10) requires evaluating conditional expectations with respect to subsets of covariates. The sensitivity analysis literature presents various methods to do this (e.g., Jansen, 1999; Saltelli et al., 2010; Sobol’, 2001), though generally under the assumption of independent \mathbf{X} which is not tenable in an insurance context. Alternative approaches are based on predictive modelling (Veiga et al., 2009); for consistent and efficient evaluation of non-linear regressions on all \mathbf{X}_S , $S \subseteq \mathcal{Q}$, through a single model see Richman and Wüthrich (2023).

The metrics introduced in Definition 5 have the following properties, which are standard in the literature on variance-based sensitivity (e.g. Saltelli et al., 2008).

Proposition 3. *The metrics PD_S , and $\widetilde{\text{PD}}_S$ satisfy the following properties.*

- i) $0 \leq \text{PD}_S(\pi)$, $\widetilde{\text{PD}}_S(\pi) \leq \text{PD}(\pi)$.
- ii) a) If $\mathbf{X}_S \perp\!\!\!\perp \Lambda(\pi, \mathbf{X})$, then $\text{PD}_S(\pi) = 0$.
 b) If $\Lambda(\pi, \mathbf{X})$ is \mathbf{X}_S -measurable, then $\text{PD}_T(\pi) = \text{PD}(\pi)$ for all $S \subseteq T \subseteq \mathcal{Q}$.
- iii) a) If $\Lambda(\pi, \mathbf{X})$ is \mathbf{X}_{S^c} -measurable, then $\widetilde{\text{PD}}_T(\pi) = 0$, for all $T \subseteq S$.
 b) If $\mathbf{X}_{S^c} \perp\!\!\!\perp \Lambda(\pi, \mathbf{X})$, then $\widetilde{\text{PD}}_S(\pi) = \text{PD}(\pi)$.
- iv) If there are functions g, h such that $\Lambda(\pi, \mathbf{X}) = g(\mathbf{X}_S) + h(\mathbf{X}_{S^c})$ and $\mathbf{X}_S \perp\!\!\!\perp \mathbf{X}_{S^c}$, then $\text{PD}_S(\pi) + \text{PD}_{S^c}(\pi) = \text{PD}(\pi)$ and $\widetilde{\text{PD}}_S(\pi) + \widetilde{\text{PD}}_{S^c}(\pi) = \text{PD}(\pi)$.

Part i) of Proposition 3 is a natural condition for stating that the metrics $\text{PD}_S(\pi)$, $\widetilde{\text{PD}}_S(\pi)$ reflect contributions to the overall proxy discrimination $\text{PD}(\pi)$. Parts ii) and iii) give conditions for the metrics taking their extremal values. The conditions are complementary, using different independence or measurability assumptions to reflect irrelevance or full relevance of \mathbf{X}_S . Finally, part iv) gives strong conditions for being able to additively decompose the proxy discrimination metric $\text{PD}(\pi)$. We discuss further properties of PD_S , $\widetilde{\text{PD}}_S(\pi)$ in Section 3.1.

The question of additivity is important for interpreting the contributions of covariates to proxy discrimination – thus, lack of additivity is a drawback of the metrics PD_S and $\widetilde{\text{PD}}_S$. Furthermore, it is widely understood that the canonical application domain of Sobol’ indices is that of independent model inputs. One can address both these issues by employing the game-theoretical concept of the *Shapley value* (Shapley, 1953) for the value functional $S \mapsto \text{Var}(\mathbb{E}[\Lambda(\pi, \mathbf{X}) \mid \mathbf{X}_S])$. While recent literature on model interpretability has focused on the use of Shapley values to derive local model explanations for given instances $\mathbf{X} = \mathbf{x}$ (Aas et al., 2021; Lundberg & Lee, 2017), we use the Shapley value for a decomposition of a global sensitivity measure, following Owen (2014), Owen and Prieur (2017), Song et al. (2016).

Definition 6. For the proxy discrimination metric PD of (7) and $\Lambda(\pi, \mathbf{X})$ as in (9), denote

$$w(S) = \text{Var}(\mathbb{E}[\Lambda(\pi, \mathbf{X}) \mid \mathbf{X}_S]), \quad S \subseteq \mathcal{Q}.$$

We define the Shapley attribution of the covariate X_i to proxy discrimination as the metric,

$$\text{PD}_i^{\text{sh}}(\pi) = \frac{1}{\text{Var}(\pi(\mathbf{X}))} \frac{1}{q} \sum_{S \subseteq \mathcal{Q} \setminus \{i\}} \binom{q-1}{|S|}^{-1} (w(S \cup \{i\}) - w(S)). \tag{12}$$

As the Shapley value is a well-known concept across literatures, we do not review its properties here. Its key features are that, by its construction (12), it reflects variable interactions, while achieving an additive attribution, $\sum_{i=1}^q \text{PD}_i^{\text{sh}}(\pi) = \text{PD}(\pi)$. While Definition 6 calculates the Shapley value with respect to $w(S) = \text{Var}(\mathbb{E}[\Lambda(\pi, \mathbf{X}) \mid \mathbf{X}_S])$ and is thus based on PD_S , a result identical to (12) is obtained if instead we use as value function the alternative metric $\widetilde{\text{PD}}_S$ (Song et al., 2016). Given its combinatoric construction, calculation of the Shapley value in models with many inputs is computationally intensive – we return to this issue in case study of Section 4.

3. Structural properties, price adjustments, and local measures

3.1. Structural properties

The ideas of demographic unfairness and proxy discrimination discussed in Section 2 related to the statistical properties and construction of pricing functionals π . Here we associate such properties with structural properties of the data generating process, that is, with features of the joint distribution $\mathbb{P}(Y, \mathbf{X}, D)$. First, we define a number of such properties.

Definition 7 (Structural properties).

- P1 $\mathbf{X} \perp\!\!\!\perp D$ (independence)
- P2 $Y \perp\!\!\!\perp D \mid \mathbf{X}$ (\mathbf{X} -sufficiency)
- P3 $\mu(\mathbf{X}, D) = \mu(\mathbf{X})$ (weak \mathbf{X} -sufficiency)
- P4 $\sigma(\mathbf{X}) \subseteq \sigma(D)$ (\mathbf{X} -irrelevance)
- P5 $Y \perp\!\!\!\perp \mathbf{X} \mid D$ (D -sufficiency)
- P6 $\mu(\mathbf{X}, D) = \mu(D)$ (weak D -sufficiency)

The formulation of properties P1-P6 is not reliant on any assumed causal relation between Y , \mathbf{X} and D . Nonetheless, one can formulate Directed Acyclical Graphs (DAG) representing causal relations such that, e.g. properties P2 or P5 are satisfied – for causal perspectives on proxy discrimination see Iturria et al. (2024), Côté et al. (2025), Tschantz (2022).

In the light of Definitions 1 – 4, the relationships summarised in Proposition 4, below, hold. These show the implications of structural properties P1-P6 for the metrics UF and PD, for either a general price $\pi(\mathbf{X})$ or, more specifically, the unawareness price $\mu(\mathbf{X})$. As the properties P1-P6 are formulated with respect to the data generating process rather than arbitrary pricing functionals, they are strong and only give rise to sufficient conditions.

Proposition 4.

- i) If P1 holds, any pricing functional $\pi(\mathbf{X})$ satisfies demographic parity and $\text{UF}(\pi) = 0$.
- ii) If any of P1, P2 or P3 hold, the unawareness price $\mu(\mathbf{X})$ avoids proxy discrimination and $\text{PD}(\mu) = 0$.
- iii) If P4 holds, any pricing functional $\pi(\mathbf{X})$ is demographically unfair and $\text{UF}(\pi) = 1$, except if $\pi(\mathbf{X}) \equiv \pi$, a constant.
- iv) If any of P4, P5 or P6 hold, any pricing functional $\pi(\mathbf{X})$ is proxy discriminatory and $\text{PD}(\pi) = 1$, except if $\pi(\mathbf{X}) \equiv \pi$, a constant.

Proposition 4, parts i)-ii), give conditions for avoiding demographic unfairness or proxy discrimination; in ii) limiting to the case of unawareness prices. Demographic fairness relates to the joint law of the pricing

functional π and the response Y , hence the strong requirement P1 arises as a natural sufficient condition. Property P2 means that Y depends on D only via X . Hence measurement of non-protected characteristics X eliminates any benefit to predictions from collecting protected characteristics D . This in turn implies P3, which means that the best-estimate prices are insensitive in D . Hence, if we restrict to unawareness prices, these conditions guarantee the absence of proxy discrimination. Conversely, parts iii)-iv) of Proposition 4 give conditions for maximal levels of demographic unfairness and proxy discrimination. Here, properties P4-P6, in different ways, mean that knowing X in addition to D adds no new information useful for predicting the claim Y .

We now turn our attention to the way that structural properties of the data generating process impact on the sensitivity of the PD measure to covariate sub-vectors X_S . In Proposition 5 below, we deal with the common case of unawareness prices, providing dependence scenarios, under which the contributions of sub-vectors of X are either zero or equal to the total level of proxy discrimination $PD(\mu)$. These scenarios hence represent cases of full relevance of X_S and irrelevance of X_{S^c} with respect to the alternative metrics PD_S and \widetilde{PD}_S .

Proposition 5.

- i) Assume that $X_{S^c} \perp\!\!\!\perp (Y, D) \mid X_S$. Then, for the unawareness price $\pi(X) = \mu(X)$ the following hold.
 - a) $PD_S(\mu) = PD(\mu)$ and $\widetilde{PD}_{S^c}(\mu) = 0$.
 - b) If additionally $X_{S^c} \perp\!\!\!\perp X_S$, then it also holds that $\widetilde{PD}_S(\mu) = PD(\mu)$ and $PD_{S^c}(\mu) = 0$.
- ii) Let the best-estimate price take the form $\mu(X, D) = g(X) + h(D)$ and assume that $Cov(g(X), h(D)) \geq 0$ and $X_{S^c} \perp\!\!\!\perp D \mid X_S$. Then, for the unawareness price $\pi(X) = \mu(X)$ the following hold.
 - a) $PD_S(\mu) = PD(\mu)$ and $\widetilde{PD}_{S^c}(\mu) = 0$.
 - b) If additionally $X_{S^c} \perp\!\!\!\perp X_S$, then it also holds that $\widetilde{PD}_S(\mu) = PD(\mu)$ and $PD_{S^c}(\mu) = 0$.

3.2. Local measurement of demographic unfairness

We introduce local (i.e., policyholder-specific) measures, based on the difference of the price-in-use and the a benchmark fair price. For demographically fair prices, a standard construction has been via optimal transport (OT) methods (Chiappa et al., 2020; Gordaliza et al., 2019). In an insurance context, such methods aim at achieving independence between prices $\pi(X)$ and the protected characteristics D by suitable variable transformations; for insurance-specific investigations see Charpentier (2024), Lindholm et al. (2024b). The broader mathematical problem of approximating a random variable with another, subject to an independence constraint, is treated by Delbaen and Majumdar (2024).

Here, to create a demographically fair benchmark price, we follow an Output OT approach. This refers to a transformation (post-processing) of prices $\pi(X)$ in order to achieve independence from D . Denote the conditional distributions of the prices by $G_d(m) = \mathbb{P}(\pi(X) \leq m \mid D = d)$, $d \in \mathfrak{D}$, and assume for simplicity that they are continuous. Then, for any continuous distribution G , we may construct the prices

$$\tilde{\pi}(X, D) = \sum_{d \in \mathfrak{D}} \mathbb{1}_{\{D=d\}} G^{-1} \circ G_d(\pi(X)). \tag{13}$$

The price $\tilde{\pi}(X, D)$ satisfies $\mathbb{P}(\tilde{\pi}(X, D) \leq m) = G(m)$ and $\mathbb{P}(\tilde{\pi}(X, D) \leq m \mid D = d) = G(m)$, $d \in \mathfrak{D}$, such that $D \perp\!\!\!\perp \tilde{\pi}(X, D) \sim G$; the construction (13) works by making the conditional distribution of prices the same on each demographic subgroup $D = d$. Note that the transformed price $\tilde{\pi}(X, D)$ explicitly depends on D – even as $\pi(X)$ does not. This is a form of direct discrimination arising in the process of engineering demographic parity; see Lindholm et al. (2024b) for more discussion of this point. Finally, to construct a demographically fair benchmark, we need to select the target distribution G . A standard choice is given by Chzhen et al. (2020)

$$G^{-1}(u) = \sum_{d' \in \mathfrak{D}} \mathbb{P}(D = d') G_{d'}^{-1}(u). \tag{14}$$

From now on we will consistently refer to the Output OT price as the construction $\tilde{\pi}(X, D)$ from (13) and (14).

We can now proceed with the definition of a local measure of demographic unfairness.

Definition 8. Consider a pricing functional π and the Output OT price $\tilde{\pi}(X, D)$. Then, for the policyholder with profile $X = x$, $D = d$, the local measure of demographic unfairness is defined as:

$$\delta_{UF}(x, d; \pi) = \pi(x) - \tilde{\pi}(x, d). \tag{15}$$

If $\pi(X) = \mu(X)$ is the unawareness price, we just write $\delta_{UF}(x, d)$.

A value of $\delta_{UF}(x, d; \pi) > 0$ implies that policyholders with attributes $X = x$, $D = d$ suffer from demographic unfairness in the sense that the price they are charged is higher than the corresponding benchmark demographically fair price. Clearly, if X and D are already independent, the prices $\pi(X)$ and $\tilde{\pi}(X, D)$ coincide such that the measure becomes zero. This is stated formally below.

Proposition 6. If P1 in Definition 7 holds, then $\delta_{UF}(x, d) = 0$.

Furthermore, it is of interest to establish conditions for the sign of the metric $\delta_{UF}(x, d; \pi)$. In the simple but common case of a binary D , this is straightforward. Denote by \leq_{st} precedence in the usual stochastic order, such that for two distributions F, G , we have that $F \leq_{st} G \iff F(x) \geq G(x)$ for all x ; this is a strong condition not allowing any crossing the distributions.

Proposition 7. Let $\mathfrak{D} = \{0, 1\}$ and $G_0 \leq_{st} G_1$.

- i) $\delta_{UF}(x, 0; \pi) \leq 0$ and $\delta_{UF}(x, 1; \pi) \geq 0$.
- ii) For x such that $\mathbb{P}(D = 0 \mid X = x) > 0$, it holds that:

$$\mathbb{E}[\delta_{UF}(X, D; \pi) \mid X = x] > 0 \iff \frac{\mathbb{P}(D = 1 \mid X = x)}{\mathbb{P}(D = 0 \mid X = x)} > \frac{-\delta_{UF}(x, 0; \pi)}{\delta_{UF}(x, 1; \pi)} = \frac{G^{-1} \circ G_0(\pi(x)) - \pi(x)}{\pi(x) - G^{-1} \circ G_1(\pi(x))} \geq 0,$$

where G is given by (14).

To interpret part i) of Proposition 7, first note that $G_0 \leq_{st} G_1$ means that the policyholders with protected attribute $D = 1$ tend to be considered as higher risk, according to best-estimate prices. So these are the policyholders for whom the use of the Output OT price (13) should confer a discount, compared with the unawareness price. For part ii), we consider a situation where for a policyholder with $X = x$ the value of D may not be known. The left-hand side of the stated condition implies that, on average, the local measure of unfairness will be positive when $\mathbb{P}(D = 1 \mid X = x) / \mathbb{P}(D = 0 \mid X = x)$ is high, such that there is a high chance that, given $X = x$, the policyholder belongs to the demographically disadvantaged group $D = 1$. Furthermore, the condition is more likely to be satisfied when the ratio $-\delta_{UF}(x, 0; \pi) / \delta_{UF}(x, 1; \pi)$ is low. That fraction becomes small if the comparative disadvantage for group $D = 1$ (denominator) becomes much higher than the comparative advantage of group $D = 0$ (numerator), given the information $X = x$. These ideas are illustrated in the following example.

Example 4. We continue from Example 3, where $\mu(X, D) = \frac{1}{2} + X + D$ and $X \sim U(0, 1)$. Note that $\mu(x, 1) > \mu(x, 0)$ for all x . However, we now allow a variety of positive and negative dependence relations between (X, D) by assuming that

$$\mathbb{P}(D = 1 \mid X = x) = \frac{1-a}{2} + ax, \quad a \in (-1, 1].$$

By setting $a = 1$, we recover the exact setting of Example 3; $0 < a < 1$ gives a weaker positive dependence, while $-1 < a < 0$ gives a negative dependence. With this modification the unawareness price changes to $\mu(X) = \frac{2-a}{2} + (a+1)X$.

We now evaluate the local measure of unfairness (15) for the unawareness price $\mu(X)$ and various levels of a . The calculations are simple but tedious and are not reported here. In Fig. 1 we plot the functions $\delta_{UF}(x, d)$ (blue for $d = 0$, red for $d = 1$), as well as their conditional

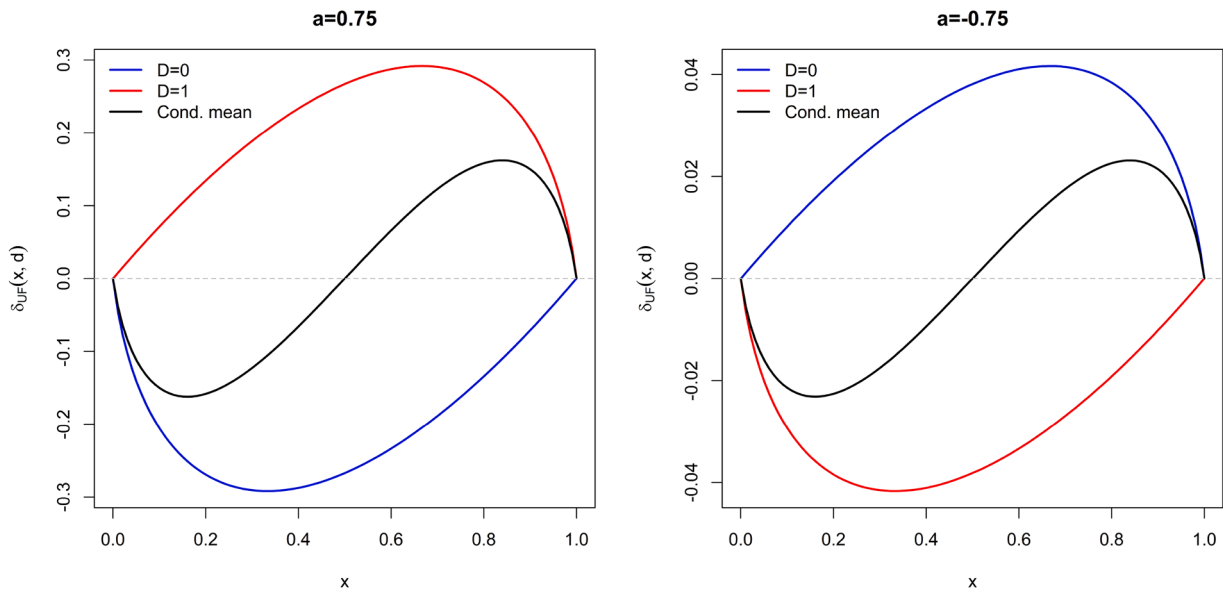


Fig. 1. Local unfairness metric $\delta_{UF}(x, d)$ (blue for $d = 0$, red for $d = 1$) and conditional mean $\mathbb{E}[\delta_{UF}(X, D) | X = x]$ (black) for $a = 0.75$ (left) and $a = -0.75$ (right). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

mean $\mathbb{E}[\delta_{UF}(X, D) | X = x]$ (black) for $a = 0.75$ (positive dependence) and $a = -0.75$ (negative dependence). While the shapes appear similar, there are two observations. First, for $a = 0.75$, the red line is above zero and the blue line below, showing that policyholders with $D = 1$ are adversely affected by demographic unfairness, while policyholders with $D = 0$ are benefiting. This pattern is reversed when the dependence of (X, D) becomes negative ($a = -0.75$). Second, the two plots are at very different scales, with the absolute value of local demographic unfairness being an order of magnitude higher in the case of positive dependence. The reason is that, for $a = 0.75$, the positive dependence of (X, D) works in the same direction as the impact of each of those two variables on claims costs. However, when dependence is negative, then the unawareness price becomes less sensitive to X (in the extreme $a \rightarrow -1$ leads to a constant $\mu(X)$). As a result, for negative dependence much smaller disparities between demographic groups emerge. This point is reinforced by Fig. 2, where we plot $\mathbb{E}[\delta_{UF}(X, D) | X = x]$ against different values of the dependence parameter a .

□

3.3. Local measurement of proxy discrimination

Analogously to the last section, we propose a local measure of proxy discrimination. Again, we need for that purpose a benchmark price that avoids proxy discrimination. Any such price takes the form (6); nonetheless, to produce a benchmark one needs to choose the values of $c, v_d, d \in \mathfrak{D}$. Here we choose values that minimise the numerator in (7), such that the benchmark is the price closest to the original price $\pi(X)$, which is free from proxy discrimination. The difference between a price and its closest approximation that is free from proxy discrimination has already been defined in Section 2.5, as the regression residual $\Lambda(X, \pi)$. Hence we re-purpose this quantity as a local measure of proxy discrimination.

Definition 9. Consider a pricing functional π and $\Lambda(X, \pi)$ in Eq. (9). Then, for the policyholder with profile $X = x$, the local measure of proxy discrimination is defined as:

$$\delta_{PD}(x; \pi) = \Lambda(x, \pi). \tag{16}$$

If $\pi(X) = \mu(X)$ is the unawareness price, we just write $\delta_{PD}(x)$.

We state Proposition 8 below without proof – for each property it is easy to show that the unawareness price is already free from proxy discrimination, such that $\Lambda(X, \pi)$ is identically zero.

Proposition 8. If any one of P1, P2 or P3 in Definition 7 holds, then $\delta_{PD}(x) = 0$.

We conclude this section with a continuation of our running example.

Example 5. We continue from Example 4, considering the evaluation of our local measure of proxy discrimination. We make a qualitative argument, omitting a formal proof. Recall the forms of the best-estimate and unawareness prices, $\mu(X, D) = 1/2 + X + D$ and $\mu(X) = 1 - a/2 + (1 + a)X$ respectively. If $-1 < a \leq 0$, the slope of the unawareness price in X is less or equal to that of the best-estimate price. As a result $\Lambda(X, \mu) = 0$ and there is no proxy discrimination. However, when $0 < a \leq 1$, proxy discrimination arises. The closest approximation to $\mu(X)$ that avoids proxy discrimination is the price $\mu^*(X) = 1 + X$, which reduces the slope in X to 1. Consequently we have

$$\delta_{PD}(x) = \begin{cases} 0, & -1 \leq a \leq 0, \\ -\frac{a}{2} + ax, & 0 < a \leq 1, \end{cases}$$

showing how the metric increases linearly with x , at a slope determined by the dependence parameter a . Because of the positive dependence, policyholders with $x > 0.5$ are implicitly inferred to have protected attribute $D = 1$ and hence are disadvantaged in the sense of proxy discrimination; the reverse happens for $x < 0.5$.

4. Case study: Insurance prices and ethnicity in a real-world motor portfolio

We now apply the metrics developed in this paper to a real-world non-life motor insurance dataset, previously discussed in Lindholm et al. (2024a). The dataset consists of policyholder and vehicle covariates, as well as exposure and claims data for a large portfolio, observed over a single year of exposure. Moreover, the data include protected information $D \in \mathfrak{D} = \{1, 2, 3, 4, 5\}$, relating to the ethnicity code of the policyholder, as defined in the jurisdiction in which the policies were written. The anonymous insurer contributing the data – a large multi-national insurance company – records this information at the time of underwriting the policy. The ethnicity data are not used by the insurer for pricing, but

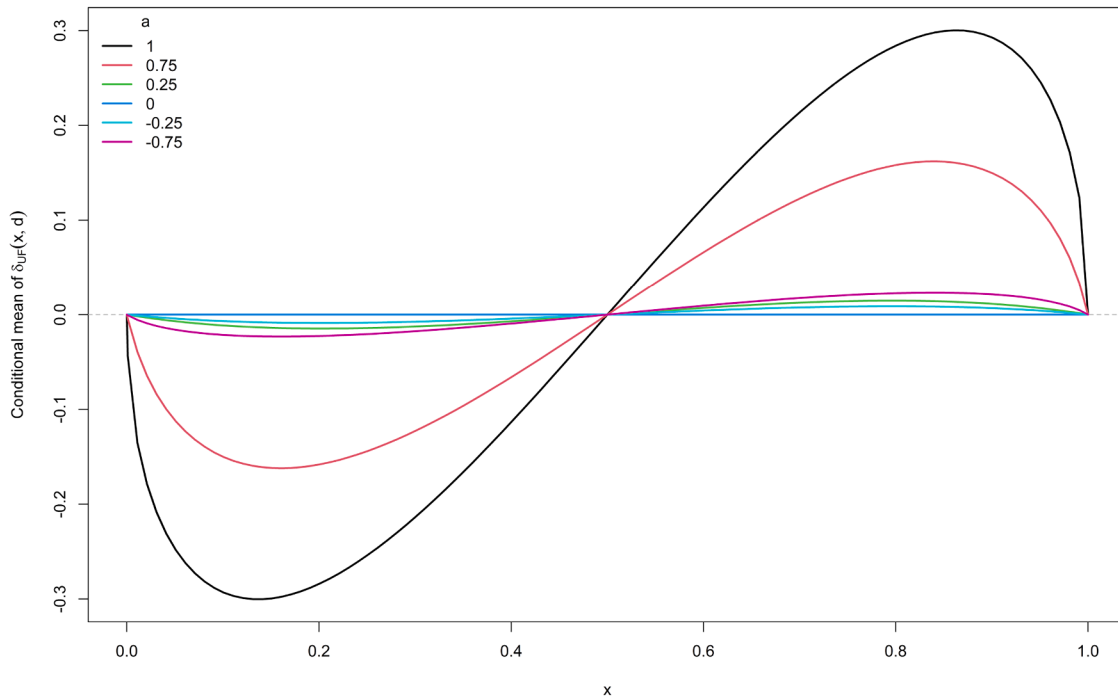


Fig. 2. Conditional mean of the local unfairness metric $\mathbb{E}[\delta_{UF}(X, D) | X = x]$ for the unawareness price and a range of dependence parameters a .

Table 1
Summary of real-world non-life insurance claims, exposures and frequencies of claims according to the ethnicity codes $D \in \mathfrak{D}$.

Ethnicity code	Number of claims	Exposure	Frequency
1	5, 223	14, 317	36.48%
2	965	3, 925	24.59%
3	3, 354	14, 363	23.35%
4	5, 249	20, 240	25.93%
5	26, 817	112, 667	23.80%

Table 2
Demographic unfairness metric UF and proxy discrimination metric PD, for the motor dataset and for the three pricing rules considered. Bootstrap standard errors are in parentheses.

Price	UF	PD
Best-estimate	0.0892 (0.00095)	0.00924 (0.0000443)
Unawareness	0.0639 (0.00079)	0.00277 (0.0000203)
Discrimination-free	0.0533 (0.00070)	0

are collected to enable the insurer to monitor and report insurance penetration. To preserve confidentiality of the contributing company and to ensure that no commercially sensitive information is provided, the ethnicity categories relating to the levels of D have not been used in our analysis; instead we use an integer code randomly assigned to each category.

The data relate to a period close to the turn of the century and contain 165,511 years of exposure of comprehensive motor insurance policies, with 41,608 claims arising in the same period of exposure. The claims can be assumed to be fully run-off, i.e., no incurred but not reported adjustments need to be made, and they relate to property (motor-hull), liability (third-party property and/or bodily injury) and other associated coverages; we do not disclose the exact coverages to avoid disclosing potentially commercially sensitive information. Moreover, information on the excesses and deductibles (which influence the frequency) in this portfolio are also not disclosed, thus, the information shown is not commercially useful.

Table 1 shows a summary of the claims, exposures and frequencies for each ethnicity code $D \in \mathfrak{D}$ in the data. The portfolio exhibits substantial heterogeneity in the size of demographic groups. The largest group corresponds to $D = 5$, while other groups range from $D = 2$ with the smallest population to $D = 4$ with approximately one fifth of the size of the largest group. Analysis of claims experience reveals broadly similar frequencies across groups, with the notable exception of policyholders with $D = 1$ who display markedly higher claim frequencies.

The vector of non-protected characteristics $\mathbf{X} = (X_1, \dots, X_q)^\top$ comprises standard rating factors used in non-life insurance pricing, including policyholder and driver demographics (age and gender), vehicle characteristics, and geographical information (note that, within the context of this application, we do not treat gender as a protected characteristic). For each policy i , where $1 \leq i \leq n$, we observe the tuple $(Y_i, \mathbf{X}_i, D_i, v_i)$, where v_i represents the exposure period (the fraction of the year for which the policy was active) and serves as a weight in the calculation of claim frequencies.

Throughout this application we compare three types of price: (a) best-estimate prices (1); (b) unawareness prices (2); and (c) discrimination-free prices (6) under the canonical choices $c = 0$, $v_d = \mathbb{P}(D = d)$ as in Lindholm et al. (2022) – note that these choices are exogenous and not the product of any optimisation. These data were modelled by Lindholm et al. (2024a), using a plain-vanilla Feed-forward Neural Network (FNN) with embedding layers for the categorical covariates and standard training procedures to regularize the network. An overview of the network structure and fitting process is given in the Supplementary materials; for further details we refer the interested reader to Lindholm et al. (2024a) and its appendices. Here, using the predictions from the fitted models, we estimate the unfairness metric UF and the proxy discrimination metric PD, using empirical versions of Eqs. (4) and (7).

Table 2 shows the estimated values of the two measures, for the best-estimate, unawareness and discrimination-free prices, with bootstrap standard deviations in parentheses. We observe that the best-estimate price exhibits the highest level of the unfairness metric UF, with $UF = 0.0892$, indicating that approximately 9% of price variability

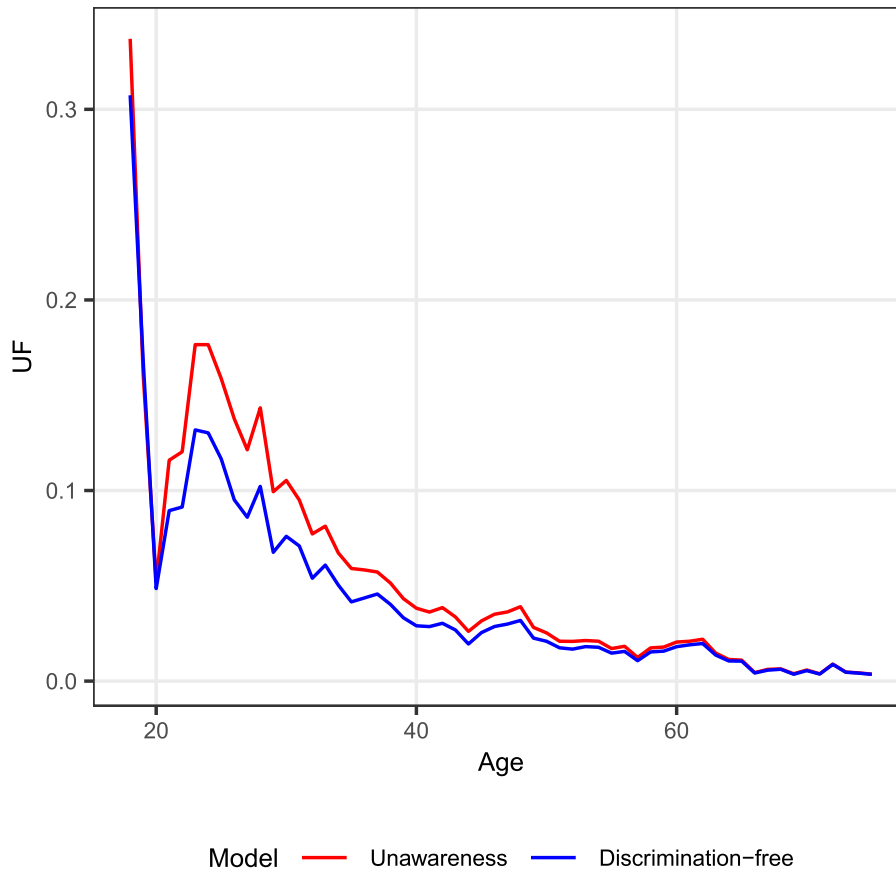


Fig. 3. Unfairness metric $UF_{|Age}(a)$ evaluated for the unawareness prices, evaluated on sub-portfolios split by age (Eq. (17)).

can be attributed to systematic differences between ethnic groups. The unawareness price yields an intermediate value of $UF = 0.0639$, while the lowest level of the unfairness measure, $UF = 0.0533$, comes from the discrimination-free prices. This ordering is not surprising. First, the unawareness price leads to a reduction in demographic unfairness, since it eliminates the direct impact of D on prices. Nonetheless, demographic unfairness persists, because of the statistical dependence of $\mu(X)$ and D . Second, beyond removing the direct impact of D on predictions, discrimination-free prices aim to also remove indirect proxying impacts. Nonetheless, it is not obvious that the latter effect holds more generally, given that demographic unfairness and proxy discrimination are deeply different concepts, and there are no guarantees that using a discrimination-free price will also lead to a reduction in (let alone an elimination of) demographic unfairness – see Lindholm et al. (2024b) for an extensive discussion.

In Fig. 3 we take a more granular view, by evaluating UF for the unawareness and discrimination-free prices, on different sub-portfolios arising from segmentation by age. We select age given its importance in terms of potential discriminatory effects, as identified in Lindholm et al. (2024a). Specifically, we depict the function

$$UF_{|Age}(a) := \frac{\text{Var}(E[\pi(X) | D, \text{Age} = a] | \text{Age} = a)}{\text{Var}(\pi(X) | \text{Age} = a)}. \tag{17}$$

Note that this process is not directly related to model explainability methods, such as PDP and ALE plots (Apley & Zhu, 2020): here we are not looking at drivers of unfairness but identify demographic subgroups that may be most adversely impacted. The demographic unfairness metric exhibits its highest values for young drivers around age 20, with UF reaching approximately 0.3 for both unawareness and discrimination-free prices. This effect diminishes substantially with age, stabilizing at much lower levels (around 0.05) for drivers over 40. While both pricing approaches show similar patterns, the unawareness price consistently

demonstrates slightly higher levels of unfairness compared to the discrimination-free price, with the difference being most pronounced for drivers of ages 20-40.

Table 2 also presents the proxy discrimination metrics PD for the three pricing approaches. The standard best-estimate price shows the highest level of proxy discrimination with $PD = 0.00924$. Note that in this case, the PD measure also implicitly reflects some of the direct discrimination arising from using D as a rating factor. For unawareness prices, the value of the metric is lower, at $PD = 0.00277$, now reflecting only proxy discrimination. Finally, by construction, the discrimination-free price achieves $PD = 0$.

The relatively low, though still statistically significant, value of the PD metric for the unawareness price, suggests that proxy discrimination represents a relatively small portion of the overall price variation in this portfolio. This is consistent with the nature of modern motor insurance pricing, where key rating factors like vehicle characteristics and driver age have strong direct effects on expected claims, beyond any relation they may have with protected characteristics. Nonetheless, even such low levels of proxy discrimination may be considered problematic from a regulatory perspective, particularly given the systematic nature of the bias they represent. In particular, heterogeneity in the portfolio means that substantial discriminatory effects may be suffered by a specific sub-population of policyholders. To investigate this further, we evaluate PD on subsets of the portfolio, by age and ethnicity, that is, we plot

$$PD_{|Age, D}(a, d) := \frac{\min_{c \in \mathbb{R}, v \in \mathcal{V}} E \left[(\pi(X) - c - \sum_{d' \in \mathcal{D}} \mu(X, d') v_{d'})^2 \mid \text{Age} = a, D = d \right]}{\text{Var}(\pi(X) \mid \text{Age} = a, D = d)}. \tag{18}$$

Thus, the results reflect the comparative extent to which different ethnicity/age sub-populations are impacted by proxy discrimination, with reference to variation within their own group. The results are

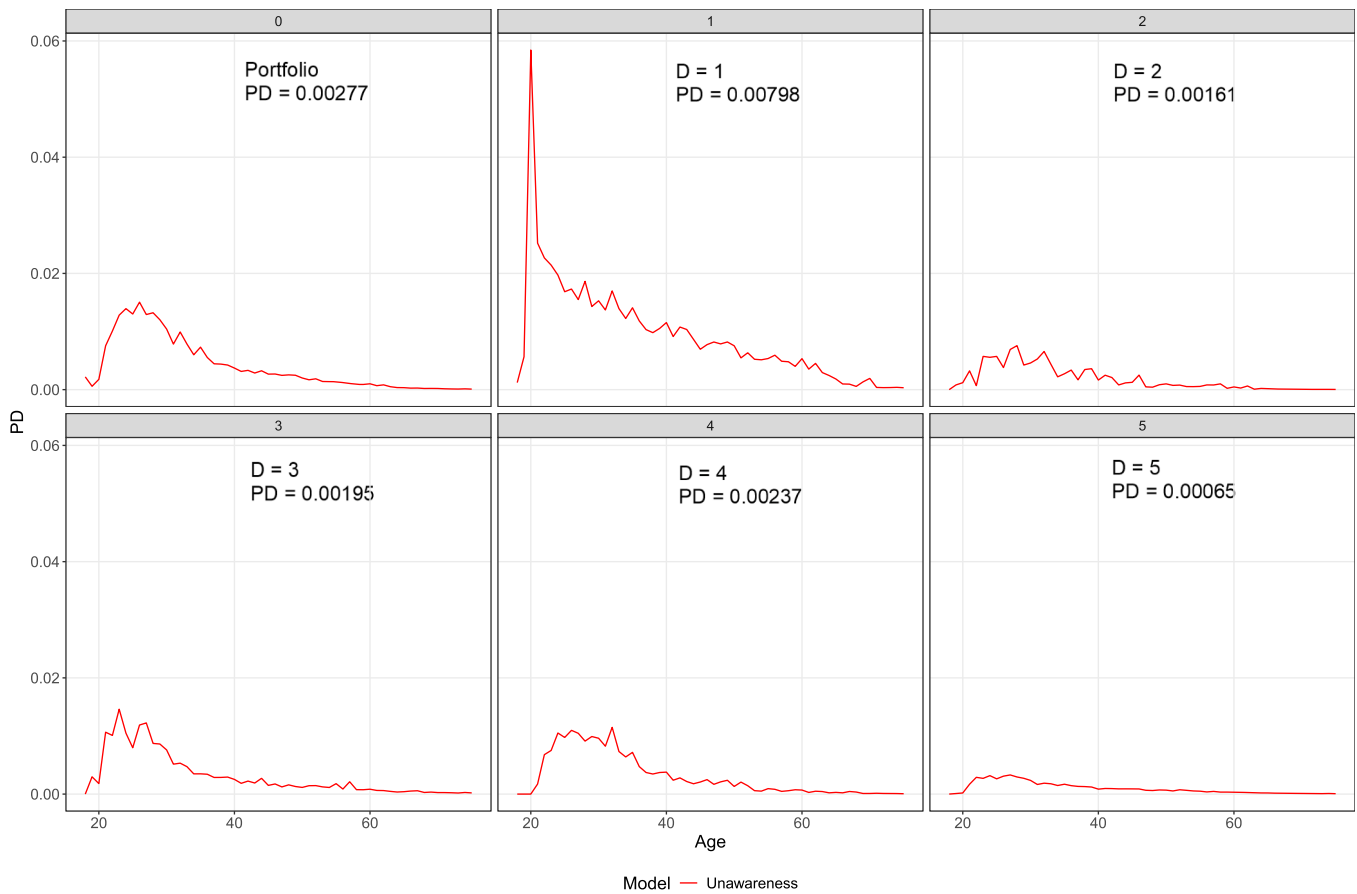


Fig. 4. Proxy discrimination metric $PD_{|Age, D}(a, d)$ evaluated for the age variable, for each ethnicity code separately and for the portfolio as a whole (Eq. (18)).

summarised Fig. 4; panels correspond, first, to the portfolio as a whole and, then, to individual ethnicity codes. Within each panel we state the value of the PD metric and show graphically PD by age. It is seen that the proxy discrimination metric shows pronounced heterogeneity across both age and ethnicity groups. Most notably, ethnicity code 1 exhibits the highest levels of proxy discrimination, with a sharp peak for young drivers around age 20. While other ethnicity groups also show some proxy discrimination effects, these are generally of lower magnitude. The age-related pattern of declining discrimination is consistent across all ethnic groups, though the rate and extent of this decline varies. These findings align with and provide additional insight to the observations in Lindholm et al. (2024a), who noted the vulnerability of young drivers with ethnicity code 1 in this dataset, to discriminatory pricing effects.

While Fig. 4 shows how the heterogeneity of proxy-discriminatory effects across the portfolio, it does not provide an interpretation of the absolute scale of the PD metric. For example, the portfolio shows $PD = 0.00277$ – but is that a lot? We investigate this issue in the Supplementary materials – answering this question in the affirmative – where we benchmark the value of the PD metric to user-defined model distortions.

We now demonstrate the practical application of the PD attribution methodology of Section 2.5 using real insurance data. Specifically, we use attributions to identify the most important (single) covariates in our dataset, which contribute most substantially to proxy discrimination when the unawareness price is used. We calculate both the first-order (PD) and total sensitivity (\widehat{PD}_T) metrics defined in Section 2.5 for each covariate. The numerical implementation follows the following process:

- i) Estimate the regression functions $x_i \mapsto \mathbb{E}[\Lambda(X, \mu) | X_i = x_i]$ and $x_{-i} \mapsto \mathbb{E}[\Lambda(X, \mu) | X_{-i} = x_{-i}]$. The conditional expectations are estimated using a random forest model (Breiman, 2001), fitted to a 50,000-record

subset of the dataset. Random forests are well-known for approximating well non-linear relationships while being relatively robust to over-fitting. Model fitting takes place in R using the ranger package with default hyperparameters (Wright & Ziegler, 2017), having found those to produce robust out-of-sample performance.

- ii) Given the regression estimates, form samples of the random vectors $\mathbb{E}[\Lambda(X, \mu) | X_i]$ and $\mathbb{E}[\Lambda(X, \mu) | X_{-i}]$ by evaluating conditional expectations empirically over all policies.
- iii) Calculate the measures (10) for $S = \{X_i\}$ for all i , by evaluating the variances with respect to the empirical distribution of X .

Fig. 5 shows these results for each covariate, where, for commercial reasons, we have disguised the covariate name (except for the age). We observe that the first-order and total sensitivity metrics produce different orderings of the covariates, though there is a broad agreement between them regarding the most significant ones. Furthermore, the attribution analysis reveals substantial variation in the extent to which different covariates contribute to proxy discrimination. Several covariates demonstrate notable proxying effects, with covariate 17 showing the strongest first-order contribution to proxy discrimination.

Finally, we consider local measurements of demographic unfairness and proxy discrimination using the tools developed in Section 3. First, we illustrate the local measure of demographic unfairness by estimating empirical versions of the distributions underlying (13), calculating $\delta_{UF}(x, d)$ w.r.t. the unawareness prices predicted for the real-world portfolio and then plotting this quantity against the age covariate, as well as against covariate 17 identified above. This analysis is shown in Fig. 7.

The total sensitivity measures generally exceed their first-order counterparts and, furthermore, total sensitivities differentiate less between covariates. This indicates substantial interaction effects (including statistical dependence) between covariates. Age demonstrates moderate

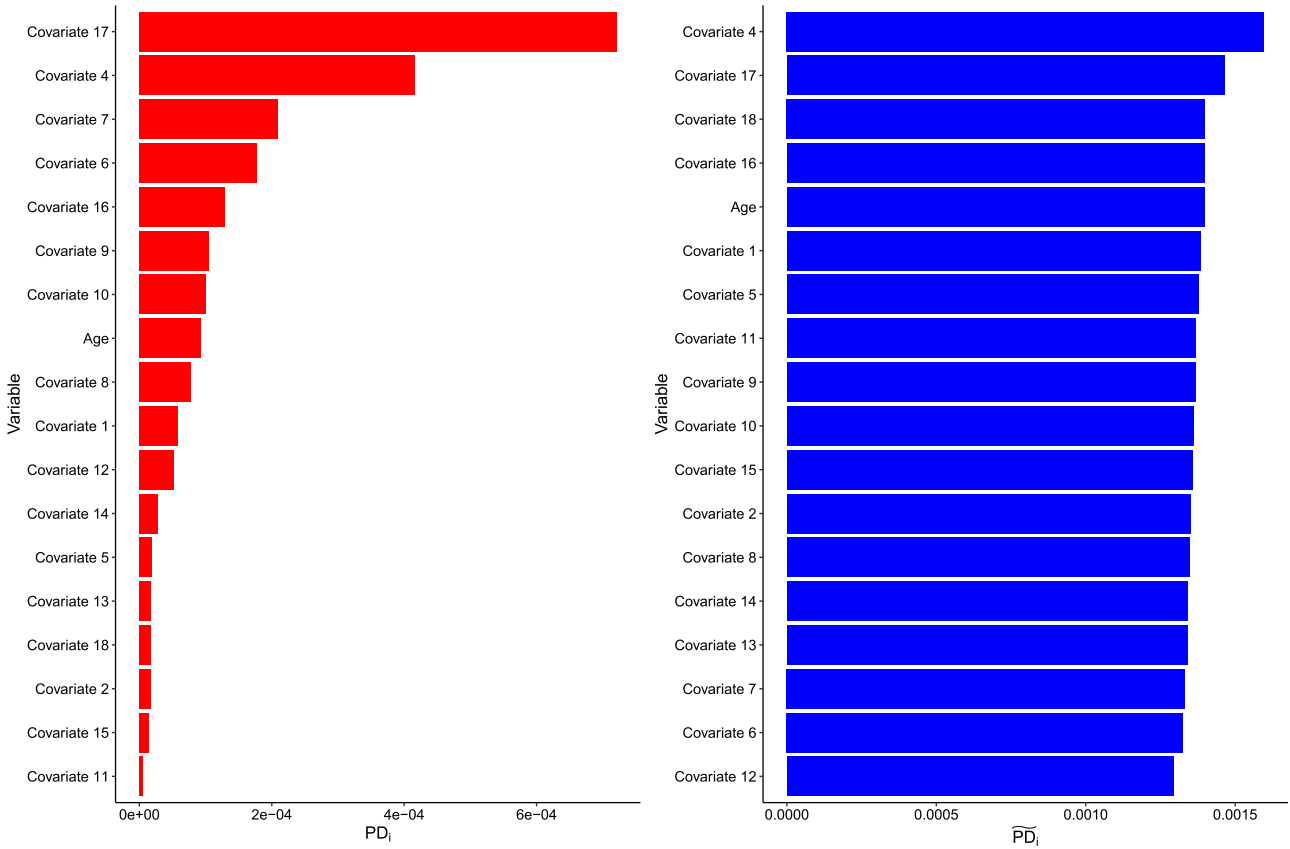


Fig. 5. First-order (PD_i) and total sensitivity (\bar{PD}_i) of unawareness prices, for each covariate in the real world dataset.

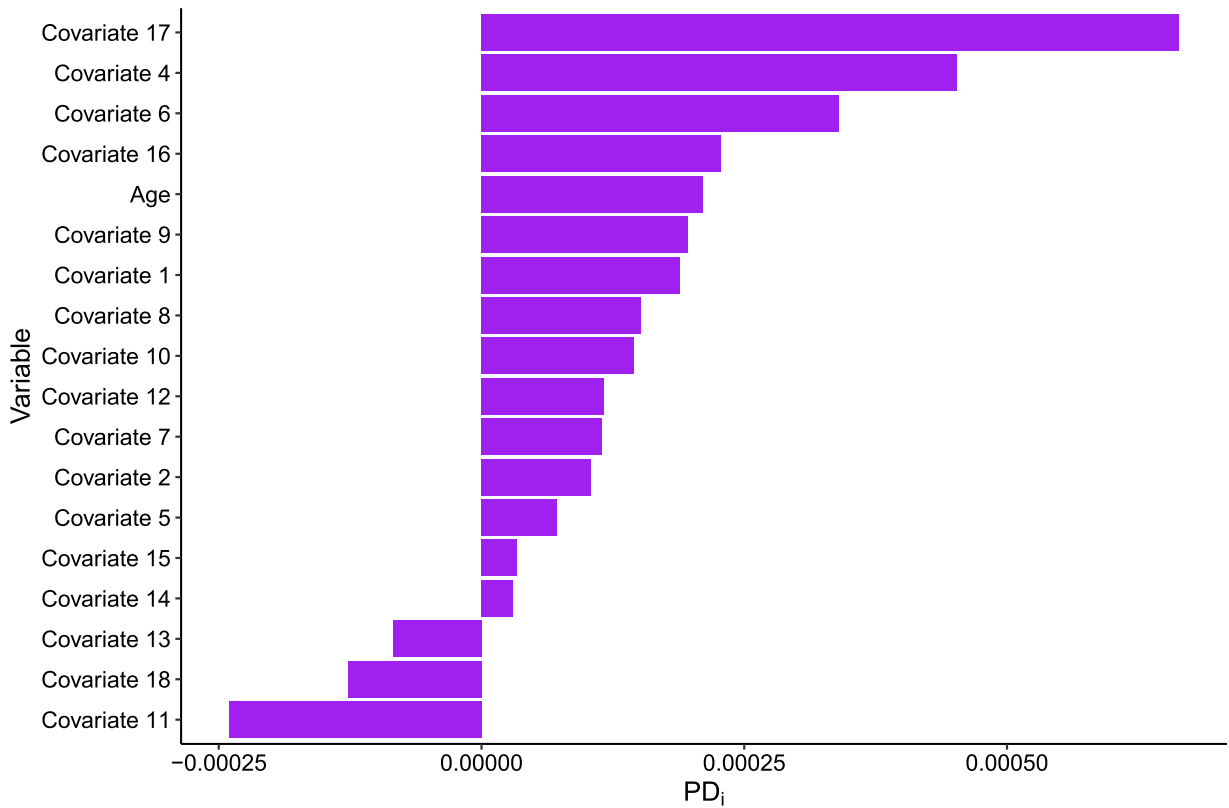


Fig. 6. CEN-SHAP attributions of proxy discrimination, implementing Shapley values (12).

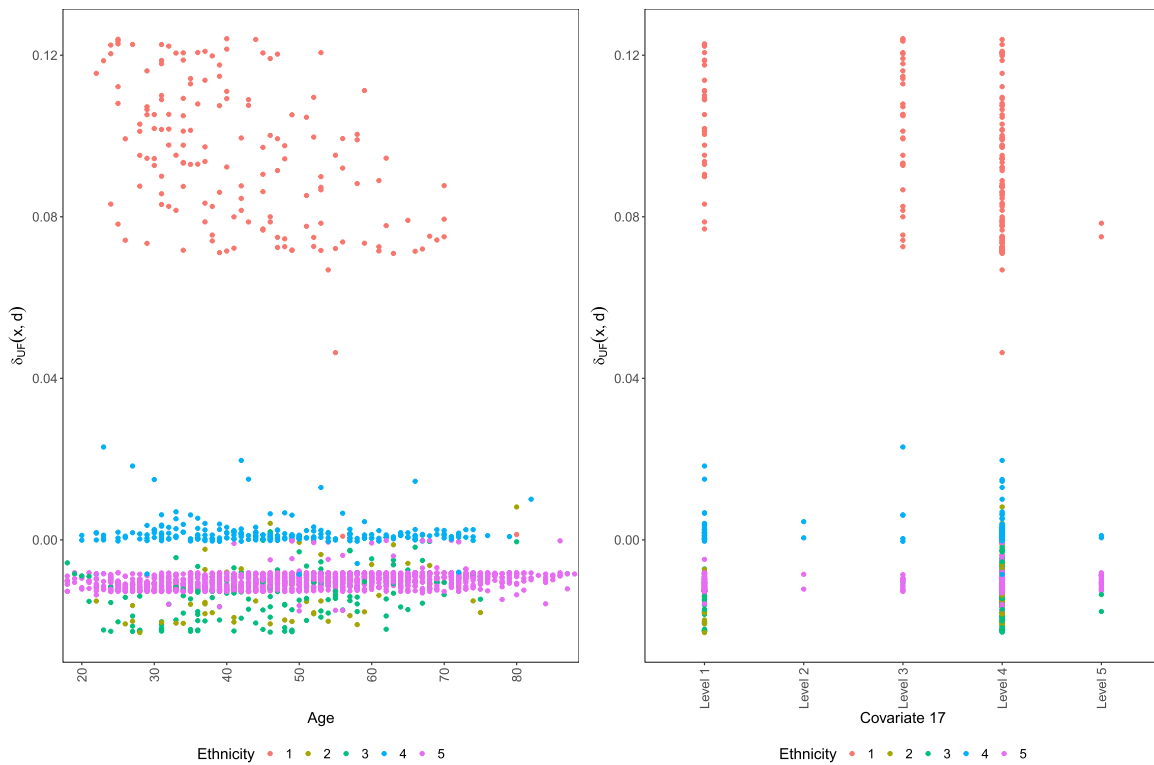


Fig. 7. Local unfairness metric $\delta_{UF}(x, d)$ plotted against the values of age (left panel) and covariate 17 (right panel).

Table 3
Conditional distribution of covariate 17, given ethnicity code (the table rows sum to 1).

Ethnicity Code	Levels of covariate 17				
	Level 1	Level 2	Level 3	Level 4	Level 5
1	0.201	0.004	0.171	0.607	0.017
2	0.153	0.003	0.016	0.820	0.009
3	0.156	0.001	0.009	0.823	0.011
4	0.137	0.002	0.018	0.832	0.011
5	0.124	0.002	0.006	0.858	0.011

proxy effects by both measures, suggesting that while demographic variables naturally correlate with ethnicity, their contribution to proxy discrimination is not necessarily dominant. Some covariates, e.g., 4 and 12, show notably different rankings between first-order and total sensitivity measures, highlighting the complex interplay between rating factors in insurance portfolios.

In Table 3, we consider covariate 17, to which the value of the PD metric is most strongly attributed using first-order sensitivities, and report its conditional distribution given each ethnicity code. Marked differences in the distribution of this variable across ethnic groups can be observed – in particular the conditional distribution of the covariate is different for ethnicity code 1, compared to other ethnicities. For instance, Level 4 represents 60.7% of ethnicity code 1 but ranges between 82.0% and 85.8% for other ethnicity codes. Correspondingly, for ethnicity code 1, there is a higher frequency of observing Level 1 of covariate 17. The different distribution for ethnicity code 1 is consistent with the higher proxy discrimination observed for this ethnicity in the earlier analysis (Fig. 4) together with the higher observed claim rates for this ethnicity shown in Table 1.

We now refine the attribution analysis of proxy discrimination by a dependence-respecting approximation to the Shapley value (12). To address the combinatoric nature of the Shapley value, we follow the well-established KernelSHAP approach of Lundberg and Lee (2017).

This still leaves the problem of evaluating the value functional $w(S)$ in Definition 6 for arbitrary coalitions S , which in turn relies on calculating the conditional expectation $\mathbb{E}[\Lambda(X, \mu) | X_S]$ for all S and instances $X_S = x_S$. For this, we use the Conditional Expectation Network (CEN) approach introduced in Richman and Wüthrich (2023). The structure of the CEN allows an efficient evaluation of predictions, conditional on arbitrary subsets of predictors. Hence, by fitting a CEN surrogate model with inputs X and response $\Lambda(X, \mu)$, we are able to calculate $\mathbb{E}[\Lambda(X, \mu) | X_S]$ for any coalition S , without the need for, e.g., model refitting. In contrast to interventional SHAP implementations, this approach does not ignore feature dependence.

We call the resulting approximate Shapley values *CEN-SHAP attributions* and depict them in Fig. 6. The largest contributors align well with the first-order attributions in Fig. 5; Age is moderate, while covariate 17 is the most prominent one. Additionally, the construction of the Shapley value allows for some model inputs, such as covariates 11 and 18, to have a negative contribution to proxy discrimination. These are covariates with low PD but high \bar{PD} , revealing their impact via interactions with other variables. The negative CEN-SHAP attributions indicate that these variables actually mitigate proxy discrimination. While we do not have access to the true joint distribution of X , one possible explanation is that these variables intervene between the sensitive attribute D and the response Y , thus reducing the scope of proxy discrimination by other covariates. For the age covariate (left panel), we observe clear patterns of demographic unfairness that vary both by age and ethnicity. Policyholders with ethnicity 1 (red points) experience consistently positive values of $\delta_{UF}(x, d)$, particularly in the 20-50 age range where values reach up to 0.12, indicating these policyholders are charged higher prices compared to the demographically fair benchmark. In contrast, policyholders with ethnicities 3-5 (green, blue and purple points) show predominantly negative or near-zero values of $\delta_{UF}(x, d)$, suggesting they benefit from demographic unfairness in pricing. The magnitude of unfairness appears to decrease with age across all ethnic groups, with the spread of $\delta_{UF}(x, d)$ values notably smaller above age 60.

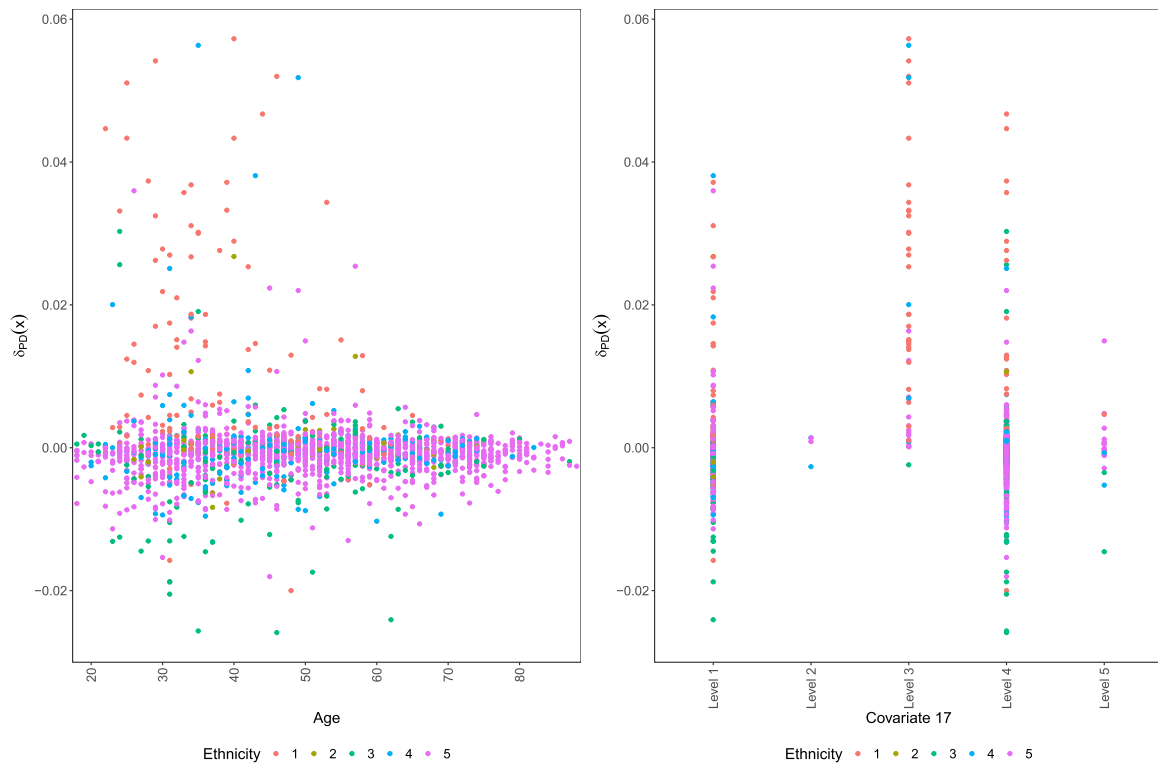


Fig. 8. Local proxy discrimination metric $\delta_{PD}(x)$ plotted against the values of age (left panel) and covariate 17 (right panel).

The right panel shows the relationship between $\delta_{UF}(x, d)$ and covariate 17, which takes discrete levels 1-5, as analyzed in Table 3. Here we observe that the magnitude of unfairness varies substantially across different levels of this covariate. The effect is particularly pronounced for ethnicity 1, where level 4 shows the highest concentration of positive $\delta_{UF}(x, d)$ values. This suggests that the interaction between ethnicity and covariate 17 is an important driver of demographic unfairness in the portfolio. These patterns suggest that demographic unfairness is not uniform across the real world portfolio but rather concentrated in particular combinations of age, ethnicity and other rating factors.

A different picture emerges when we consider the local measure of proxy discrimination, $\delta_{PD}(x)$, shown in Fig. 8. The left panel shows $\delta_{PD}(x)$ plotted against age. There is a notable concentration of positive values for ethnicity 1 (red points) in the younger age ranges, particularly between ages 20-40, suggesting these segments are most affected by proxy discrimination, which is a finding in line with Lindholm et al. (2024a). The effect appears to diminish with age, with values converging closer to zero beyond age 60. Unlike the demographic unfairness measure, we observe both positive and negative values of $\delta_{PD}(x)$ across all ethnic groups, indicating a more complex pattern of discriminatory effects.

The right panel displays $\delta_{PD}(x)$ against the levels of covariate 17, revealing distinct clustering patterns. Level 3 shows the highest concentration of discriminatory effects, with values reaching up to 0.06 particularly for ethnicity 1. This suggests that this level of covariate 17 may be particularly effective at proxying protected characteristics, as discussed above. These visualizations demonstrate how the local measure $\delta_{PD}(x)$ can identify specific segments where proxy discrimination manifests most strongly, while also highlighting areas where pricing appears to avoid such effects. The patterns suggest that proxy discrimination, while present, operates in a more nuanced way than demographic un-

fairness. This is because the latter quantifies systematic differences in outcomes across groups without considering the level of risk within each subset of the portfolio, while the former involves the implicit inference of protected characteristics from other variables within the risk rating process.

5. Conclusions

We propose measures of demographic unfairness and proxy discrimination that can be widely applied to decision-making contexts where algorithmic predictions are used and these forms of unfairness are of concern. The measure of demographic unfairness is already present in the literature (Bénésse et al., 2024), while the measure of proxy discrimination is new. For that measure, we also propose methods for attributing any proxying effects to the different covariates, including an efficient implementation of the relevant Shapley value. These measures are global, in the sense that they quantify unfairness or discrimination across a portfolio. In addition to studying the properties of these measures, we develop related local measures, which allow a quantification of demographic unfairness and proxy discrimination at a more granular level.

The proposed methods were successfully applied to a real-world insurance claims data set, where ethnicity is the protected characteristic. Importance analysis of the proxy discrimination metric revealed the key covariates generating discriminatory effects, through implicit inference of ethnicity. We note that the importance of those variables was not a priori known to us – in other words, the analysis led to a deeper understanding of the portfolio. Furthermore, given the heterogeneity of the portfolio, application of the PD measure to sub-populations of policyholders provides a more granular understanding of the groups mostly affected by proxy discrimination.

CRedit authorship contribution statement

Mathias Lindholm: Writing – original draft, Software, Methodology, Conceptualization; **Ronald Richman:** Writing – original draft, Software, Methodology, Conceptualization; **Andreas Tsanakas:** Writing – original draft, Software, Methodology, Conceptualization; **Mario V. Wüthrich:** Writing – original draft, Software, Methodology, Conceptualization.

Data availability

The real insurance data used in the numerical illustrations is not allowed to be shared due the confidentiality reasons.

Declaration of competing interest

None.

Acknowledgment

Parts of this research was carried out while M.V. Wüthrich was a KAW guest professor at Stockholm University.

Supplementary material

Supplementary material associated with this article can be found in the online version at [10.1016/j.ejor.2026.01.021](https://doi.org/10.1016/j.ejor.2026.01.021).

References

- Aas, K., Jullum, M., & land, A. L. (2021). Explaining individual predictions when features are dependent: More accurate approximations to shapley values. *Artificial Intelligence*, *298*, 103502.
- Apley, D. W., & Zhu, J. (2020). Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *82*(4), 1059–1086.
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, *104*(3), 671–732.
- Bénesse, C., Gamboa, F., Loubes, J.-M., & Boissin, T. (2024). Fairness seen as global sensitivity analysis. *Machine Learning*, *113*(5), 3205–3232.
- Bertsimas, D., Farias, V. F., & Trichakis, N. (2013). Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, *61*(1), 73–87.
- Bock, K. W. D., Coussement, K., Caigny, A. D., owiński, R. S., Baesens, B., Boute, R. N., Choi, T.-M., Delen, D., Kraus, M., Lessmann, S., & et al. (2024). Explainable ai for operational research: A defining framework, methods, applications, and a research agenda. *European Journal of Operational Research*, *317*(2), 249–272.
- Borgonovo, E., Clemente, G. P., & Rabitti, G. (2024). Why insurance regulators need to require sensitivity settings of internal models for their approval. *Finance Research Letters*, *60*, 104859.
- Borgonovo, E., Figalli, A., Plischke, E., & Savaré, G. (2025). Global sensitivity analysis via optimal transport. *Management Science*, *71*(5), 3809–3828.
- Borgonovo, E., & Plischke, E. (2016). Sensitivity analysis: A review of recent advances. *European Journal of Operational Research*, *248*(3), 869–887.
- Breiman, L. (2001). Random forests. *Machine Learning*, *45*, 5–32.
- Carrizosa, E., Halskov, T., & Morales, D. R. (2026). Wasserstein support vector machine: Support vector machines made fair. *European Journal of Operational Research*, *329*(2), 641–652.
- Carrizosa, E., Kurishchenko, K., & Morales, D. R. (2025). On enhancing the explainability and fairness of tree ensembles. *European Journal of Operational Research*, *323*(2), 599–608.
- Charpentier, A. (2024). Insurance, biases, discrimination and fairness. Springer.
- Chiappa, S., Jiang, R., Stepleton, T., Pacchiano, A., Jiang, H., & Aslanides, J. (2020). A general approach to fairness with optimal transport. *Proceedings of the AAAI Conference on Artificial Intelligence*, *34*, 3633–3640.
- Chzhen, E., Denis, C., Hebiri, M., Oneto, L., & Pontil, M. (2020). Fair regression with Wasserstein barycenters. *Advances in Neural Information Processing Systems*, *33*, 7321–7331.
- Cook, T., Greenall, A., & Sheehy, E. (2022). Discriminatory pricing: Exploring the 'ethnicity penalty' in the insurance market. Technical Report Citizens Advice.
- Côté, O., Côté, M.-P., & Charpentier, A. (2025). A fair price to pay: Exploiting causal graphs for fairness in insurance. *Journal of Risk and Insurance*, *92*(1), 33–75.
- Delbaen, F., & Majumdar, C. (2024). Approximation with independent variables. In *Peter Carr Gedenkschrift: Research advances in mathematical finance* (pp. 311–327). World Scientific.
- Frees, E. W., & Huang, F. (2023). The discriminating (pricing) actuary. *North American Actuarial Journal*, *27*(1), 2–24.
- Glasserman, P., & Li, M. (2025). Should bank stress tests be fair? *Management Science*, *71*(1), 262–278.
- Gordaliza, P., Barrio, E. D., Fabrice, G., & Loubes, J.-M. (2019). Obtaining fairness using optimal transport theory. In *International conference on machine learning* PMLR (pp. 2357–2365).
- Hiabu, M., Meyer, J. T., & Wright, M. N. (2023). Unifying local and global model explanations by functional decomposition of low dimensional structures. In *International conference on artificial intelligence and statistics* PMLR (pp. 7040–7060).
- Hurlin, C., Pérignon, C., & Saurin, S. (2024). The fairness of credit scoring models. *Management Science*, <https://doi.org/10.1287/mnsc.2022.03888>.
- Iturria, C. A. A., Hardy, M., & Marriott, P. (2024). A discrimination-free premium under a causal framework. *North American Actuarial Journal*, *24*(4), 801–821.
- Jansen, M. J. (1999). Analysis of variance designs for model output. *Computer Physics Communications*, *117*(1-2), 35–43.
- Kleinberg, J., Mullainathan, S., & Raghavan, M. (2016). Inherent trade-offs in the fair determination of risk scores. [arXiv:1609.05807](https://arxiv.org/abs/1609.05807)
- Komiyama, J., & Noda, S. (2024). On statistical discrimination as a failure of social learning: A multiarmed bandit approach. *Management Science*, <https://doi.org/10.1287/mnsc.2022.00893>.
- Kozdoi, N., Jacob, J., & Lessmann, S. (2022). Fairness in credit scoring: Assessment, implementation and profit implications. *European Journal of Operational Research*, *297*(3), 1083–1094.
- Kraus, M., Tschermutter, D., Weinzierl, S., & Zschech, P. (2024). Interpretable generalized additive neural networks. *European Journal of Operational Research*, *317*(2), 303–316.
- Lamboni, M., & Kucherenko, S. (2021). Multivariate sensitivity analysis and derivative-based global sensitivity measures with dependent variables. *Reliability Engineering & System Safety*, *212*, 107519.
- Lindholm, M., Richman, R., Tsanakas, A., & Wüthrich, M. V. (2024). A multi-task network approach for calculating discrimination-free insurance prices. *European Actuarial Journal*, *14*, 329–369.
- Lindholm, M., Richman, R., Tsanakas, A., & Wüthrich, M. V. (2022). Discrimination-free insurance pricing. *ASTIN Bulletin - The Journal of the IAA*, *52*(1), 55–89.
- Lindholm, M., Richman, R., Tsanakas, A., & Wüthrich, M. V. (2024). What is fair? Proxy discrimination vs. demographic disparities in insurance pricing. *Scandinavian Actuarial Journal*, *2024*(9), 935–970.
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, *30*, 4768–4777.
- Mara, T. A., & Tarantola, S. (2012). Variance-based sensitivity indices for models with dependent inputs. *Reliability Engineering & System Safety*, *107*, 115–121.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, *54*(6), 1–35.
- Owen, A. B. (2014). Sobol' indices and shapley value. *SIAM/ASA Journal on Uncertainty Quantification*, *2*(1), 245–251.
- Owen, A. B., & Prieur, C. (2017). On shapley value for measuring importance of dependent inputs. *SIAM/ASA Journal on Uncertainty Quantification*, *5*(1), 986–1002.
- Pesenti, S. M., Millosovich, P., & Tsanakas, A. (2021). Cascade sensitivity measures. *Risk Analysis*, *41*(12), 2392–2414.
- Pesenti, S. M., Millosovich, P., & Tsanakas, A. (2025). Differential quantile-based sensitivity in discontinuous models. *European Journal of Operational Research*, *322*(2), 554–572.
- Plischke, E., & Borgonovo, E. (2020). Fighting the curse of sparsity: Probabilistic sensitivity measures from cumulative distribution functions. *Risk Analysis*, *40*(12), 2639–2660.
- Prince, A. E., & Schwarcz, D. (2019). Proxy discrimination in the age of artificial intelligence and big data. *Iowa L. Rev.*, *105*, 1257.
- Rabitti, G., & Borgonovo, E. (2020). Is mortality or interest rate the most important risk in annuity models? A comparison of sensitivity analysis methods. *Insurance: Mathematics and Economics*, *95*, 48–58.
- Radovanović, S., Savić, G., Delibašić, B., & Suknović, M. (2022). Fairdea - removing disparate impact from efficiency scores. *European Journal of Operational Research*, *301*(3), 1088–1098.
- Richman, R., & Wüthrich, M. V. (2023). Conditional expectation network for SHAP. Technical Report 4514891. Accessed 2025-09-29. SSRN Working Paper.
- Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., & Tarantola, S. (2010). Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. *Computer Physics Communications*, *181*(2), 259–270.
- Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). *Global Sensitivity Analysis: the Primer*. John Wiley & Sons.
- Shapley, L. S. (1953). *A value for n-person games* (vol. II). Princeton University Press Princeton.
- Sobol', I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and Computers in Simulation*, *55*(1-3), 271–280.
- Song, E., Nelson, B. L., & Staum, J. (2016). Shapley effects for global sensitivity analysis: Theory and computation. *SIAM/ASA Journal on Uncertainty Quantification*, *4*(1), 1060–1083.
- Tschantz, M. C. (2022). What is proxy discrimination? In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency* (pp. 1993–2003).
- Vallarino, A., Rabitti, G., & Chokami, A. K. (2024). Construction of rating systems using global sensitivity analysis: A numerical investigation. *ASTIN Bulletin - The Journal of the IAA*, *54*(1), 25–45.
- Veiga, S. D., Wahl, F., & Gamboa, F. (2009). Local polynomial estimation for sensitivity analysis on models with correlated inputs. *Technometrics*, *51*(4), 452–463.
- Wright, M. N., & Ziegler, A. (2017). ranger: A fast implementation of random forests for high dimensional data in C and R. *Journal of Statistical Software*, *77*(1), 1–17.
- Wüthrich, M. V., & Merz, M. (2023). *Statistical foundations of actuarial learning and its applications*. Springer Nature.