

DISCRIMINATION-FREE INSURANCE PRICING

M. LINDHOLM, R. RICHMAN, A. TSANAKAS, M.V. WÜTHRICH

ABSTRACT. A simple formula for non-discriminatory insurance pricing is introduced. This formula is based on the assumption that certain individual (discriminatory) policyholder information is not allowed to be used for insurance pricing. The suggested procedure can be summarized as follows: First, we construct a price that is based on all available information, including discriminatory information. Thereafter, we average out the effect of discriminatory information. This averaging out is done such that discriminatory information can also not be inferred from the remaining non-discriminatory one, thus, neither allowing for direct nor for indirect discrimination.

Keywords: Discrimination, differentiation, insurance pricing, individual policy characteristics, discriminatory covariates, direct discrimination, indirect discrimination, neural networks, complex algorithmic models, causal inference, confounding.

1. INTRODUCTION

The basic question is as follows: given observed claims of insurance policies for which we have access to individual policyholder characteristics (e.g. gender), how can we construct an insurance tariff for which we can assure that it does not discriminate, say, w.r.t. gender? This is a highly relevant question, for instance, the current EU legislation requires gender neutral pricing, see [9]. This question has become even more relevant due to big data and recent developments in complex algorithmic models because such algorithms may be able to unfold, e.g., the gender from other detailed individual policyholder characteristics. For a broader overview on anti-discrimination laws we refer to Avraham et al. [2].

We aim at developing pricing formulas that shall be devoid of discrimination, while the insurer is still able to differentiate between policyholders with respect to non-discriminatory features. We assume that an insurer has access to policyholder data that can be split into discriminatory (e.g. gender, ethnicity) and non-discriminatory features (e.g. age, smoking habits). For the purposes of this paper we distinguish between *direct* and *indirect discrimination* by saying that:¹

- A pricing model *avoids direct discrimination*, if none of the discriminatory features is used as a rating factor.
- A pricing model *avoids indirect discrimination*, if it avoids direct discrimination and, furthermore, the non-discriminatory features are used in a way that does not allow implicit inference of discriminatory features from them.

Date: Version of January 16, 2020.

¹We note that the regulatory text [9] provides definitions of direct and indirect discrimination, which to an extent motivate our technical arguments. However, our definitions are drawn from an *actuarial* thinking and we do not make any claim about their correspondence to the *legal* definitions of those terms.

In this paper we make two key points. First, we demonstrate that avoiding direct discrimination does not necessarily entail avoiding indirect discrimination. Consequently, just ignoring discriminatory features in the calculation of insurance prices does not necessarily yield discrimination-free prices. Hence, “unawareness” of discriminatory features is not a solution to our problem. Second, we propose a simple pricing formula that avoids both direct and indirect discrimination. While the formula only uses non-discriminatory features as rating factors, it introduces an adjustment, which requires knowledge of policyholders’ discriminatory features.

These concepts are illustrated in the following simple example, a more sophisticated version of which is given in Section 6.

Example 1. Assume that you have observed a health insurance product and obtained the following claim counts $(n_{i,j})_{i,j=0,1}$ and claim exposures $(r_{i,j})_{i,j=0,1}$:

$n_{i,j}$	woman	man	row total	$r_{i,j}$	woman	man	row total
smoker	32	4	36	smoker	133	24	157
non-smoker	28	48	76	non-smoker	131	301	432
column total	60	52	112	column total	264	325	589

where $i = 1$ corresponds to “smoker” and $j = 1$ corresponds to “woman”. Based on the above contingency tables we may estimate the claim frequencies $\lambda_{i,j}$ by the empirical frequency $\hat{\lambda}_{i,j} = n_{i,j}/r_{i,j}$. Assume now that gender is considered a discriminatory characteristic. Thus, in order to avoid direct discrimination, its explicit influence on the calculated insurance price needs to be removed. The standard way of doing this is to consider the aggregated estimators (row sums) $\hat{\lambda}_{i,\bullet} = n_{i,\bullet}/r_{i,\bullet} = (n_{i,0} + n_{i,1})/(r_{i,0} + r_{i,1})$. This approach produces, e.g. for smokers,

$$\hat{\lambda}_{1,\bullet} = \frac{36}{157} = 0.229.$$

The estimate $\hat{\lambda}_{1,\bullet}$ (and a premium for smokers based on it), thus, can be produced by completely ignoring policyholders’ gender information. But one can note that an alternative representation of $\hat{\lambda}_{1,\bullet}$ is

$$(1) \quad \hat{\lambda}_{1,\bullet} = \hat{\lambda}_{1,1}\hat{\mathbb{P}}(\text{woman} \mid \text{smoker}) + \hat{\lambda}_{1,0}\hat{\mathbb{P}}(\text{man} \mid \text{smoker}),$$

where $\hat{\mathbb{P}}$ refers to the empirical distribution obtained from the data. Hence, the estimate $\hat{\lambda}_{1,\bullet}$ not only contains information about the influence of smoking on producing a claim, but also, via $\hat{\mathbb{P}}(\text{gender} \mid \text{smoking habits})$, about the propensity of smokers to be female or male. In our case, because smoking habits substantially differ between genders (a smoker is a woman with probability $133/157 = 85\%$, whereas a non-smoker is a woman with probability $131/432 = 30\%$), it is indeed the case that the above approach implicitly infers the likely gender from smoking habits. Therefore, indirect discrimination is present. ■

Literature review. Although an issue of key relevance for actuarial pricing, relatively little attention has been paid to the issue of discrimination-free pricing within the actuarial literature. In a discussion of the implications of the EU gender

legislation, it is suggested in Guillén [12] that covariates highly correlated with gender can be used as proxies by insurance companies, which from our perspective results in indirect discrimination. Focusing on the case of mortality pricing, Chen and Vigna [5] criticize the industry practice of deriving a unisex life table by mixing the life tables for each gender on the grounds that this does not respect the principles of actuarial fairness, which is to say that the total unisex premiums charged for the portfolio are not equal to the total premiums charged using gender specific life tables. They provide alternative approaches without this shortcoming; note that the discrimination-free price in Section 2 will reproduce the pricing formulas of Chen and Vigna [5]. The implications of unisex pricing on insurer capital requirements in the context of Solvency II are examined in Chen et al. [4], and an ALM approach to unisex pricing is taken in Burszas et al. [3], where also the concept of “gender mix risk” is discussed. Market implications of unisex tariffs are discussed in Sass and Seifried [19], see also De Jong and Ferris [7] for a discussion of adverse selection stemming from restrictions on risk classification.

The issue of indirect discrimination occurring when using the above mentioned “unawareness” approach, see also Definition 2 below, has been discussed in Pope and Sydnor [18] and Kusner et al. [15]. The procedure for discrimination-free pricing provided in Pope and Sydnor [18] is the same as our proposal in Section 2, below, but they do not provide a probabilistic justification for the procedure. Moreover, in Pope and Sydnor [18] the (critical) issue of a potential bias at the portfolio level and corrections for this are not addressed. In Aseervatham et al. [1], the method of Pope and Sydnor [18] is applied in the context of auto insurance pricing, but similarly, the issue of a bias at the portfolio level is not considered.

Relatively few examples of causal inference applied within an insurance context are known to the authors. At the renewal of insurance policies, some insurers seek to estimate policyholder demand elasticity by randomly varying renewal prices for a subset of policyholders (i.e. a form of a randomized controlled trial is conducted) and estimating the impact on the probability of renewal. Once the demand elasticities have been estimated, a profit maximizing pricing policy can be established in a practice referred to as price optimization, see e.g. Krikler et al. [13]. Also within this context, Guelman and Guillén [11] apply methods from causal inference to estimate demand elasticity functions from observational data collected by an insurer.

Statement of contribution. This paper makes several contributions. Firstly, we formally define direct and indirect discrimination, a rigorous probabilistic account of discrimination-free prices is given and the existence of these prices is examined. Secondly, we justify the discrimination-free pricing formula of Section 2 using the tools of causal inference. Thirdly, the bias correction methodologies of Section 4 are a novel contribution to the literature. Finally, the discrimination-free pricing procedure is illustrated using a machine learning algorithm in Section 6.

Organization of the paper. In Section 2 we discuss different insurance prices which comprises the *best-estimate price* considering all available information, the *unawareness price* which avoids direct discrimination, and the *discrimination-free price* which avoids direct and indirect discrimination. Moreover, in Subsection 2.3 we give a mathematical definition of indirect discrimination which is based on a change of probability measure. Special versions of discrimination-free prices can nicely be interpreted in terms of causal inference, which describes how to adjust for confounding covariates; this is discussed in Section 3. Best-estimate prices (in

terms of conditional expectations) naturally lead to an unbiased pricing system on a portfolio level. Discrimination-free pricing does not necessarily have this property. This deficiency is going to be analyzed in Section 4. In Section 5 we describe the implementation of the procedure for calculating discrimination-free prices. This is verified in more detail in Section 6 based on a longer numerical illustration. The paper ends with a number of concluding remarks in Section 7.

2. DISCRIMINATION-FREE PRICING

2.1. Definition of discrimination-free prices. We denote by $(\Omega, \mathcal{F}, \mathbb{P})$ the underlying probability space with physical probability measure \mathbb{P} . For a given portfolio of insurance policies, let \mathbf{D} denote the vector of *discriminatory covariates* (characteristics, features, explanatory variables) of policyholders, and let \mathbf{X} denote the vector of *non-discriminatory covariates*. We assume that \mathbf{X} and \mathbf{D} are random vectors on $(\Omega, \mathcal{F}, \mathbb{P})$; the randomness of these covariate vectors represents variations between policyholders within a given portfolio. An individual policyholder is obtained by conditioning on $\mathbf{X} = \mathbf{x}$, $\mathbf{D} = \mathbf{d}$ (which can be interpreted as choosing an insurance policy at random from the portfolio). For simplicity, we will denote the marginal and conditional distributions of covariates under \mathbb{P} by $\mathbf{X} \sim \mathbb{P}(\mathbf{x})$, $\mathbf{D} \sim \mathbb{P}(\mathbf{d})$ and $(\mathbf{D} | \mathbf{X} = \mathbf{x}) \sim \mathbb{P}(\mathbf{d} | \mathbf{x})$, respectively.

A policyholder claim is denoted by the random variable Y . The claim Y typically depends on (but is not fully determined by) both the discriminatory covariates \mathbf{D} and the non-discriminatory ones \mathbf{X} . Our aim is to price such a claim Y , with the resulting price being free from direct as well as indirect discrimination, according to the arguments of Section 1. A technical definition of these concepts will be given in Section 2.3, below.

In the sequel, it will be useful to assume $Y, \mathbf{X}, \mathbf{D} \in \mathcal{L}^2(\Omega, \mathcal{F}, \mathbb{P})$. This assumption is not crucial for defining discrimination-free prices, but will allow us to give more intuitive interpretations in terms of orthogonal projections and minimal distances. Our notion of price will be based on conditional expectations of Y , when conditioning on different subsets of covariates. We first introduce a number of different prices that are important for the subsequent discussions and derivations.

Definition 2 (best-estimate price). *The best-estimate price for Y w.r.t. (\mathbf{X}, \mathbf{D}) is defined by*

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

Remark 3.

- (a) We call the price $\mu(\mathbf{X}, \mathbf{D})$ “best-estimate” because it minimizes the \mathcal{L}^2 -distance of all (\mathbf{X}, \mathbf{D}) -measurable prices to Y , i.e. $\mu(\mathbf{X}, \mathbf{D})$ is the orthogonal projection of Y onto the sub-space generated by (\mathbf{X}, \mathbf{D}) .
- (b) In general, the best-estimate price is not discrimination-free, unless we are in the special case of $\mu(\mathbf{X}, \mathbf{D}) = \mu(\mathbf{X})$.
- (c) The best-estimate price is *unbiased* w.r.t. Y , that is,

$$\mu := \mathbb{E}[Y] = \mathbb{E}[\mu(\mathbf{X}, \mathbf{D})];$$

by the tower property for conditional expectations. Unbiasedness is important because it indicates that best-estimate prices achieve on average the correct price level for the portfolio.

As already discussed in Example 1, an initial attempt at achieving a discrimination-free price arises through simply ignoring discriminatory covariates \mathbf{D} .

Definition 4 (unawareness price). *The unawareness price for Y w.r.t. \mathbf{X} is defined by*

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

Remark 5.

- (a) As the price $\mu(\mathbf{X})$ does not depend explicitly on \mathbf{D} , it avoids direct discrimination. However, in general, the unawareness price will produce indirect discrimination, as was discussed in Example 1; see also Kusner et al. [15]. Specifically, we can write the unawareness price as

$$(3) \quad \mu(\mathbf{X}) = \int_{\mathbf{d}} \mu(\mathbf{X}, \mathbf{d}) \, d\mathbb{P}(\mathbf{d} \mid \mathbf{X}).$$

Indirect discrimination arises because the conditional probability $d\mathbb{P}(\mathbf{d} \mid \mathbf{X})$ enables inference of discriminatory covariates \mathbf{D} from non-discriminatory ones \mathbf{X} . Such discrimination is avoided in the special case when \mathbf{D} and \mathbf{X} are independent, since then it holds that $d\mathbb{P}(\mathbf{d} \mid \mathbf{X}) = d\mathbb{P}(\mathbf{d})$.

- (b) The price $\mu(\mathbf{X})$ minimizes the \mathcal{L}^2 -distance to Y based solely on \mathbf{X} , i.e. it is the best price w.r.t. information \mathbf{X} . At the same time, the price $\mu(\mathbf{X})$ also minimizes the \mathcal{L}^2 -distance to $\mu(\mathbf{X}, \mathbf{D})$, by a simple application of the Pythagorean theorem. Note that

$$\|\mu(\mathbf{X}) - \mu(\mathbf{X}, \mathbf{D})\|_2^2 = \mathbb{E}[\text{Var}(\mu(\mathbf{X}, \mathbf{D}) \mid \mathbf{X})],$$

which intuitively should decrease with increasing dependence between \mathbf{X} and \mathbf{D} . Hence, the quality in the approximation of $\mu(\mathbf{X}, \mathbf{D})$ using $\mu(\mathbf{X})$ should be good if \mathbf{X} essentially is a function of \mathbf{D} , i.e., if the non-discriminatory covariates \mathbf{X} allow us to almost perfectly infer the discriminatory covariates \mathbf{D} .

- (c) The unawareness price is unbiased, since

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

We now propose a price that is free of both direct and indirect discrimination.

Definition 6 (discrimination-free price). *A discrimination-free price for Y w.r.t. \mathbf{X} is defined by*

$$(4) \quad h^*(\mathbf{X}) := \int_{\mathbf{d}} \mu(\mathbf{X}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d}),$$

where the distribution $\mathbb{P}^*(\mathbf{d})$ is defined on the same range as the marginal distribution of discriminatory factors $\mathbf{D} \sim \mathbb{P}(\mathbf{d})$.

Remark 7.

- (a) The discrimination-free price (4) is obtained by averaging best-estimate prices over discriminatory covariates, using a (potentially arbitrary) marginal distribution $\mathbb{P}^*(\mathbf{d})$, where it is the imposed marginalization w.r.t. \mathbf{D} which is the crucial step (e.g. nothing prevents us from using $\mathbb{P}^*(\mathbf{D}) = \mathbb{P}(\mathbf{D})$). Given that the price $h^*(\mathbf{X})$ does not explicitly depend on \mathbf{D} , it is obviously free from direct discrimination. We argue that the averaging construction proposed in (4) also removes indirect discrimination. First, note that while (4) appears similar to (3), there is a key difference: *discrimination-free prices do not in any way*

depend on the conditional distribution $\mathbb{P}(\mathbf{d} \mid \mathbf{X})$ – hence they do not enable to infer the discriminatory covariates from the non-discriminatory ones. This will be further discussed in Section 2.3 and verified in the case study of Section 6. In the special case of \mathbf{X} and \mathbf{D} being independent, it follows that $h^*(\mathbf{X}) = \mu(\mathbf{X})$.

- (b) Definition 6 is designed to *remove* the possible explanatory power that \mathbf{X} may have for \mathbf{D} ; it does *not* assume independence between \mathbf{X} and \mathbf{D} in the given portfolio. This point will be made more precise in Section 2.3, and in Section 2.4 we discuss existence of discrimination-free prices as well as alternative interpretations of $h^*(\mathbf{X})$.
- (c) Definition 6 can also be motivated by arguments from causal inference. Specifically, formulas like (4) are used to quantify the direct causal effect of \mathbf{X} on Y . Our distribution-free price corresponds to the (expected value version) of the so-called “back-door” adjustment formula in causal inference settings, when seeing \mathbf{D} as a potential confounder of \mathbf{X} ; these ideas will be discussed in more detail in Section 3. Furthermore, formula (4) using the choice $\mathbb{P}^*(\mathbf{d}) = \mathbb{P}(\mathbf{d})$ corresponds precisely to the Partial Dependence Plot (PDP) introduced by Friedman [10], see also Zhao and Hastie [20]; in particular, our estimated model approach in Section 5 will correspond to formula (53) in Friedman [10].
- (d) Prices obtained using (4) will in general *not* be unbiased, since

$$(5) \quad \mu = \mathbb{E}[Y] \neq \mathbb{E}[h^*(\mathbf{X})] = \int_{\mathbf{x}, \mathbf{d}} \mu(\mathbf{x}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d})d\mathbb{P}(\mathbf{x}),$$

even for the special choice $\mathbb{P}^*(\mathbf{d}) = \mathbb{P}(\mathbf{d})$. For that reason, the need emerges for portfolio level price corrections, which will be discussed in Section 4.

- (e) Note that, given the potential arbitrariness of \mathbb{P}^* , evaluation of discrimination-free prices only requires knowledge of the mapping $(\mathbf{x}, \mathbf{d}) \mapsto \mu(\mathbf{x}, \mathbf{d})$, where $\mu(\mathbf{x}, \mathbf{d})$ may be an (algorithmically derived implicit) regression function. Nevertheless, as pointed out in the previous remark, if one aims to correct a potential bias of $h^*(\mathbf{X})$, it is necessary to consider (estimate under) the “real-world” probability measure \mathbb{P} .

2.2. Choice of weighting distributions for discriminatory covariates. From Definition 6 it is clear that the distribution $\mathbb{P}^*(\mathbf{d})$ can be chosen rather freely. A simple choice is to set $\mathbb{P}^*(\mathbf{d}) = \mathbb{P}(\mathbf{d})$, that is, average in (4) w.r.t. the marginal distribution of the discriminatory characteristics in the portfolio. This choice is supported by causal inference (CI) arguments in the settings of Section 3, below. Therefore, we introduce the special notation for this choice

$$(6) \quad h^{(CI)}(\mathbf{X}) := \int_{\mathbf{d}} \mu(\mathbf{X}, \mathbf{d}) \, d\mathbb{P}(\mathbf{d}).$$

We illustrate how $h^{(CI)}(\mathbf{X})$ is evaluated in the context of Example 1.

Example 8. Recall that in Example 1 we argued that aggregated estimators (row sums) $\hat{\lambda}_{i, \bullet}$ are discriminatory because gender may be inferred from smoking habits. The price $h^{(CI)}(\mathbf{X})$ proposes to remove this effect by replacing the conditional probability $\mathbb{P}(\text{gender} \mid \text{smoking habits})$ with $\mathbb{P}(\text{gender})$. This implies that the frequency

estimate for smokers, $\widehat{\lambda}_{1,\bullet}$, should be replaced by

$$\begin{aligned}\widehat{\lambda}_{1,\bullet}^{(CI)} &= \widehat{\lambda}_{1,1}\widehat{\mathbb{P}}(\text{woman}) + \widehat{\lambda}_{1,0}\widehat{\mathbb{P}}(\text{man}) \\ &= \frac{32}{133} \cdot \frac{264}{589} + \frac{4}{24} \cdot \frac{325}{589} \\ &= 0.200 < 0.229 = \widehat{\lambda}_{1,\bullet}.\end{aligned}$$

■

Furthermore, it is useful to consider the extrema of discrimination-free prices. Consider the following distribution-free prices

$$\begin{aligned}h^{(+)}(\mathbf{X}) &:= \sup_{\mathbb{P}^*} \int_{\mathbf{d}} \mu(\mathbf{X}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d}), \\ h^{(-)}(\mathbf{X}) &:= \inf_{\mathbb{P}^*} \int_{\mathbf{d}} \mu(\mathbf{X}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d}).\end{aligned}$$

$h^{(+)}(\mathbf{X})$ and $h^{(-)}(\mathbf{X})$ correspond to the essential supremum and infimum over \mathbf{d} in the range of \mathbf{D} , respectively. Thus, for non-discriminatory covariates $\mathbf{X} = \mathbf{x}$, this immediately gives us

$$h^{(-)}(\mathbf{x}) \leq h^*(\mathbf{x}), h^{(CI)}(\mathbf{x}), \mu(\mathbf{x}) \leq h^{(+)}(\mathbf{x}).$$

Moreover, for the bias property we get the following relationship

$$\int_{\mathbf{x}} h^{(-)}(\mathbf{x}) \, d\mathbb{P}(\mathbf{x}) \leq \mathbb{E}[h^*(\mathbf{X})], \mu \leq \int_{\mathbf{x}} h^{(+)}(\mathbf{x}) \, d\mathbb{P}(\mathbf{x}).$$

By definition $h^{(+)}(\mathbf{x})$ corresponds to the “worst” (or most “prudent”) price, and has been discussed in the context of unisex pricing in Chen and Vigna [5].

Finally, recall that the distribution-free price (4) in general is biased. An alternative possibility for the choice of $\mathbb{P}^*(\mathbf{d})$ is to additionally require unbiasedness in (5). In the simple case of a binary discriminatory covariate like gender in Example 8, this reduces to choosing $\mathbb{P}^*(\text{woman})$ such that we obtain an identity in (5). A more general construction of unbiased prices via choices of $\mathbb{P}^*(\mathbf{d})$ is presented in Section 4, below.

A special case corresponds to an additive best-estimate price, in the sense that $\mu(\mathbf{X}, \mathbf{D}) = \mu_1(\mathbf{X}) + \mu_2(\mathbf{D})$. Then, the simple choice $\mathbb{P}^*(\mathbf{d}) = \mathbb{P}(\mathbf{d})$ is appealing, as it provides an unbiased price. Note that

$$h^{(CI)}(\mathbf{X}) = \int_{\mathbf{d}} \mu_1(\mathbf{X}) \, d\mathbb{P}(\mathbf{d}) + \int_{\mathbf{d}} \mu_2(\mathbf{d}) \, d\mathbb{P}(\mathbf{d}) = \mu_1(\mathbf{X}) + \mathbb{E}[\mu_2(\mathbf{D})],$$

which implies

$$\mathbb{E}[h^{(CI)}(\mathbf{X})] = \mathbb{E}[\mu_1(\mathbf{X})] + \mathbb{E}[\mu_2(\mathbf{D})] = \mathbb{E}[\mu(\mathbf{X}, \mathbf{D})] = \mu.$$

2.3. Revisiting direct and indirect discrimination. In this section, following the development of our ideas so far, we provide more technical definitions of prices that avoid direct and indirect discrimination.²

Choose an arbitrary probability measure \mathbb{P}^* on the measurable space (Ω, \mathcal{F}) such that $Y \in \mathcal{L}^1(\Omega, \mathcal{F}, \mathbb{P}^*)$. Choose a (sub-)vector \mathbf{Z} of the covariates (\mathbf{X}, \mathbf{D}) and define the $(\mathbb{P}^*, \mathbf{Z})$ -conditional-expectation price by

$$\mu^*(\mathbf{Z}) := \mathbb{E}^*[Y \mid \mathbf{Z}],$$

²(Fast) readers not familiar with changes of probability measures may skip this section.

where \mathbb{E}^* denotes the expectation under \mathbb{P}^* .

Definition 9. *A price avoids direct discrimination, if it can be written as*

$$\mu^*(\mathbf{Z}) = \mathbb{E}^*[Y \mid \mathbf{Z}],$$

where \mathbf{Z} is $\sigma(\mathbf{X})$ -measurable, and where the expectation is taken w.r.t. a probability measure \mathbb{P}^* on (Ω, \mathcal{F}) such that $Y \in \mathcal{L}^1(\Omega, \mathcal{F}, \mathbb{P}^*)$.

Remark 10.

- (a) Definition 9 says that a price avoids direct discrimination if it can be written as a measurable function of the non-discriminatory covariates \mathbf{X} . For $\mathbf{Z} = \mathbf{X}$ we receive maximal use of non-discriminatory information (relative to \mathbb{P}^*), therefore, we typically work with $\mathbf{Z} = \mathbf{X}$.
- (b) The choice $\mathbb{P}^* = \mathbb{P}$ (and $\mathbf{Z} = \mathbf{X}$) provides the unawareness price $\mu(\mathbf{X})$ of Definition 4 which, thus, avoids direct discrimination.
- (c) Importantly, under the choice $\mathbb{P}^* = \mathbb{P}$, the unawareness price $\mu(\mathbf{X})$ can be calculated without explicit knowledge of $\mu(\mathbf{X}, \mathbf{D})$ – hence it does *not* require collection of discriminatory policyholder information. This also applies if we need to estimate $\mu(\mathbf{X})$ from data, see (17) below.
- (d) Note that by using standard Radon-Nikodym arguments (assuming existence of relevant quantities), Definition 9 can be re-written according to

$$\mu^*(\mathbf{Z}) = \mathbb{E}[Y^* \mid \mathbf{Z}],$$

where $Y^* = Y d\mathbb{P}^*(Y \mid \mathbf{Z})/d\mathbb{P}(Y \mid \mathbf{Z})$. That is, the use of \mathbb{P}^* will in general correspond to using another actuarial pricing principle than the fair (pure risk premium) pricing principle under \mathbb{P} . Another observation, which is a direct consequence of the above, relating to a potential bias, is that

$$\mathbb{E}[\mu^*(\mathbf{Z})] = \mathbb{E}[\mathbb{E}[Y^* \mid \mathbf{Z}]] = \mathbb{E}[Y^*].$$

These ideas will not be pursued further other than that we will give an example when \mathbb{P}^* is chosen such that $\mathbb{E}[Y] = \mathbb{E}[Y^*]$, where it is possible to interpret Y^* in terms of a certain type of risk-adjustment, see Section 4.

Now, indirect discrimination can be defined.

Definition 11. *A price $\mu^*(\mathbf{Z})$ that avoids direct discrimination is said to avoid indirect discrimination if \mathbf{Z} and \mathbf{D} are independent under \mathbb{P}^* .*

Independence under \mathbb{P}^* effects the decoupling of discriminatory covariates from non-discriminatory ones, for specific policyholders. Thus, according to Definition 11, a price that avoids indirect discrimination satisfies

$$(7) \quad \mu^*(\mathbf{Z}) = \int_{\mathbf{d}} \mu^*(\mathbf{Z}, \mathbf{d}) d\mathbb{P}^*(\mathbf{d} \mid \mathbf{Z}) = \int_{\mathbf{d}} \mu^*(\mathbf{Z}, \mathbf{d}) d\mathbb{P}^*(\mathbf{d}),$$

where $\mu^*(\mathbf{Z}, \mathbf{d}) = \mathbb{E}^*[Y \mid \mathbf{Z}, \mathbf{D} = \mathbf{d}]$.

Remark 12.

- (a) It is important to note that the independence in Definition 11 is an artifice of the introduced probability measure \mathbb{P}^* under which insurance is priced and does *not* generally reflect the actual observed dependence between \mathbf{X} and \mathbf{D} .

- (b) For $\mathbf{Z} = \mathbf{X}$, the calculation that avoids indirect discrimination is based on the knowledge of $\mu^*(\mathbf{X}, \mathbf{D})$, see (7) – hence it *requires* collection of discriminatory policyholder information. In fact, one of the most critical problems in practice is that discriminatory information is often incomplete, e.g. about ethnicity, which typically results in indirect discrimination.
- (c) In statistical applications we usually use the conditional probability measure $\mathbb{P}(y | \mathbf{X}, \mathbf{D})$ to model the claim Y , given the covariates (\mathbf{X}, \mathbf{D}) . The reason for this choice is that Y , given (\mathbf{X}, \mathbf{D}) , is observed under the real world measure \mathbb{P} , which allows for direct estimation of the regression function, see also Section 5 below,

$$(\mathbf{x}, \mathbf{d}) \mapsto \mu(\mathbf{x}, \mathbf{d}).$$

We could even insist on the measure choice $\mathbb{P}(y | \mathbf{X}, \mathbf{D})$ for the claim Y , given (\mathbf{X}, \mathbf{D}) , because it implies that we maintain the (causal) structure of how the covariates impact the response. This then motivates the choice

$$\mathbb{P}^*(y, \mathbf{x}, \mathbf{d}) = \mathbb{P}(y | \mathbf{X} = \mathbf{x}, \mathbf{D} = \mathbf{d}) \mathbb{P}^*(\mathbf{x}) \mathbb{P}^*(\mathbf{d}),$$

for $\mathbf{Z} = \mathbf{X}$ in Definition 11, which in view of (7) results in the discrimination-free price

$$\mu^*(\mathbf{X}) = \int_{\mathbf{d}} \mu(\mathbf{X}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d} | \mathbf{X}) = \int_{\mathbf{d}} \mu(\mathbf{X}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d}) = h^*(\mathbf{X}).$$

Thus, the discrimination-free price of Definition 6 does neither allow for direct nor for indirect discrimination.

- (d) Linking to Remark 7(e), we need to know the real world measure \mathbb{P} in order to study unbiasedness w.r.t. $\mu = \mathbb{E}[Y]$. Moreover, the actual portfolio that we hold is described by $\mathbf{Z} \sim \mathbb{P}(\mathbf{z})$ which motivates the choice $\mathbb{P}^*(\mathbf{z}) = \mathbb{P}(\mathbf{z})$ for studying unbiasedness. Thus, unbiasedness under Definition 11 holds if

$$\mu = \mathbb{E}[\mu^*(\mathbf{Z})] = \int_{\mathbf{d}, \mathbf{z}} \mu^*(\mathbf{z}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d}) \, d\mathbb{P}(\mathbf{z}).$$

We come back to this in Section 4.

2.4. Existence of discrimination-free prices. We have not yet discussed existence of discrimination-free prices according to Definition 6 and the possibility of avoiding indirect discrimination according to Definition 11. This is done in the present section.

We emphasize that data (and the related statistical models) play a crucial role in discrimination-free considerations:

- Indirect discrimination may be the result of incomplete discriminatory information, see Remark 12(b).
- Indirect discrimination may be the result of nonexistent or insufficient information of certain parts of the population.

In this section we discuss the second item which can enter in different ways. A first one, which is probably less relevant in insurance, is that not all parts of the population are equally represented in the development of the statistical model. For instance, there is research in image recognition to discover malignant melanoma (skin cancer). If this research is mainly based on images of people with light skin, the corresponding model will likely fail to discover malignant melanoma for people with dark skin. This is a form of discrimination resulting from *insufficient data* of certain parts of the population.

In the current section we rather discuss *nonexistent data* of certain parts of the population. The meaning and implications of nonexistent data are going to be discussed in more detail. We start with an example. Assume that the discriminatory covariates \mathbf{D} correspond to gender and the non-discriminatory ones \mathbf{X} to education. Education could be in the ordinal form “secondary school degree”, “high school degree” or “university degree”, but information about education could also be received in the following categorical form “Catholic college degree”, “public college degree” or “girls college degree”. Per definition the last label “girls college degree” contains as only gender “female”. This implies that

$$\mathbb{P}(\mathbf{X} = \text{girls college degree}, \mathbf{D} = \text{man}) = 0,$$

thus, the event $A = \{\mathbf{X} = \text{girls college degree}, \mathbf{D} = \text{man}\} \in \mathcal{F}$ is a null set w.r.t. \mathbb{P} . In many cases, we do not model responses Y on null sets. Therefore, neither Y on A may be specified in our model nor the conditional expectation $\mu(\text{girls college degree}, \text{man}) = \mathbb{E}[Y \mid A]$ may be determined. But this implies that we cannot evaluate the discrimination-free price

$$h^*(\mathbf{X}) = \int_{\mathbf{d}} \mu(\mathbf{X}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d}),$$

if $\mathbb{P}^*(\mathbf{d})$ has positive probability mass on both genders. In the current situation, the problem may be solved by setting $\mathbb{P}^*(\mathbf{D} = \text{woman}) = 1$ which gives the discrimination-free price $h^*(\mathbf{X}) = \mu(\mathbf{X}, \text{woman})$.

If the education information \mathbf{X} has an additional level “boys college degree”, the above solution will not work because we have a second \mathbb{P} -null set $B = \{\mathbf{X} = \text{boys college degree}, \mathbf{D} = \text{woman}\} \in \mathcal{F}$ which makes it impossible to choose a distribution $\mathbb{P}^*(\mathbf{d})$ such that the discrimination-free price $h^*(\mathbf{X})$ is well-defined.

The simple solution to this problem is to drop the education information, that is, choose a smaller covariate set. This is equivalent to choosing a true subset \mathbf{Z} of \mathbf{X} in Definition 11. In practice, we often try to inter- or extrapolate the model assumptions for Y . This is reasonable if unavailable information corresponds to ordinary variables (and responses have some continuity in these covariates). In certain cases it may also be justified for nominal variables by, for example, postulating a multiplicative influence structure of covariates, say, women are $x\%$ better than men regardless of the attended college. In our situation such an assumption can be made, but it cannot be verified because of a missing control group.

Proposition 13. *Assume there exists a product measure $\mathbb{P}^*(\mathbf{x})\mathbb{P}^*(\mathbf{d})$ on (Ω, \mathcal{F}) which is absolutely continuous w.r.t. the probability measure $\mathbb{P}(\mathbf{x}, \mathbf{d})$ of the covariates (\mathbf{X}, \mathbf{D}) . Then, there exists a price $\mu^*(\mathbf{X})$ that avoids indirect discrimination.*

Proof. Absolute continuity implies that every $\mathbb{P}(\mathbf{x}, \mathbf{d})$ -null set is also a $\mathbb{P}^*(\mathbf{x})\mathbb{P}^*(\mathbf{d})$ -null set. Therefore, $\mu(\mathbf{X}, \mathbf{D})$ is well-defined on all sets where (\mathbf{X}, \mathbf{D}) has positive $\mathbb{P}^*(\mathbf{x})\mathbb{P}^*(\mathbf{d})$ -probability mass. Since the latter is a product measure we can calculate the discrimination-free price $h^*(\mathbf{X})$ by integrating $\mu(\mathbf{X}, \mathbf{d})$ over $d\mathbb{P}^*(\mathbf{d} \mid \mathbf{X}) = d\mathbb{P}^*(\mathbf{d})$, see also (7). This completes the proof. \square

3. CAUSAL INFERENCE AND DISCRIMINATION

The purpose of this section is to justify the discrimination-free price of Definition 6 in a causal inference setting. We try to give these arguments in a pedagogical

(though slightly informal) way. For a rigorous treatment we refer to Pearl [16] and Pearl et al. [17, Ch. 3.1].

The starting point of causal inference is a hypothesis of a potential covariate relationship which may be described in terms of a directed graph \mathfrak{G} . The directed graph \mathfrak{G} consists of a set of *nodes* corresponding to the characteristics, including the response Y , and *directed edges* – “arrows” – indicating directions of potential influence between the characteristics (including the response Y). This informal definition can be made precise, but it is most easily understood by the example given in Figure 1 (lhs).



FIGURE 1. (lhs) Causal diagram described by \mathfrak{G} ; (rhs) causal diagram altered according to the do-operation $\text{do}(\mathbf{X} = \mathbf{x})$.

The directed graph \mathfrak{G} in Figure 1 (lhs) is an example of a directed acyclic graph (DAG), where acyclic means that when following the direction of the edges, the graph does not contain any loops. For a precise description, see Pearl et al. [17, Ch. 1.4]. The graphical representation given in Figure 1 (lhs) corresponds to a situation where the discriminatory characteristics \mathbf{D} may influence both directly on Y , but also indirectly via \mathbf{X} . Hence, Figure 1 (lhs) tells us that \mathbf{D} is a *confounder* w.r.t. the effect of \mathbf{X} on Y .

Figure 1 (lhs) already captures a large number of realistic insurance pricing situations. For instance, in view of Example 1, we may identify smoking habits by \mathbf{X} and the gender by the discriminatory factors \mathbf{D} . Differences in smoking habits between men and women can be expressed by a directed edge $\mathbf{D} \rightarrow \mathbf{X}$. There are intrinsic differences between men and women when it comes to health outcomes (e.g. women cannot get prostate cancer), this is described by $\mathbf{D} \rightarrow Y$. Moreover, smoking in itself may cause health problems $\mathbf{X} \rightarrow Y$. This situation is exactly covered by the DAG \mathfrak{G} in Figure 1 (lhs).

We remark that the subsequent causality arguments also hold true for more general causal diagrams, for instance, if we have additional (unmeasured) confounding characteristics \mathbf{U} that enter the causal diagram as illustrated in Figures 10-12, given in the appendix. For pedagogical reasons, we restrict ourselves to the causal diagram in Figure 1 (lhs).

Since the directed edges in the DAG \mathfrak{G} do not act fully deterministically, and to make probabilistic statements, we need to endow the DAG \mathfrak{G} with a probability measure \mathbb{P} that describes the randomness involved. This probability measure \mathbb{P} should be Markovian on \mathfrak{G} which, colloquially speaking, means that all nodes have independent noisy background variables, see Pearl et al. [17, Ch. 3.2.1]; for noisy proxy variables we also refer to Kuroki and Pearl [14]. This Markovian DAG $(\mathfrak{G}, \mathbb{P})$ leads in our example of Figure 1 to the following factorization, see (3.5)-(3.6) in

[17],³

$$(8) \quad \begin{aligned} \mathbb{P}(Y, \mathbf{X}, \mathbf{D}) &= \mathbb{P}(Y \mid \mathbf{X}, \mathbf{D}) \mathbb{P}(\mathbf{X}, \mathbf{D}) \\ &= \mathbb{P}(Y \mid \mathbf{X}, \mathbf{D}) \mathbb{P}(\mathbf{X} \mid \mathbf{D}) \mathbb{P}(\mathbf{D}). \end{aligned}$$

With this setup in place, one way to approach non-discriminatory pricing is to require the following:

“Given that a policyholder has the set of characteristics $\mathbf{X} = \mathbf{x}$, what is the expected value of Y , after removing the possible confounding effects of discriminatory covariates \mathbf{D} ?”

The idea of the causal (non-confounded) effect⁴ of \mathbf{X} on Y can be captured precisely by a modification of the DAG \mathfrak{G} given in Figure 1. Specifically, as illustrated on the rhs of Figure 1, remove all directed edges to \mathbf{X} and set the value of \mathbf{X} to \mathbf{x} . The removal of the directed edge $\mathbf{D} \rightarrow \mathbf{X}$ allows us to consider only the (direct) causal effect of $\mathbf{X} = \mathbf{x}$ on Y . Note that this operation is different to conditioning. When conditioning on $\mathbf{X} = \mathbf{x}$, the distribution of \mathbf{D} is generally affected; but in the modified graph on the rhs of Figure 1, changes in \mathbf{x} do not influence \mathbf{D} . This is precisely the desired effect of removing the ability of inferring discriminatory covariates from non-discriminatory ones, as was discussed in Remark 12(a) in Section 2. The above intervention of removing all directed edges to \mathbf{X} and of fixing $\mathbf{X} = \mathbf{x}$ is denoted by the so-called do-operator $\text{do}(\mathbf{X} = \mathbf{x})$, see Pearl et al. [17, Ch. 3.2.1]. Henceforth, a price that takes into account only the causal (non-confounded) effect of \mathbf{X} on Y can be defined by

$$(9) \quad \mathbb{E}[Y \mid \text{do}(\mathbf{X} = \mathbf{x})],$$

where the probability $\mathbb{P}(Y \mid \text{do}(\mathbf{X} = \mathbf{x}))$ still needs to be defined.

The do-operation $\text{do}(\mathbf{X} = \mathbf{x})$ going from the lhs to the rhs in the DAG \mathfrak{G} of Figure 1 fulfills the so-called back-door criterion relative to the ordered pair (\mathbf{X}, Y) , see Definition 3.3.1 in Pearl et al. [17]. This means that the do-operation $\text{do}(\mathbf{X} = \mathbf{x})$ blocks every path between \mathbf{X} and Y that contains an arrow into \mathbf{X} . Under these properties Theorem 3.3.2 (back-door adjustment) of Pearl et al. [17] states that the causal effect of \mathbf{X} on Y is identifiable in the Markovian DAG $(\mathfrak{G}, \mathbb{P})$ and given by the formula

$$(10) \quad \mathbb{P}(Y \mid \text{do}(\mathbf{X} = \mathbf{x})) = \int_{\mathbf{d}} \mathbb{P}(Y \mid \mathbf{X} = \mathbf{x}, \mathbf{D} = \mathbf{d}) \, d\mathbb{P}(\mathbf{d}).$$

The following proposition is an easy consequence of definition (9), the definition of the best-estimate price $\mu(\mathbf{X}, \mathbf{D})$ and formula (10).

Proposition 14. *Assume that the do-operation $\text{do}(\mathbf{X} = \mathbf{x})$ fulfills the back-door criterion relative to the ordered pair (\mathbf{X}, Y) in the Markovian DAG $(\mathfrak{G}, \mathbb{P})$. It holds that*

$$\mathbb{E}[Y \mid \text{do}(\mathbf{X} = \mathbf{x})] = \int_{\mathbf{d}} \mu(\mathbf{x}, \mathbf{d}) \, d\mathbb{P}(\mathbf{d}) = h^{(CI)}(\mathbf{x}),$$

³Formula (8) is nothing else than the Bayes' rule on the Markovian DAG which more formally may either be understood in terms of probabilities of events or in terms of conditional distributions.

⁴For our intended application, the interpretation of “causal” will here be restricted to the sub-population which chooses to buy insurance cover. No explicit consideration will be taken w.r.t. which characteristics that will make individuals belong to this sub-population. However, it is often reasonable to assume that the underlying causal relations should be the same for the general population and the insured sub-population, although the effects are not the same.

where $h^{(CI)}(\mathbf{x})$ is defined by (6).

Remark 15.

- (a) Proposition 14 justifies the discrimination-free price $h^{(CI)}(\mathbf{X})$ of equation (6) under specific Markovian DAG assumptions, and motivates the choice $\mathbb{P}^*(\mathbf{d}) = \mathbb{P}(\mathbf{d})$ in Definition 6. In particular, this discrimination-free price is obtained by adjusting for possible confounding of \mathbf{D} w.r.t. \mathbf{X} , thus, only measuring the direct causal effect of \mathbf{X} on Y .
- (b) Proposition 14 does not say anything about the quality of the induced price which, of course, will depend on the specific model being used – it merely states that there is no confounding, given that the assumptions of the corollary are fulfilled.
- (c) It is possible to extend the covariate relations described by Figure 1 by including unmeasured characteristics \mathbf{U} , see Figures 10 and 11, below. Furthermore, recall that the ambition is *not* to obtain an as accurate attribution of the effect of \mathbf{X} as possible (this would be the unawareness price of Definition 4), but the focus rather is on ascertaining that the resulting price is devoid of any influence from discriminatory characteristics. Thus, if it can be argued that the covariate relations described in Figure 11 (or any subgraph thereof) can be mapped into an actually observable graph corresponding to that of Figure 12 (or any subgraph thereof), then Proposition 14 applies again. That is, pricing based on (6) in the situation described in Figure 12 will be discrimination-free, but the quality of the induced price may be poor because we do not account for differences described by \mathbf{U} .

4. ATTRIBUTION OF TOTAL PORTFOLIO PREMIUM TO INDIVIDUAL POLICIES

The difficulty that we still have to deal with is that, in general, a discrimination-free price has a bias, see (5). This bias needs to be corrected because otherwise the premium for the entire portfolio may not be at the appropriate level. There is no canonical way of correcting for this potential bias, the only requirement is that the bias correction should be discrimination-free too, which excludes complex cost allocation mechanisms.

The portfolio bias of the \mathbb{P}^* -discrimination-free price is defined by

$$B^* := \mu - \mathbb{E}[h^*(\mathbf{X})] = \mathbb{E}[Y] - \int_{\mathbf{x}, \mathbf{d}} \mu(\mathbf{x}, \mathbf{d}) \, d\mathbb{P}^*(\mathbf{d})d\mathbb{P}(\mathbf{x}).$$

Simple bias corrections arise from taking rather different positions. An egalitarian position is taken by distributing the portfolio bias B^* uniformly across the entire portfolio, regardless of any non-discriminatory covariates \mathbf{X} . This motivates the *uniformly adjusted \mathbb{P}^* -discrimination-free price* defined by

$$(11) \quad \pi^{*,u}(\mathbf{X}) := h^*(\mathbf{X}) + B^*.$$

Moreover, if we do not consider any covariates (neither discriminatory nor non-discriminatory ones) we are back in the situation of a homogeneous situation where we charge the same (constant) premium μ to every policyholder. A drawback of the uniformly adjusted price (11) is that it may result in negative prices for certain covariate values \mathbf{X} .

A different position is to allocate the bias B^* by differentiating w.r.t. \mathbf{X} in a still discrimination-free fashion (avoiding any inference of \mathbf{D} from \mathbf{X}). A natural

way is to allocate the total premium proportionally to $h^*(\mathbf{X})$, resulting in the *proportionally adjusted \mathbb{P}^* -discrimination-free price*

$$(12) \quad \pi^{*,P}(\mathbf{X}) := h^*(\mathbf{X}) \frac{\mu}{\mu - B^*}.$$

In the remainder of this section we discuss a more sophisticated approach which chooses the distribution $\mathbb{P}^*(\mathbf{d})$ specifically such that the discrimination-free price $h^*(\mathbf{X})$ is unbiased, i.e. $B^* = 0$. As there may be many such distributions that may satisfy this condition, an additional criterion for choosing $\mathbb{P}^*(\mathbf{d})$ is needed. A standard criterion is to choose the measure \mathbb{P}^* as close as possible to the physical distribution $\mathbb{P}(\mathbf{d})$, subject to the unbiasedness constraint on $h^*(\mathbf{X})$. If the relative entropy (Kullback–Leibler divergence) is chosen as the divergence measure, then standard results can be applied, see Csiszár [6]. Specifically, note that the unbiasedness condition $\mathbb{E}[h^*(\mathbf{X})] = \mu$ is equivalent to

$$(13) \quad \mathbb{E}^*[\zeta(\mathbf{D})] = \mu,$$

where we define

$$\zeta(\mathbf{t}) := \mathbb{E}[\mu(\mathbf{X}, \mathbf{t})].$$

Note, we assume existence of distributions $\mathbb{P}^*(\mathbf{d})$ that fulfill (13) which in view of Section 2.4 needs to be made. The distribution $\mathbb{P}^*(\mathbf{d})$ that minimizes the relative entropy with respect to $\mathbb{P}(\mathbf{d})$ subject to⁵ (13) is given by

$$\mathbb{P}^*(\mathbf{d}) = \mathbb{E} \left[\mathbf{1}_{\{\mathbf{D} \leq \mathbf{d}\}} \frac{e^{\beta \zeta(\mathbf{D})}}{\mathbb{E}[e^{\beta \zeta(\mathbf{D})}]} \right],$$

for a suitably chosen parameter β . Hence, the premium for a policyholder with non-discriminatory covariates $\mathbf{X} = \mathbf{x}$ is defined by (subject to existence)

$$\pi^{*,KL}(\mathbf{x}) := h^*(\mathbf{x}) = \mathbb{E} \left[\mu(\mathbf{x}, \mathbf{D}) \frac{e^{\beta \zeta(\mathbf{D})}}{\mathbb{E}[e^{\beta \zeta(\mathbf{D})}]} \right],$$

which in light of Remark 10(d) may be thought of in terms of $h^*(\mathbf{x}) = \mathbb{E}[Y^* | \mathbf{X} = \mathbf{x}]$. To ease interpretation of this formula, and in order to make the connection to Remark 10(d) more clear, let $\mathbf{D} = D$ be one-dimensional and $\mu(\mathbf{x}, d) \geq 0$ be increasing in d . Then, for $\beta > 0$, we have

$$\begin{aligned} \pi^{*,KL}(\mathbf{x}) &= \mathbb{E} \left[\mu(\mathbf{x}, D) \frac{e^{\beta \zeta(D)}}{\mathbb{E}[e^{\beta \zeta(D)}]} \right] \\ &= \mathbb{E}[\mu(\mathbf{x}, D)] + \text{Cov} \left[\mu(\mathbf{x}, D), \frac{e^{\beta \zeta(D)}}{\mathbb{E}[e^{\beta \zeta(D)}]} \right] \\ &\geq \mathbb{E}[\mu(\mathbf{x}, D)] = h^{(CI)}(\mathbf{x}), \end{aligned}$$

which corresponds to the situation where the choice $\mathbb{P}^* = \mathbb{P}$ would produce a negative bias (under-pricing). The calculation of $\pi^{*,KL}(\mathbf{x})$ assigns a higher premium to policyholders with covariates $\mathbf{X} = \mathbf{x}$ such that $\mu(\mathbf{x}, D)$ is more volatile (this can be made rigorous but is best visible in approximation (14), below). This represents policies for which lack of information on discriminatory covariates matters more, in the sense that there is a higher sensitivity to the uncertainty induced by not using

⁵Since relative entropy is not symmetric, for the sake of completeness and to avoid ambiguity, following the notation of Csiszár [6], the relative entropy of $\mathbb{P}^*(\mathbf{d})$ relative $\mathbb{P}(\mathbf{d})$ is denoted $I(\mathbb{P}^*(\mathbf{d}) || \mathbb{P}(\mathbf{d}))$.

the discriminatory factor D . One can thus view the bias correction in $\pi^{*,KL}(\mathbf{x})$ as an implicit discrimination-free) risk load.

For small β we have approximation

$$(14) \quad \begin{aligned} h^*(\mathbf{x}) &\approx \mathbb{E}[\mu(\mathbf{x}, \mathbf{D})] + \beta \text{Cov}[\mu(\mathbf{x}, \mathbf{D}), \zeta(\mathbf{D})] \\ &= \mathbb{E}[\mu(\mathbf{x}, \mathbf{D})] + \beta \sqrt{\text{Var}[\mu(\mathbf{x}, \mathbf{D})] \text{Var}[\zeta(\mathbf{D})]} \text{Corr}[\mu(\mathbf{x}, \mathbf{D}), \zeta(\mathbf{D})]. \end{aligned}$$

5. ESTIMATED MODEL

All previous discussion and derivations of discrimination-free prices and indirect discrimination were conducted under the assumption that the “true” probabilistic model underlying the portfolio $(Y, \mathbf{X}, \mathbf{D})$ is known, represented by the physical measure \mathbb{P} . In practice, an *estimated model* is used because, typically, the data generating mechanism is unknown.

Specifically, one starts from data

$$\mathcal{S} = \{(y_1, \mathbf{x}_1, \mathbf{d}_1), \dots, (y_n, \mathbf{x}_n, \mathbf{d}_n)\},$$

assuming that $(y_i, \mathbf{x}_i, \mathbf{d}_i)$ are i.i.d. realisations of $(Y, \mathbf{X}, \mathbf{D}) \sim \mathbb{P}$.

Subsequently, a regression model (in the broader sense of regression models) is chosen

$$(15) \quad \hat{\mu} : (\mathbf{x}, \mathbf{d}) \mapsto \hat{\mu}(\mathbf{x}, \mathbf{d}) = \hat{\mu}(\mathbf{x}, \mathbf{d}; \boldsymbol{\theta}),$$

which typically differs from the (true) best-estimate price functional $(\mathbf{x}, \mathbf{d}) \mapsto \mu(\mathbf{x}, \mathbf{d})$, given in Definition 2, but which should mimic $\mu(\mathbf{x}, \mathbf{d})$ in the best possible way. One may specify a fixed functional form for $\hat{\mu}$ in (15) or, in a wider sense, one can specify an algorithm that generates the mapping (15) from the data \mathcal{S} . In either case, $\hat{\mu}$ will still depend on unknown parameters $\boldsymbol{\theta}$ that have to be estimated from the data \mathcal{S} (using a given objective function) yielding estimate $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}(\mathcal{S})$.

The resulting \mathcal{S} -calibrated regression function

$$(16) \quad (\mathbf{x}, \mathbf{d}) \mapsto \hat{\mu}(\mathbf{x}, \mathbf{d}; \hat{\boldsymbol{\theta}})$$

then provides the approximation to the best-estimate price functional $(\mathbf{x}, \mathbf{d}) \mapsto \mu(\mathbf{x}, \mathbf{d})$. Note that (16) provides an estimate of the best price and, obviously, this estimate is in general discriminatory because it explicitly considers the discriminatory covariate values \mathbf{d} . Moreover, since we use the data \mathcal{S} which has been generated under the physical measure \mathbb{P} , the regression function (16) also needs to be understood under the physical measure \mathbb{P} , we refer to Remark 12(c).

The unawareness price functional $\mathbf{x} \mapsto \mu(\mathbf{x})$ can be approximated in an analogous manner by just dropping \mathbf{d} in (15) and (16), resulting in an estimated regression function

$$(17) \quad \mathbf{x} \mapsto \bar{\mu}(\mathbf{x}; \hat{\boldsymbol{\vartheta}}),$$

where the functional forms $\hat{\mu}$ and $\bar{\mu}$ may differ as well as their parameters $\boldsymbol{\theta}$ and $\boldsymbol{\vartheta}$, respectively. We emphasize that typically $\bar{\mu}(\cdot; \hat{\boldsymbol{\vartheta}})$ will indirectly discriminate w.r.t. \mathbf{d} because in the estimation process of $\hat{\boldsymbol{\vartheta}}$ we implicitly use the covariate combinations $(\mathbf{x}_i, \mathbf{d}_i)$ which (empirically) contain the dependencies $\mathbb{P}(\mathbf{d} \mid \mathbf{x})$ that

allow inference of \mathbf{D} from \mathbf{X} . The estimated unawareness price $\bar{\mu}(\mathbf{x}; \hat{\boldsymbol{\theta}})$ can also be interpreted as an approximation to

$$\mathbb{E}[\hat{\mu}(\mathbf{X}, \mathbf{D}; \hat{\boldsymbol{\theta}}) \mid \mathbf{X} = \mathbf{x}; \mathcal{S}],$$

using the tower property argument (under the physical measure \mathbb{P}).

Typically, also $\mathbb{P}(\mathbf{d})$ is not known. Usually, it is estimated by the empirical distribution $(n_{\mathbf{d}}/n)_{\mathbf{d}}$, where $n_{\mathbf{d}}/n$ is the observed relative frequency of the discriminatory covariate \mathbf{d} in \mathcal{S} . This motivates the discrimination-free price

$$(18) \quad \hat{h}^{(CI)}(\mathbf{x}) = \sum_{\mathbf{d}} \hat{\mu}(\mathbf{x}, \mathbf{d}; \hat{\boldsymbol{\theta}}) \frac{n_{\mathbf{d}}}{n},$$

where we use the estimated best-estimate price functional (16). The price (18) is discrimination-free in the sense of Definition 6, i.e. the discrimination-free property is not affected by the fact that we work with an estimated model. While potential estimation error may result in prices $\hat{h}^{(CI)}(\mathbf{x})$ that are not very close to $h^{(CI)}(\mathbf{x})$, the property of discrimination-freeness is preserved within the selected model; we explore this in more detail in Section 6. When choosing the structure of the regression function $\hat{\mu}$ in (15), we typically require existence of the discrimination-free price (18) in the sense of Proposition 13.

6. NUMERICAL ILLUSTRATION

6.1. Model and alternative pricing rules. We present a simple health insurance example, demonstrating our proposed approach to discrimination-free pricing. The example we present satisfies the causal relations of Figure 1, such that discrimination-free prices can be understood as reflecting direct (unconfounded) causal effects (in an insured sub-population).

Let $\mathbf{D} = D$ correspond to the single discriminatory characteristic “gender”, that is, $D \in \{\text{woman}, \text{man}\}$. Furthermore, let $\mathbf{X} = (X_1, X_2)'$, where $X_1 \in \{15, \dots, 80\}$ denotes the age of the policyholder, and $X_2 \in \{\text{non-smoker}, \text{smoker}\}$; below we assume that smoking habits are gender related. We consider three different types of health costs: birthing related health costs only affecting women between ages 20 and 40 (type 1), cancer related health costs with a higher frequency for smokers and also for women (type 2), and health costs due to other disabilities (type 3). For simplicity, we only consider claim counts, assuming deterministic claim costs for the three different claim types. Moreover, we model all individuals as independent, having the same exposure ($= 1$). We assume that the claim counts for the different claim types are described by independent Poisson GLMs with canonical (i.e. log-) link function. The three different types of claims are governed by the following log-frequencies

$$(19) \quad \log \lambda_1(\mathbf{X}, D) := \alpha_0 + \alpha_1 1_{\{X_1 \in [20, 40]\}} 1_{\{D = \text{woman}\}},$$

$$(20) \quad \log \lambda_2(\mathbf{X}, D) := \beta_0 + \beta_1 X_1 + \beta_2 1_{\{X_2 = \text{smoker}\}} + \beta_3 1_{\{D = \text{woman}\}},$$

$$(21) \quad \log \lambda_3(\mathbf{X}, D) := \gamma_0 + \gamma_1 X_1,$$

based on the joint non-discriminatory and discriminatory covariates (\mathbf{X}, D) . The deterministic claims costs of the different claim types are given by $(c_1, c_2, c_3) = (0.5, 0.9, 0.1)$ for claims of type 1, type 2, and type 3, respectively.

The best-estimate price (considering all covariates) of Definition 2 is given by

$$\mu(\mathbf{X}, D) = c_1 \lambda_1(\mathbf{X}, D) + c_2 \lambda_2(\mathbf{X}, D) + c_3 \lambda_3(\mathbf{X}, D).$$

This best-estimate price is illustrated in Figure 2 for the parameter values $(\alpha_0, \alpha_1) = (-40, 38.5)$, $(\beta_0, \beta_1, \beta_2, \beta_3) = (-2, 0.004, 0.1, 0.2)$, and $(\gamma_0, \gamma_1) = (-2, 0.01)$. The plots on the lhs of Figure 2 refer to smokers ($X_2 = \text{smoker}$), while those on the rhs to non-smokers ($X_2 = \text{non-smoker}$). The solid black lines give the best-estimate prices $\mu(\mathbf{X}, D)$ for women and the solid red lines for men. Obviously, these best-estimate prices discriminate between genders.

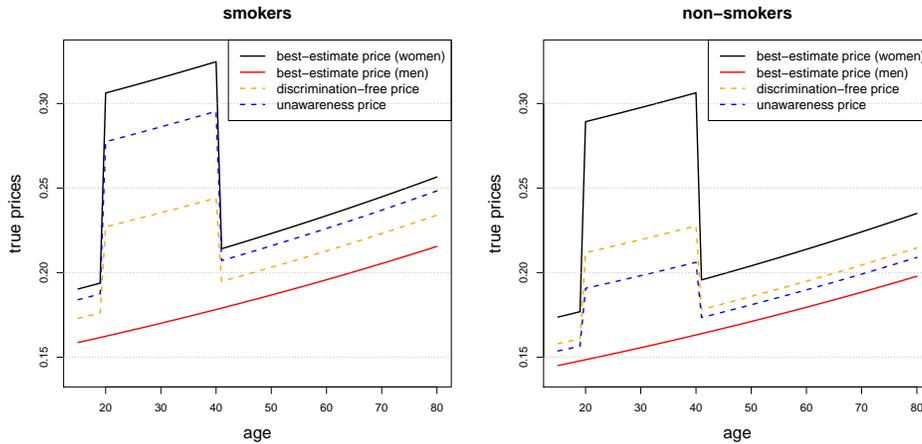


FIGURE 2. True model: (lhs) smokers and (rhs) non-smokers with solid black and red lines giving the best-estimate prices for women and men, respectively. The dotted orange lines show the discrimination-free prices and the dotted blue lines show the unawareness prices.

Next, we calculate the discrimination-free price defined in Definition 6 for $\mathbb{P}^*(d) = \mathbb{P}(d)$, see (6), motivated by Proposition 14. It is given by

$$h^{(CI)}(\mathbf{X}) = \sum_{d \in \{\text{woman, man}\}} (c_1 \lambda_1(\mathbf{X}, d) + c_2 \lambda_2(\mathbf{X}, d) + c_3 \lambda_3(\mathbf{X}, d)) \mathbb{P}(D = d).$$

For the calculation of this discrimination-free price we need the gender proportions within our population. We set $\mathbb{P}(D = \text{woman}) = 0.45$. The orange dotted lines in Figure 2 provide the resulting discrimination-free prices for smokers and non-smokers. Note that these are identical for men and women, such that all price differences can be described solely by different ages X_1 and smoking habits X_2 , irrespective of gender D . Moreover, the smoking habits do not allow us to infer the gender; note that in the exposition so far, it has not been necessary to describe how smoking habits vary by gender.

We compare this discrimination-free price to the unawareness price obtained by simply dropping the gender covariate D from the calculations (Definition 4). Thus, we calculate

$$\mu(\mathbf{X}) = c_1 \mathbb{E}[\lambda_1(\mathbf{X}, D) \mid \mathbf{X}] + c_2 \mathbb{E}[\lambda_2(\mathbf{X}, D) \mid \mathbf{X}] + c_3 \mathbb{E}[\lambda_3(\mathbf{X}, D) \mid \mathbf{X}]$$

The unawareness price requires additional information about the following conditional probabilities

$$(22) \quad \mathbb{P}(D = d \mid \mathbf{X} = \mathbf{x}) = \frac{\mathbb{P}(D = d, \mathbf{X} = \mathbf{x})}{\mathbb{P}(\mathbf{X} = \mathbf{x})} = \frac{\mathbb{P}(D = d, X_2 = x_2)}{\mathbb{P}(X_2 = x_2)},$$

the last equality following from assuming that the age variable X_1 is independent from the random vector (X_2, D) . In addition, we set $\mathbb{P}(D = \text{woman} \mid X_2 = \text{smoker}) = 0.8$ and $\mathbb{P}(X_2 = \text{smoker}) = 0.3$. The former assumption tells us that smokers are more likely women; this is similar to Example 1. As a consequence, X_2 has explanatory power to predict the gender D , and the unawareness price will therefore be indirectly discriminatory against women. These unawareness prices are illustrated by the blue dotted lines in Figure 2. The blue dotted line lies above the discrimination-free price (orange) for smokers (Figure 2, lhs) and below for non-smokers (rhs). Thus, the unawareness price implicitly allocates a higher price to women because smokers are more likely women.

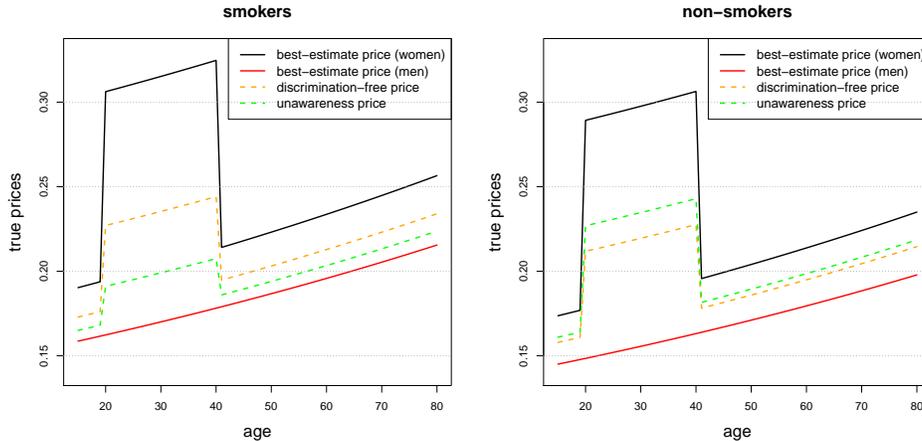


FIGURE 3. True model: (lhs) smokers and (rhs) non-smokers with solid black and red lines giving the best-estimate prices for women and men, respectively. The dotted orange lines show the discrimination-free prices and with dotted green lines show the unawareness prices, for an alternative assumption on $\mathbb{P}(D = \text{woman} \mid \text{smoker})$.

The latter indirect gender discrimination is easily verified by an alternative assumption, namely, that smokers are more likely men, say, $\mathbb{P}(D = \text{woman} \mid X_2 = \text{smoker}) = 0.2$. The resulting unawareness prices are plotted by the dotted green lines in Figure 3. We observe that the smokers are below the discrimination-free price (orange dotted line), and for non-smokers we have the opposite sign. That is, in this case women are again indirectly discriminated through their (non-)smoking habits, again serving as a proxy for the explanatory variable of gender. The break-even point is $\mathbb{P}(D = \text{woman} \mid X_2 = \text{smoker}) = 0.45 = \mathbb{P}(D = \text{woman})$ because in this case D and X_2 are independent, which prevents from indirect discrimination, and the unawareness price and the distribution-free price are equal.

6.2. Application on estimated models. The previous discussion has been based on the knowledge of the model generating the data. We now address the more realistic situation where the model needs to be estimated. To this effect, we simulate data from $(\mathbf{X}, D, Y) \sim \mathbb{P}$ consistently with the given model assumptions, and subsequently calibrate a neural network regression model to the simulated data.

Specifically, we choose a health insurance portfolio of size $n = 100,000$, and simulate claim counts from the Poisson GLMs (19), (20), and (21), with the choice $\mathbb{P}(D = \text{woman} \mid X_2 = \text{smoker}) = 0.8$. An age distribution for X_1 is also needed for the simulation – the chosen probability weights are shown in Figure 4. We assume that age X_1 is independent from gender D and smoking habits X_2 , as in (22).

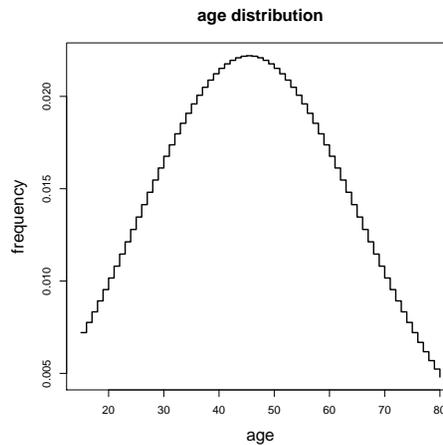


FIGURE 4. The age frequency used for both genders and smoking habits to simulate the data.

LISTING 1. Simulated health insurance data.

```

1 'data.frame':  100000 obs. of  6 variables:
2 $ N1: int  0 0 0 0 0 0 0 0 0 0 ...
3 $ N2: int  0 0 1 0 0 1 0 0 2 0 ...
4 $ N3: int  0 1 0 0 1 0 0 0 0 0 ...
5 $ X1: num  36 57 70 49 63 27 41 58 16 34 ...
6 $ X2: num  0 0 1 0 0 1 0 0 1 1 ...
7 $ D : num  0 1 1 0 0 1 0 0 1 1 ...

```

Listing 1 gives an excerpt of the simulated data. We have the three covariates X_1 (age), X_2 (smoking habit) and D (gender) on lines 5-7, and lines 2-4 illustrate the number of claims N_1 , N_2 and N_3 , separated by claim types. The proportion of women in this simulated data is 0.4505 which is close to the true value of $\mathbb{P}(D = \text{woman}) = 0.45$. Our first aim is to fit a regression model to this data, under the assumptions that individual policies are independent, and that the different claim types are independent and Poisson distributed. Beside this, we do not make any structural assumption about the regression functions, but we try to infer them from the data using neural networks. The independence assumption between the

claim counts N_1 , N_2 and N_3 motivates modeling them separately. Thus, we will fit three different neural networks to model λ_1 , λ_2 and λ_3 , respectively. As we do not use prior knowledge on the data generating process, we will feed all covariates (X_1, X_2, D) to each of the three networks.

LISTING 2. Neural network architecture used to infer λ_1 , λ_2 and λ_3 .

```

1 Design <- layer_input(shape = c(3), dtype = 'float32', name = 'Design')
2 #
3 Network = Design %>%
4   layer_dense(units=15, activation='relu', name='hidden1') %>%
5   layer_dense(units=15, activation='relu', name='hidden2') %>%
6   layer_dense(units=1, activation='exponential', name='Network')
7 #
8 model <- keras_model(inputs = c(Design), outputs = c(Network))
9 model %>% compile(loss = 'poisson', optimizer = 'adam')

```

Listing 2 illustrates the chosen neural network architecture, using the R library `keras`, with which the three regression functions (19)-(21) are estimated. We choose neural networks of depth 2 having 15 neurons in both hidden layers, the rectified linear unit (ReLU) activation function, and the canonical link under the Poisson assumption. Moreover, we select the Poisson deviance loss as our objective function. This network involves 316 weights that need to be calibrated. We train these weights of the three networks over 1000 epochs on batches of size 20,000.

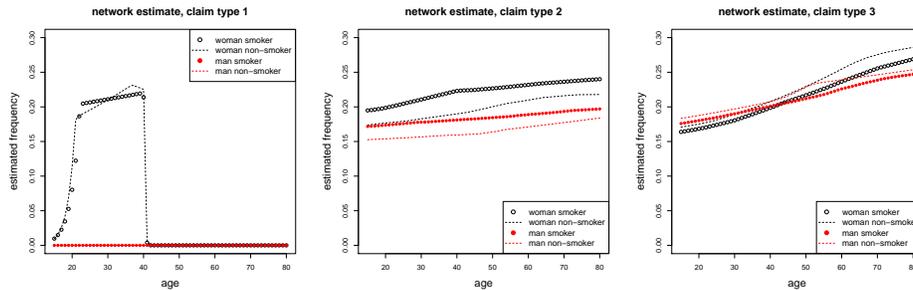


FIGURE 5. Estimated regression functions $\hat{\lambda}_1(\mathbf{X}, D)$ (lhs), $\hat{\lambda}_2(\mathbf{X}, D)$ (middle) and $\hat{\lambda}_3(\mathbf{X}, D)$ (rhs) using the neural network architecture of Listing 2.

Figure 5 illustrates the estimates $\hat{\lambda}_1(\mathbf{X}, D)$, $\hat{\lambda}_2(\mathbf{X}, D)$ and $\hat{\lambda}_3(\mathbf{X}, D)$ of the three regression functions (19), (20) and (21) respectively, obtained by fitting the three neural networks. The lhs of that figure gives claim type 1 which is birthing related. We see a rather accurate shape, with smoking habits correctly ignored, and men not affected by these claims. Figure 5 (middle) gives the cancer related frequencies. Also here we receive the same order w.r.t. gender and smoking habits as in (20). Finally, the rhs illustrates all remaining claims. As, by (21) claims frequencies should not depend on gender and smoking habits, the variation between lines indicates that the regression model captures a spurious effect.

Using these estimated frequencies, we calculate the estimated best-estimate price, (16),

$$\hat{\mu}(\mathbf{X}, D; \hat{\theta}) = c_1 \hat{\lambda}_1(\mathbf{X}, D) + c_2 \hat{\lambda}_2(\mathbf{X}, D) + c_3 \hat{\lambda}_3(\mathbf{X}, D),$$

and its discrimination-free counterpart (18),

$$\widehat{h}^{(CI)}(\mathbf{x}) = \sum_d \widehat{\mu}(\mathbf{x}, d; \widehat{\boldsymbol{\theta}}) \frac{n_d}{n},$$

with empirical proportions $n_{\text{woman}}/n = 1 - n_{\text{man}}/n = 0.4505$. These prices are illustrated in Figure 6: black lines give best-estimate prices for women, red lines for men, and with the orange dotted lines showing the discrimination-free counterparts. Comparing Figures 2 and 6 we conclude that the resulting true prices and estimated prices are rather similar. Of course, by construction the resulting discrimination-free price is gender neutral within the estimated model, and in our case close to the theoretical one.

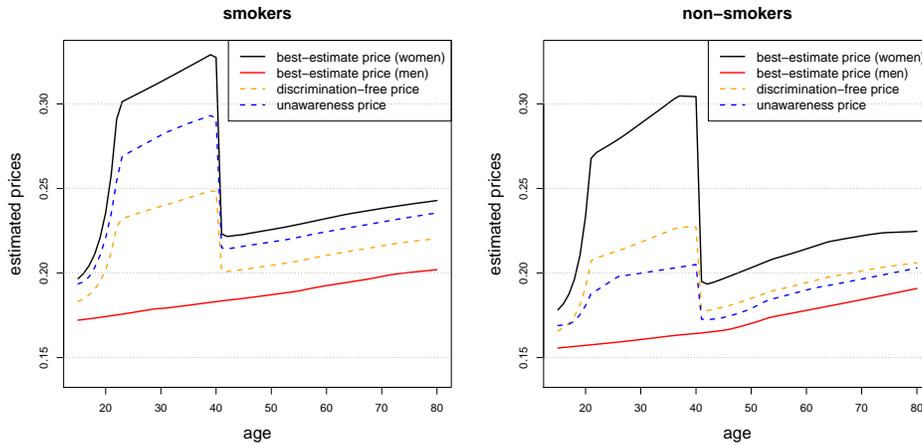


FIGURE 6. Estimated neural network model: (lhs) smokers and (rhs) non-smokers with solid black and red lines giving the best-estimate prices for women and men, respectively. The dotted orange lines show the discrimination-free prices and the dotted blue lines show the unawareness prices.

We indicate what happens if we drop the gender variable D from the very beginning, i.e. if we train the networks only on the covariates $\mathbf{X} = (X_1, X_2)$ as considered in (17). We choose exactly the same network architecture as in Listing 2 except that we modify the input dimension on line 1 from 3 for (\mathbf{X}, D) to 2 for \mathbf{X} . These networks involve 301 weights that need to be trained. The resulting estimated regression functions $\widehat{\lambda}_1(\mathbf{X})$, $\widehat{\lambda}_2(\mathbf{X})$ and $\widehat{\lambda}_3(\mathbf{X})$, ignoring gender information D , are illustrated in Figure 7. The left-hand side shows that we can no longer distinguish between gender, however, smokers are more heavily punished for birthing related costs, which is an undesired indirect discrimination effect against women because they are more often among the group of smokers (note that the y -scales in Figures 5 and 7 are the same). Finally, merging the different claim types provides the estimated unawareness prices (when first dropping D) as illustrated by the blue dotted lines in Figure 6, which can be compared with the blue dotted lines in Figure 2.

In our final analysis we illustrate that the (non-)discrimination property does not depend on the quality of the regression model (15) chosen. We choose a poor model

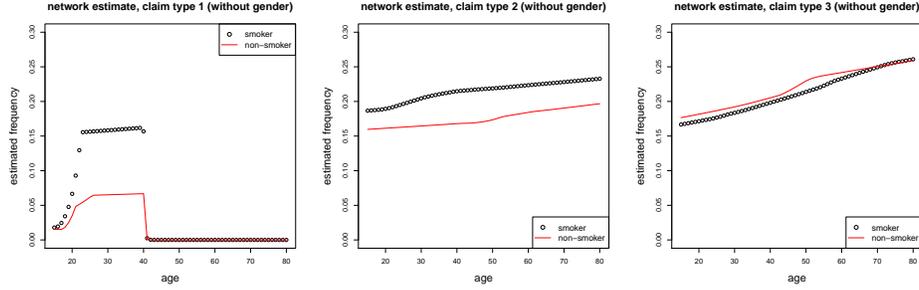


FIGURE 7. Estimated regression functions $\hat{\lambda}_1(\mathbf{X})$ (lhs), $\hat{\lambda}_2(\mathbf{X})$ (middle) and $\hat{\lambda}_3(\mathbf{X})$ (rhs) using neural networks and ignoring the gender information D .

(compared to the neural network above) by just assuming GLMs for $j = 1, 2, 3$

$$(23) \quad (\mathbf{x}, d) \mapsto \log \hat{\lambda}_j^{\text{GLM}}(\mathbf{x}, d) = \theta_0^{(j)} + \theta_1^{(j)} x_1 + \theta_2^{(j)} 1_{\{x_2=\text{smoker}\}} + \theta_3^{(j)} 1_{\{d=\text{woman}\}}.$$

This model will perform well for $j = 2, 3$, see (20)-(21), but it will perform poorly for $j = 1$, see (19). This is because such a model has difficulties capturing the highly non-linear birthing related effects, as seen in Figure 8 (lhs).

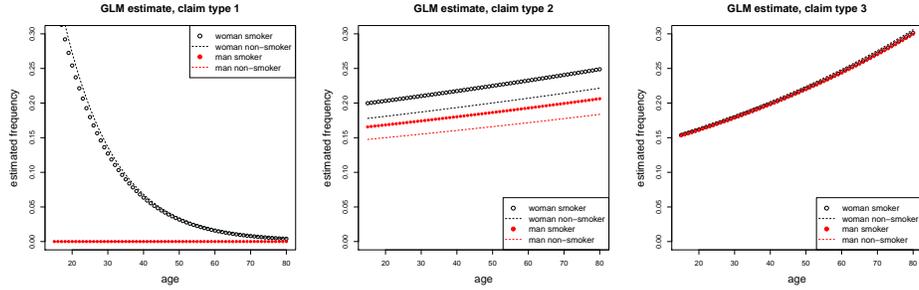


FIGURE 8. GLM estimated regression functions $\hat{\lambda}_1^{\text{GLM}}(\mathbf{X}, D)$ (lhs), $\hat{\lambda}_2^{\text{GLM}}(\mathbf{X}, D)$ (middle) and $\hat{\lambda}_3^{\text{GLM}}(\mathbf{X}, D)$ (rhs).

In Figure 9 we present the resulting best-estimate prices, unawareness prices, and discrimination-free prices (in orange dotted lines), as estimated using the GLM. The first observation is that the resulting prices are a poor approximation to the true prices of Figure 2, the latter assuming full knowledge of the true model. However, the general discrimination behavior is the same in both figures, namely, that the unawareness price discriminates indirectly by learning the gender D from smoking habits X_2 . This is illustrated by the relative positioning of blue and orange dotted lines, with smokers more heavily charged for birthing related costs due to the fact that smokers are more likely women.

The last step is to adjust for a potential bias, by considering the observed bias

$$\hat{B}^* := \frac{1}{n} \sum_{i=1}^n y_i - \frac{1}{n} \sum_{i=1}^n \hat{h}^{(CI)}(\mathbf{x}_i),$$

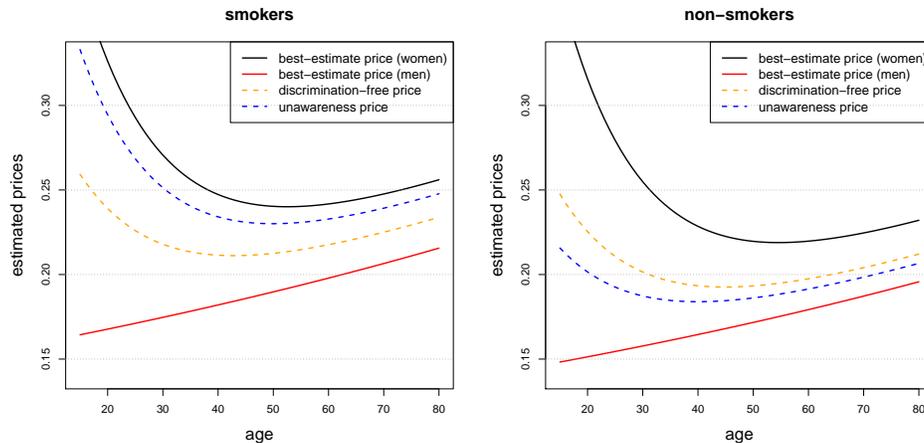


FIGURE 9. Estimated GLM: (lhs) smokers and (rhs) non-smokers with solid black and red lines giving the best-estimate prices for women and men, respectively. The dotted orange lines show the discrimination-free prices and the dotted blue lines show the unawareness prices.

and allocating it to the discrimination-free prices $\hat{h}^{(CI)}(\mathbf{x}_i)$, see Section 4. We refrain from doing this explicitly here. This finishes the example.

Remark. There is one issue that has not been considered so far, and which has been mentioned in the EU legislation [8], footnote (1) to Article 2.2(14) – life and health underwriting. Namely, we have implicitly assumed that the measurements of the non-discriminatory covariates are independent of the discriminatory characteristics. If we think of gender as a discriminatory covariate, this is not necessarily the case because, for instance, the waist to hip ratios naturally live on different scales for different genders, but they may still have the same impact on health related questions. This implies that non-discriminatory covariates may need pre-processing w.r.t. discriminatory ones, such that the resulting measurements for different discriminatory characteristics are comparable.

7. CONCLUDING REMARKS

The aim of this paper is to provide

- (a) an actuarial perspective of direct and indirect discrimination;
- (b) a demonstration that the omission of discriminatory information leads to indirect discrimination in prices;
- (c) a proposal for a simple formula that generates discrimination-free prices;
- (d) methods that ensure unbiasedness of discrimination-free prices on portfolio level, of course, these considerations also apply when transforming the actuarial tariff into a commercial one; and
- (e) a discussion on the role of the data in obtaining discrimination-free prices.

The initial starting point to this paper has been an actuarial one. We have intentionally avoided a deeper discussion on “fairness”, and, consequently, how

fairness may be measured. For more on these topics, we refer to Kusner et al. [15] and the references therein. Moreover, we have also not commented on which factors should be viewed as discriminatory – this is a societal decision that goes way beyond our actuarial discussion, see e.g. Avraham et al. [2]. We (only) provide tools to implement such decisions. Moreover, any such implementation should be analyzed in terms of possible systemic implications.

An important point worth being stressed once more is that in order to be able to calculate discrimination-free prices one needs to have access to *all discriminatory characteristics* – otherwise it is not possible to properly adjust for the discriminatory characteristics’ influence. When it comes to gender, this is generally unproblematic, but if we want to adjust for, e.g., religious or philosophical beliefs, this is information which in general is not readily available. Customers may perceive it as weird to be approached with questions concerning this type of seemingly unrelated (and possible sensitive) information. A concrete example is discussed in De Jong and Ferries [7], where sexual preference is discussed as a risk factor relating to AIDS. Another comment concerning sensitive characteristics is the risk of obtaining untruthful answers which makes the collected data unreliable, see De Jong and Ferries [7].

An important position taken in the present paper concerns the role of the overall price prediction (on portfolio level). We have argued that the overall price prediction may be carried out using all available information, including the discriminatory one. Given this, it is the allocation of this overall cost that may introduce discrimination, and the discrimination-free pricing may be thought of as an allocation that avoids this. This in turn implies that it is primarily of interest to assess the predictive performance of the overall price prediction, since the discrimination-free (price and its induced) allocation *will in general be biased*. Still, there is nothing that prevents us from assessing the quality of the allocation using e.g. bootstrap techniques and cross-validation.

The argumentation used in the present paper has focused directly on how to obtain a discrimination-free price (based on the above allocation argument). This has lead us to a procedure which tells us how to adjust the best-estimate price (obtained using all available information) to arrive at a discrimination-free price. In a statistical sense, this is interpreted as a “discrimination-free point estimate”. A different line of thought instead could be that we try to develop a full statistical model that is discrimination-free, i.e. sacrificing predictive power by appropriately disregarding direct and indirect discrimination, this would result in a full statistical model that provides discrimination-free responses. An example of this approach in a life insurance context are the gender neutral intensities discussed in e.g. Chen and Vigna [5]. The main reason for considering prices directly is that we believe that this approach is closer to actuarial thinking, and because maximal predictive accuracy is a desirable feature in risk management.

REFERENCES

- [1] Vijay Aseervatham, Christoph Lex, and Martin Spindler. How do unisex rating regulations affect gender differences in insurance premiums? *The Geneva Papers on Risk and Insurance-Issues and Practice*, 41(1):128–160, 2016.
- [2] Ronen Avraham, Kyle D Logue, and Daniel B Schwarcz. Understanding insurance anti-discrimination laws. *Southern California Law Review*, 87(2):195–274, 2014.

- [3] Sandy Bruszas, Barbara Kaschützke, Raimond Maurer, and Ivonne Siegelin. Unisex pricing of German participating life annuities—Boon or bane for customer and insurance company? *Insurance: Mathematics and Economics*, 78:230–245, 2018.
- [4] An Chen, Montserrat Guillén, and Elena Vigna. Solvency requirement in a unisex mortality model. *ASTIN Bulletin: The Journal of the IAA*, 48(3):1219–1243, 2018.
- [5] An Chen and Elena Vigna. A unisex stochastic mortality model to comply with EU Gender Directive. *Insurance: Mathematics and Economics*, 73:124–136, 2017.
- [6] Imre Csiszár. I -divergence geometry of probability distributions and minimization problems. *The Annals of Probability*, 3(1):146–158, 1975.
- [7] Piet De Jong and Shauna Ferris. Adverse selection spirals. *ASTIN Bulletin: The Journal of the IAA*, 36(2):589–628, 2006.
- [8] European Commission. Guidelines on the application of COUNCIL DIRECTIVE 2004/113/EC to insurance, in the light of the judgment of the Court of Justice of the European Union in Case C-236/09 (Test-Achats). *Official Journal of the European Union*, C11:1–11, 2012.
- [9] European Council. COUNCIL DIRECTIVE 2004/113/EC – implementing the principle of equal treatment between men and women in the access to and supply of goods and services. *Official Journal of the European Union*, L 373:37–43, 2004.
- [10] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232, 2001.
- [11] Leo Guelman and Montserrat Guillén. A causal inference approach to measure price elasticity in automobile insurance. *Expert Systems with Applications*, 41(2):387–396, 2014.
- [12] Montserrat Guillén. Sexless and beautiful data: from quantity to quality. *Annals of Actuarial Science*, 6(2):231–234, 2012.
- [13] Samuel Krikler, Dan Dolberger, and Jacob Eckel. Method and tools for insurance price and revenue optimisation. *Journal of Financial Services Marketing*, 9(1):68–79, 2004.
- [14] Manabu Kuroki and Judea Pearl. Measurement bias and effect restoration in causal inference. *Biometrika*, 101(3):423–437, 2014.
- [15] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. In *Advances in Neural Information Processing Systems*, pages 4066–4076, 2017.
- [16] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009.
- [17] Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.
- [18] Devin G Pope and Justin R Sydnor. Implementing anti-discrimination policies in statistical profiling models. *American Economic Journal: Economic Policy*, 3(3):206–31, 2011.
- [19] Jörn Sass and Frank Thomas Seifried. Insurance markets and unisex tariffs: is the European Court of Justice improving or destroying welfare? *Scandinavian Actuarial Journal*, 2014(3):228–254, 2014.
- [20] Qingyuan Zhao and Trevor Hastie. Causal interpretations of black-box models. *Journal of Business & Economic Statistics*, (just-accepted):1–19, 2019.

APPENDIX A. FURTHER CAUSAL DIAGRAMS

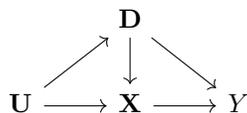


FIGURE 10. Causal diagram governed by DAG \mathfrak{G} considering unmeasured characteristics \mathbf{U} .

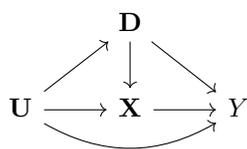


FIGURE 11. Causal diagram governed by DAG \mathfrak{G} considering unmeasured characteristics \mathbf{U} .

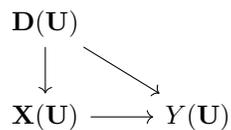


FIGURE 12. Causal diagram corresponding to Fig. 11 when the dependence of the unmeasured characteristics \mathbf{U} are suppressed.