

Fairness-Accuracy Trade-Offs: A Causal Perspective

Drago Plečko, Elias Bareinboim

Department of Computer Science, Columbia University
dp3144@columbia.edu, eb@cs.columbia.edu

Abstract

With the widespread adoption of AI systems, many of the decisions once made by humans are now delegated to automated systems. Recent works in the literature demonstrate that these automated systems, when used in socially sensitive domains, may exhibit discriminatory behavior based on sensitive characteristics such as gender, sex, religion, or race. In light of this, various notions of fairness and methods to quantify discrimination have been proposed, also leading to the development of numerous approaches for constructing fair predictors. At the same time, imposing fairness constraints may decrease the utility of the decision-maker, highlighting a tension between fairness and utility. This tension is also recognized in legal frameworks, for instance in the disparate impact doctrine of Title VII of the Civil Rights Act of 1964 – in which specific attention is given to considerations of *business necessity* – possibly allowing the usage of proxy variables associated with the sensitive attribute in case a high-enough utility cannot be achieved without them. In this work, we analyze the tension between fairness and accuracy from a causal lens for the first time. We introduce the notion of a path-specific excess loss (PSEL) that captures how much the predictor’s loss increases when a causal fairness constraint is enforced. We then show that the total excess loss (TEL), defined as the difference between the loss of predictor fair along all causal pathways vs. an unconstrained predictor, can be decomposed into a sum of more local PSELs. At the same time, enforcing a causal constraint often reduces the disparity between demographic groups. Thus, we introduce a quantity that summarizes the fairness-utility trade-off, called the causal fairness/utility ratio, defined as the ratio of the reduction in discrimination vs. the excess in the loss from constraining a causal pathway. This quantity is particularly suitable for comparing the fairness-utility trade-off across different causal pathways. Finally, as our approach requires causally-constrained fair predictors, we introduce a new neural approach for causally-constrained fair learning. Our approach is evaluated across multiple real-world datasets, providing new insights into the tension between fairness and accuracy.

1 Introduction

Automated decision-making systems based on machine learning and artificial intelligence are now commonly implemented in various critical sectors of society such as hiring,

This paper is part of the Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI 2025).

university admissions, law enforcement, credit assessments, and health care (Khandani, Kim, and Lo 2010; Mahoney and Mohen 2007; Brennan, Dieterich, and Ehret 2009). These technologies now significantly influence the lives of individuals and are frequently used in high-stakes settings (Topol 2019; Berk et al. 2021; Taddeo and Floridi 2018). As these systems replace or augment human decision-making processes, concerns about fairness and bias based on protected attributes such as race, gender, or religion have become a prominent consideration in the ML literature. The available data used to train automated systems may contain past and present societal biases as an imprint and therefore has the potential to perpetuate or even exacerbate discrimination against protected groups. This is highlighted by reports on biases in systems for sentencing (Angwin et al. 2016), facial recognition (Buolamwini and Gebru 2018), online ads (Sweeney 2013; Datta, Tschantz, and Datta 2015), and system authentication (Sanburn 2015), among many others. Despite the promise of AI to enhance human decision-making, the reality is that these technologies can also reflect or worsen societal inequalities. As alluded to before, the issue does not arise uniquely from the usage of automated systems; human-driven decision-making has long been analyzed in a similar fashion. Evidence of bias in human decision-making is abundant, including studies on the gender wage gap (Blau and Kahn 1992, 2017) and racial disparities in legal outcomes (Sweeney and Haney 1992; Pager 2003). Therefore, without proper care about fairness and transparency of the new generation of AI systems, it is unclear what its impact will be on the historically discriminated groups.

Within the growing literature on fair machine learning, a plethora of fairness definitions have been proposed. Commonly considered statistical criteria, among others, include demographic parity (independence (Darlington 1971)), equalized odds (separation (Hardt, Price, and Srebro 2016)), and calibration (sufficiency (Chouldechova 2017)). These definitions, however, have been shown as mutually incompatible (Barocas and Selbst 2016; Kleinberg, Mullainathan, and Raghavan 2016). Despite a number of proposals, there is still a lack of consensus on what the appropriate measures of fairness are, and how statistical notions of fairness could incorporate moral values of the society at large. For this reason, a number of works explored the causal approaches to fair machine learning (Kusner et al. 2017; Kilbertus et al. 2017; Nabi

and Shpitser 2018; Zhang and Bareinboim 2018b,a; Wu et al. 2019; Chiappa 2019; Plečko and Meinshausen 2020), and an in-depth discussion can be found in (Plečko and Bareinboim 2024). The main motivation for doing so is that the causal approach may allow the system designers to attribute the observed disparities between demographic groups to the causal mechanisms that underlie and generate them in the first place. In this way, by isolating disparities transmitted along different causal pathways, one obtains a more fine-grained analysis, and the capability to decide which causal pathways are deemed as unfair or discriminatory. More fundamentally, such considerations also form the basis of the legal frameworks for assessing discrimination in the United States and Europe. For instance, in the context of employment law, the disparate impact doctrine within the Title VII of the Civil Rights Act of 1964 (Act 1964) disallows any form of discrimination that results in a too large of a disparity between groups of interest. A core aspect of this doctrine, however, is the notion of *business necessity* (BN) or job-relatedness. Considerations of business necessity may allow variables correlated with the protected attribute to act as a proxy, and the law does not necessarily prohibit their usage due to their relevance to the business itself (or more broadly the utility of the decision-maker). Often, the wording that is used is that to argue business necessity in front of a court of law, the plaintiff needs to demonstrate that “there is no practice that is less discriminatory and achieves the same utility” (Els 1993). This concept illustrates the tension between fairness and utility, and demonstrates that we cannot be oblivious to considerations of utility from a legal standpoint. Moreover, demonstrating that a sufficient loss in accuracy results from imposing a fairness constraint has been previously used to justify business necessity considerations in some rulings of the European Court of Justice (Adams-Prassl, Binns, and Kelly-Lyth 2023; Weerts et al. 2023), emphasizing the importance of the topic studied in this paper.

Related work. We mention three parts of related literature. First, we mention the literature exploring fairness-utility trade-offs, such as (Corbett-Davies et al. 2017). The essential argument is that an unconstrained predictor always achieves a greater or equal utility than a constrained one. Many works find that introducing fairness constraints reduces utility (Mitchell et al. 2021), and propose ways of handling the fairness-utility trade-off (Fish, Kun, and Lelkes 2016).

However, other works in this literature still seem to be divided on the issue of whether trade-offs exist. For instance, (Rodolfa, Lamba, and Ghani 2021) finds that fairness and utility trade-offs are negligible in practice, while others argue that such trade-offs need not even exist (Maity et al. 2020; Dutta et al. 2020). Naturally, the implications on the predictor’s utility will strongly depend on the exact type of the fairness constraint that is enforced, and works that do not find a trade-off often focus on equality of odds (Hardt, Price, and Srebro 2016) ($\hat{Y} \perp\!\!\!\perp X \mid Y$) or (multi)calibration (Chouldechova 2017) ($Y \perp\!\!\!\perp X \mid \hat{Y}$). Notably, the former metric always allows for the perfect predictor $\hat{Y} = Y$, and thus in settings with good predictive power, the cost of enforcing this constraint may indeed be negligible. The latter metric

allows for the L_2 -optimal prediction score, and improving miscalibration may sometimes also yield improvements in utility by serving as a type of regularization.

Finally, in the causal fairness literature, the tension between fairness and utility has been largely unexplored. Some exceptions include works such as (Nilforoshan et al. 2022) which shows that for a decision policy that satisfies a causal fairness constraints, it is almost always possible to find another decision that has a higher utility and the same total variation (TV) measure. (Plečko and Bareinboim 2024a) performs a causal explanation on a decision score used for constructing a policy, and discusses how disparities in the decision score may influence utility. Our main aim of this paper is to fill in this gap, and provide a systematic way of analyzing the fairness-accuracy trade-off from a causal lens, and show that fairness and utility are almost always in a trade-off.

1.1 Motivating Example

We illustrate our approach in a simple linear setting:

Example 1 (Linear Fairness-Accuracy Causal Trade-Offs). *Consider variables X, W, Y behaving according to the following linear system of equations:*

$$X \leftarrow \text{Bernoulli}(0.5) \quad (1)$$

$$W \leftarrow \beta X + \epsilon_w \quad (2)$$

$$Y \leftarrow \alpha X + \gamma W + \epsilon_y, \quad (3)$$

where $\epsilon_w \sim N(0, \sigma_w^2)$, $\epsilon_y \sim N(0, \sigma_y^2)$. Variable X is the protected attribute, and Y is the outcome of interest. The causal diagram of Eqs. 1-3 is shown in Fig. 2 (with the Z set empty). Attribute X can influence Y along two different pathways: the direct path $X \rightarrow Y$, and the indirect path $X \rightarrow W \rightarrow Y$. Therefore, for considering fairness, we focus on fair linear predictors \hat{Y}^S of the form

$$\hat{Y}^S = \hat{\alpha}_S X + \hat{\gamma}_S W, \quad (4)$$

where the predictor \hat{Y}^S removes effects in the set S , with S ranging in $\{\emptyset, DE, IE, \{DE, IE\}\}$ (DE, IE stand for direct and indirect effects, and any subset of these could be removed). For instance, the optimal predictor \hat{Y}^\emptyset , which is not subject to any fairness constraints, has the coefficients

$$\hat{\alpha}_\emptyset = \alpha, \hat{\gamma}_\emptyset = \gamma, \quad (5)$$

which are the ordinary least squares (OLS) coefficients. Therefore, its mean-squared error (MSE) can be computed as $\mathbb{E}[Y - \hat{Y}^\emptyset]^2 = \sigma_y^2$. The DE -fair predictor \hat{Y}^{DE} , which has the direct effect constrained to zero, has coefficients

$$\hat{\alpha}_{DE} = 0, \hat{\gamma}_{DE} = \gamma. \quad (6)$$

The fully-fair predictor, labeled $\hat{Y}^{\{DE, IE\}}$, has both direct and indirect effects constrained to zero, and thus has coefficients

$$\hat{\alpha}_{\{DE, IE\}} = 0, \hat{\gamma}_{\{DE, IE\}} = 0. \quad (7)$$

Thus, the corresponding MSE values for \hat{Y}^{DE} and $\hat{Y}^{\{DE, IE\}}$ can be computed as (see Appendix A for computation details):

$$\mathbb{E}[Y - \hat{Y}^{DE}]^2 = \sigma_y^2 + \frac{\alpha^2}{2}, \quad (8)$$

$$\mathbb{E}[Y - \hat{Y}^{\{DE, IE\}}]^2 = \sigma_y^2 + \frac{\alpha^2 + \gamma^2 \beta^2 + \alpha \gamma \beta}{2} + \gamma^2 \sigma_w^2. \quad (9)$$

Our goal is to decompose the total excess loss (TEL) originating from imposing the fairness constraints, defined as:

$$TEL := \underbrace{\mathbb{E}[Y - \hat{Y}^{\{DE, IE\}}]^2}_{\text{fully-fair predictor's loss}} - \underbrace{\mathbb{E}[Y - \hat{Y}^{\emptyset}]^2}_{\text{unconstrained loss}}. \quad (10)$$

TEL measures the excess loss (in terms of the increase in the MSE, compared to the unconstrained predictor) originating from the removal of the direct and indirect effects. This quantifies the excess loss traded off for enforcing fairness constraints. Our goal is to decompose the TEL to obtain path-specific contributions originating from the removal of direct and indirect effects as follows:

$$TEL = \underbrace{\mathbb{E}[Y - \hat{Y}^{DE}]^2 - \mathbb{E}[Y - \hat{Y}^{\emptyset}]^2}_{\text{Term I = excess DE loss}} \quad (11)$$

$$+ \underbrace{\mathbb{E}[Y - \hat{Y}^{\{DE, IE\}}]^2 - \mathbb{E}[Y - \hat{Y}^{DE}]^2}_{\text{Term II = excess IE loss}} \quad (12)$$

$$= \underbrace{\frac{\alpha^2}{2}}_{\text{Term I}} + \underbrace{\frac{\gamma^2\beta^2 + \alpha\gamma\beta}{2}}_{\text{Term II}} + \gamma^2\sigma_w^2. \quad (13)$$

Term I is the direct effect excess loss, incurred by constraining the direct effect to 0. Term II is the indirect effect excess loss, incurred by constraining the indirect effect to 0.

At the same time, enforcing fairness constraints may reduce the disparity between groups. For any predictor \hat{Y} , we can measure the disparity using the difference in conditional expectations, called the total variation measure, defined as $TV_{x_0, x_1}(\hat{y}) = \mathbb{E}[\hat{Y} | x_1] - \mathbb{E}[\hat{Y} | x_0]$. This measure is sometimes also referred to as the parity gap. Similarly as for the TEL, we compute the decrease in group disparity associated with removing DE and IE, by comparing the TV measures of the fully-fair $\hat{Y}^{\{DE, IE\}}$ and unconstrained \hat{Y}^{\emptyset} , and computing the TV difference (TVD, for short), defined as:

$$TVD = TV(\hat{Y}^{\{DE, IE\}}) - TV(\hat{Y}^{\emptyset}) \quad (14)$$

$$= \underbrace{(\mathbb{E}[\hat{Y}^{\{DE, IE\}} | x_1] - \mathbb{E}[\hat{Y}^{\{DE, IE\}} | x_0])}_{\text{disparity after removing DE, IE}} \quad (15)$$

$$- \underbrace{(\mathbb{E}[\hat{Y}^{\emptyset} | x_1] - \mathbb{E}[\hat{Y}^{\emptyset} | x_0])}_{\text{disparity before removing DE, IE}} = -\alpha - \beta\gamma. \quad (16)$$

The TVD metric again decomposes into the contributions along direct and indirect effects:

$$TVD = \underbrace{TV(\hat{Y}^{DE}) - TV(\hat{Y}^{\emptyset})}_{\text{Term A = disparity reduction of DE}} \quad (17)$$

$$+ \underbrace{TV(\hat{Y}^{\{DE, IE\}}) - TV(\hat{Y}^{DE})}_{\text{Term B = disparity reduction of IE}}. \quad (18)$$

Terms A and B can be computed as:

$$\text{Term A} = (\mathbb{E}[\hat{Y}^{DE} | x_0] - \mathbb{E}[\hat{Y}^{DE} | x_0]) \quad (19)$$

$$- (\mathbb{E}[\hat{Y}^{\emptyset} | x_0] - \mathbb{E}[\hat{Y}^{\emptyset} | x_0]) = -\alpha \quad (20)$$

$$\text{Term B} = (\mathbb{E}[\hat{Y}^{\{DE, IE\}} | x_1] - \mathbb{E}[\hat{Y}^{\{DE, IE\}} | x_0]) \quad (21)$$

$$- (\mathbb{E}[\hat{Y}^{DE} | x_0] - \mathbb{E}[\hat{Y}^{DE} | x_0]) = -\beta\gamma. \quad (22)$$

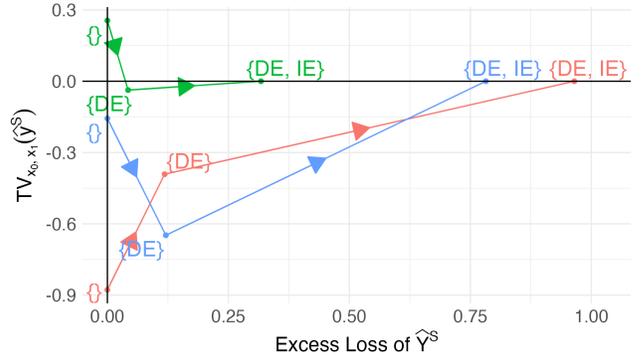


Figure 1: Total Variation (TV) measures vs. Excess Loss. Different colors represent trajectories obtained for different randomly sampled linear SCMs.

Based on the excess loss and the reduction in disparity resulting from constraining a causal path to zero, we can quantify the fairness/utility trade-off through a causal lens. Prototypical instances for predictors \hat{Y}^S are visualized in Fig. 1, for three randomly drawn triples (α, β, γ) . Predictor \hat{Y}^{\emptyset} is optimal and thus always has 0 excess loss (and hence lies on the vertical axis). Fully-fair predictor $\hat{Y}^{\{DE, IE\}}$ removes both direct and indirect effects, and thus always has the TV measure equal to 0. Therefore, $\hat{Y}^{\{DE, IE\}}$ always lies on the horizontal axis (corresponding to $TV = 0$). The slopes in the plot between \hat{Y}^{\emptyset} , \hat{Y}^{DE} , and $\hat{Y}^{DE, IE}$ (indicated by arrows), geometrically capture the tension between excess loss and reducing discrimination upon imposing a constraint. These slopes are computed as the ratio

$$\frac{TV \text{ difference}}{\text{Excess Loss}}, \quad (23)$$

a quantity that we call the causal fairness-utility ratio (CFUR). Based on Eqs. 11-12 and 19-22 we can compute

$$CFUR(DE) = -\frac{1}{\alpha}, \quad (24)$$

$$CFUR(IE) = -\frac{2\beta\gamma}{\gamma^2\beta^2 + \alpha\gamma\beta + 2\gamma^2\sigma_w^2}, \quad (25)$$

summarizing the fairness-utility trade-off for each path. \square

The above example illustrates how in the simple linear case we can attribute the increased loss from imposing fairness constraints to the specific causal pathway in question. It also shows we can compute the associated change in the disparity between groups, quantified by the TV measure $\mathbb{E}[\hat{Y} | x_1] - \mathbb{E}[\hat{Y} | x_0]$. In this paper, we generalize the approach from the above example to a non-parametric setting, with the following key contributions:

- (i) We introduce the notion of a *path-specific excess loss* (PSEL) associated with imposing a fairness constraint along a causal path (Def. 3), and we prove how the total excess loss (TEL) can be decomposed into a sum of path-specific excess losses (Thm. 1),

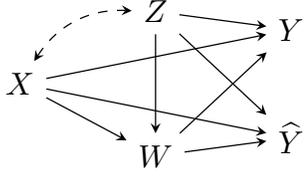


Figure 2: Standard Fairness Model, with the protected attribute X , set of confounders Z , set of mediators W , outcome Y , and a predictor \hat{Y} .

- (ii) We develop an algorithm for attributing path-specific excess losses to different causal paths (Alg. 1), allowing the system designer to explain how the total excess loss is affected by different fairness constraints. In this context, we show the equivalence of Alg. 1 with a Shapley value (Shapley et al. 1953) approach (Prop 3),
- (iii) For purposes of applying Alg. 1, a key requirement is the construction of causally-fair predictors \hat{Y}^S that remove effects along pathways in S . We introduce a novel Lagrangian formulation of the optimization problem for such \hat{Y}^S (Def. 5) accompanied with a training procedure for learning the predictor (Alg. 2),
- (iv) We introduce the causal fairness/utility ratio (CFUR, Def. 4) that summarizes how much the group disparity can be reduced per fixed cost in terms of excess loss. We compute CFURs on a range of real-world datasets, and demonstrate that from a causal viewpoint fairness and utility are almost always in tension.

1.2 Preliminaries

We use the language of structural causal models (SCMs) (Pearl 2000). An SCM is a tuple $\mathcal{M} := \langle V, U, \mathcal{F}, P(u) \rangle$, where V, U are sets of endogenous (observable) and exogenous (latent) variables, respectively, \mathcal{F} is a set of functions f_{V_i} , one for each $V_i \in V$, where $V_i \leftarrow f_{V_i}(\text{pa}(V_i), U_{V_i})$ for some $\text{pa}(V_i) \subseteq V$ and $U_{V_i} \subseteq U$. $P(u)$ is a strictly positive probability measure over U . Each SCM \mathcal{M} is associated to a causal diagram \mathcal{G} (Bareinboim et al. 2022) over the node set V where $V_i \rightarrow V_j$ if V_i is an argument of f_{V_j} , and $V_i \leftrightarrow V_j$ if the corresponding U_{V_i}, U_{V_j} are not independent. An instantiation of the exogenous variables $U = u$ is called a *unit*. By $Y_x(u)$ we denote the potential response of Y when setting $X = x$ for the unit u , which is the solution for $Y(u)$ to the set of equations obtained by evaluating the unit u in the submodel \mathcal{M}_x , in which all equations in \mathcal{F} associated with X are replaced by $X = x$. For more details about the causal inference background, we refer the reader to (Pearl 2000; Bareinboim et al. 2022; Plečko and Bareinboim 2024). Throughout the paper, we assume a specific cluster causal diagram \mathcal{G}_{SFM} known as the standard fairness model (SFM) (Plečko and Bareinboim 2024) over endogenous variables $\{X, Z, W, Y, \hat{Y}\}$ shown in Fig. 2. The SFM consists of the following: *protected attribute*, labeled X (e.g., gender, race, religion), assumed to be binary; the set of *confounding* variables Z , which are not causally influenced by the attribute

X (e.g., demographic information, zip code); the set of *mediator* variables W that are possibly causally influenced by the attribute (e.g., educational level or other job-related information); the *outcome* variable Y (e.g., GPA, salary); the *predictor* of the outcome \hat{Y} (e.g., predicted GPA, predicted salary). The SFM encodes the assumptions typically used in the causal inference literature about the lack of hidden confounding. The availability of the SFM and the implied assumptions are a possible limitation of the paper, while we note that partial identification techniques for bounding effects can be used for relaxing them (Zhang, Tian, and Bareinboim 2022). Based on the SFM, we will use the following causal fairness measures:

Definition 1 (Population-level Causal Fairness Measures (Pearl 2001; Plečko and Bareinboim 2024)). *The natural direct, indirect, and spurious effects are defined as*

$$NDE_{x_0, x_1}(y) = P(y_{x_1, W_{x_0}}) - P(y_{x_0}) \quad (26)$$

$$NIE_{x_1, x_0}(y) = P(y_{x_1, W_{x_0}}) - P(y_{x_1}) \quad (27)$$

$$NSE_x(y) = P(y | x) - P(y_x). \quad (28)$$

The NDE in Eq. 26 compares the potential outcome $Y_{x_1, W_{x_0}}$, in which Y responds to $X = x_1$ along the direct path, while W is set at the value it would attain naturally when responding to $X = x_0$, against the potential outcome Y_{x_0} where $X = x_0$ along both direct and indirect paths. In this way, the NDE measures the variations induced by changing $x_0 \rightarrow x_1$ along the direct causal path, quantifying direct discrimination. Similarly, the NIE in Eq. 27 compares $Y_{x_1, W_{x_0}}$ vs. Y_{x_1} , and thus captures variations induced by considering a change $x_1 \rightarrow x_0$ along the indirect causal path (note that both $Y_{x_1, W_{x_0}}$ and Y_{x_1} respond to $X = x_1$ along the direct path, so only indirect variations are induced when taking the difference). Finally, the NSE in Eq. 28 compares $Y | X = x$ vs. Y_x . In the former, due to conditioning on $X = x$, the distribution over the set of confounders Z (in Fig. 2) changes according to this conditioning, while in the potential outcome Y_x , the distribution of Z does not change, since $X = x$ is set by intervention. Therefore, taking the difference captures the spurious effect of X on Y , along the backdoor path $X \leftarrow Z \rightarrow Y$. A causally-fair predictor for a subset S of the above measures is defined as:

Definition 2 (Causally Fair Predictor (Plečko and Bareinboim 2024b)). *The optimal causally S -fair predictor \hat{Y}^S with respect to a loss function L and pathways in S is the solution to the following optimization problem:*

$$\arg \min_f \mathbb{E} L(Y, f(X, Z, W)) \quad (29)$$

$$s.t. \quad NDE_{x_0, x_1}(f) = NDE_{x_0, x_1}(y) \cdot \mathbb{1}(DE \notin S) \quad (30)$$

$$NIE_{x_1, x_0}(f) = NIE_{x_1, x_0}(y) \cdot \mathbb{1}(IE \notin S) \quad (31)$$

$$NSE_{x_0}(f) = NSE_{x_0}(y) \cdot \mathbb{1}(SE \notin S) \quad (32)$$

$$NSE_{x_1}(f) = NSE_{x_1}(y) \cdot \mathbb{1}(SE \notin S). \quad (33)$$

The definition of \hat{Y}^S has a straightforward interpretation. For any pathway in the set S , the corresponding causal effect should be 0, as proposed in the path-specific causal fairness literature (Nabi and Shpitser 2018; Chiappa 2019). However,

importantly, pathways that are not in S also need to be constrained – the effect of X on \hat{Y} along these paths should not change compared to the true outcome Y (Plecko and Bareinboim 2024b). For instance, if the direct path is not in S (meaning it is considered to be non-discriminatory), then we expect to have $\text{NDE}_{x_0, x_1}(\hat{y}) = \text{NDE}_{x_0, x_1}(y)$ (and similarly for IE, SE). This form of constraint on discriminatory pathways ensures that no undesirable bias amplification occurs along non-discriminatory pathways.

2 Path-Specific Excess Loss

In this section, we introduce the concept of a path-specific excess loss, and then demonstrate how the total excess loss can be decomposed into path-specific excess losses.

Definition 3 (Path-Specific Excess Loss). *Let $L(\hat{Y}, Y)$ be a loss function and \hat{Y}^S the optimal causally S -fair predictor with respect to L . Define the path-specific excess loss (PSEL) of a pair S, S' as:*

$$\text{PSEL}(S \rightarrow S') = \mathbb{E}[L(\hat{Y}^{S'}, Y)] - \mathbb{E}[L(\hat{Y}^S, Y)]. \quad (34)$$

The quantity $\text{PSEL}(\emptyset \rightarrow \{D, I, S\})$ is called the total excess loss (TEL).

The total excess loss computes the increase in the loss for the totally constrained predictor $\hat{Y}^{\{D, I, S\}}$ with direct, indirect, and spurious effects removed compared to the unconstrained predictor \hat{Y}^\emptyset .¹ In the following theorem, we show that the total excess loss can be decomposed as a sum of path-specific excess losses. All proofs are provided in Appendix B (for supplements, see full paper version at <https://arxiv.org/abs/2405.15443>):

Theorem 1 (Total Excess Loss Decomposition). *The total excess loss $\text{PSEL}(\emptyset \rightarrow \{D, I, S\})$ can be decomposed into a sum of path-specific excess losses as follows:*

$$\text{PSEL}(\emptyset \rightarrow \{D, I, S\}) = \text{PSEL}(\emptyset \rightarrow \{D\}) \quad (35)$$

$$+ \text{PSEL}(\{D\} \rightarrow \{D, I\}) \quad (36)$$

$$+ \text{PSEL}(\{D, I\} \rightarrow \{D, I, S\}). \quad (37)$$

Remark 2 (Non-Uniqueness of Decomposition). *The decomposition in Thm. 1 is not unique. In particular, the $\text{PSEL}(\emptyset \rightarrow \{D, I, S\})$ can be decomposed as*

$$\text{PSEL}(\emptyset \rightarrow \{S_1\}) + \text{PSEL}(\{S_1\} \rightarrow \{S_1, S_2\}) \quad (38)$$

$$+ \text{PSEL}(\{S_1, S_2\} \rightarrow \{D, I, S\}) \quad (39)$$

for any choice of $S_1, S_2 \in \{D, I, S\}$ with $S_1 \neq S_2$. Therefore, six different decompositions exist (three choices for S_1 , two for S_2).

Fig. 3 provides a graphical overview of all the possible path-specific excess losses. In the left side, we start with $S = \emptyset$ and the predictor \hat{Y}^\emptyset . Then, we can add any of $\{D, I, S\}$ to the S -set, to obtain the predictors \hat{Y}^D, \hat{Y}^I , or \hat{Y}^S , and

¹Other classifiers that remove only subsets of the causal paths between X and Y may be considered fair, depending on considerations of business necessity (Plecko and Bareinboim 2024). The rationale developed in this paper can be adapted to such settings.

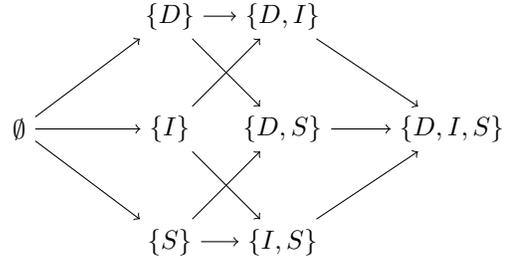


Figure 3: Graphical representation $\mathcal{G}_{\text{PSEL}}$.

so on. The graph representing all the possible states \hat{Y}^S and transitions between pairs $(\hat{Y}^S, \hat{Y}^{S'})$ shown in Fig. 3 is labeled $\mathcal{G}_{\text{PSEL}}$. There are six paths starting from \emptyset and ending in $\{D, I, S\}$. In Alg. 1, we introduce a procedure that sweeps over all the edges and paths in $\mathcal{G}_{\text{PSEL}}$ to compute path-specific excess losses, while also computing the change in the TV measure between groups in order to track the reduction in discrimination. In the main body of the paper we discuss fairness-utility trade-offs when considering direct, indirect, and spurious effects. In Appendix F we described how to adapt all the results to a general setting considering more granular path-specific effects. Formally, for any edge (S, S') in the graph $\mathcal{G}_{\text{PSEL}}$, the value of $\text{PSEL}(S \rightarrow S')$ is computed, together with the difference in the TV measure (TVD, for short) from the transition $S \rightarrow S'$, defined through the following expression

$$\text{TVD}(S \rightarrow S') = \underbrace{\mathbb{E}[Y^{S'} | x_1] - \mathbb{E}[Y^{S'} | x_0]}_{\text{TV after removing } S' \setminus S} \quad (42)$$

$$- \underbrace{\mathbb{E}[Y^S | x_1] - \mathbb{E}[Y^S | x_0]}_{\text{TV before removing } S' \setminus S}. \quad (43)$$

The quantities $\text{PSEL}(S \rightarrow S')$ and $\text{TVD}(S \rightarrow S')$ are naturally associated with the effect that was removed, i.e., $S' \setminus S$. In this context, we mention a connection with previous work (Zhang and Bareinboim 2018b), which provides a way of quantifying direct, indirect, and spurious effects of the attribute X on the outcome Y . In Appendix C, we show that our approach of quantifying the change in the TV measure through the TVD quantity in practice closely corresponds to methods for decomposing the TV measure into its direct, indirect, and spurious contributions (Zhang and Bareinboim 2018b).

As there are multiple ways of reaching the set $\{D, I, S\}$ from \emptyset in $\mathcal{G}_{\text{PSEL}}$, each of the corresponding effects (direct, indirect, spurious) will be associated with a number of different PSELs and TVDs (generally, note that the complexity is exponential in the number of causal paths included). In Eqs. 40-41 inside the algorithm, we compute the average PSEL and TVD across all the edges that are associated with a specific effect S_i . This simple intuition, corresponding to taking an average across all the possible decompositions of the total excess loss (Eq. 38), turns out to be equivalent to a Shapley value (Shapley et al. 1953) of a suitably chosen value function:

Algorithm 1: Path-Specific Excess Loss Attributions

Input: data \mathcal{D} , predictors \hat{Y}^S for S -sets $\subseteq \{D, I, S\}$

- 1 **foreach** edge $(S, S') \in \mathcal{G}_{\text{PSEL}}$ **do**
- 2 compute the path-specific excess loss of $S' \setminus S$, given by $\text{PSEL}(S \rightarrow S')$
- 3 compute the TV measure difference of $S' \setminus S$, written $\text{TVD}(S \rightarrow S')$ given by $\text{TV}_{x_0, x_1}(\hat{Y}^{S'}) - \text{TV}_{x_0, x_1}(\hat{Y}^S)$
- 4 **foreach** causal path $S_i \in \{D, I, S\}$ **do**
- 5 compute the average path-specific excess loss and TV difference across all paths $\emptyset \rightarrow \dots \rightarrow \{D, I, S\}$ in $\mathcal{G}_{\text{PSEL}}$
- 6 $\text{APSEL}(S_i)$ is computed as

$$\frac{1}{3!} \sum_{\substack{\pi \in \mathcal{G}_{\text{PSEL}}: \\ \emptyset \text{ to } \{D, I, S\}}} \text{PSEL}(\pi^{<S_i} \rightarrow \pi^{<S_i} \cup S_i) \quad (40)$$
- 7 $\text{ATVD}(S_i)$ is computed as

$$\frac{1}{3!} \sum_{\substack{\pi \in \mathcal{G}_{\text{PSEL}}: \\ \emptyset \text{ to } \{D, I, S\}}} \text{TVD}(\pi^{<S_i} \rightarrow \pi^{<S_i} \cup S_i), \quad (41)$$

where $\pi^{<S_i}$ denotes the set of causal paths that precede S_i in π .
- 8 **return** set of $\text{PSEL}(S \rightarrow S')$, $\text{TVD}(S \rightarrow S')$, attributions $\text{APSEL}(S_i)$, $\text{ATVD}(S_i)$

Proposition 3 (PSEL Attribution as Shapley Values). *Let the functions $f_1(S)$, $f_2(S)$ be defined as:*

$$f_1(S) = \text{PSEL}(\emptyset \rightarrow S) \quad (44)$$

$$f_2(S) = \text{TVD}(\emptyset \rightarrow S). \quad (45)$$

The Shapley value $\phi^k(S_i)$ for the effect $S_i \in \{D, I, S\}$ and function f_k , is computed as

$$\sum_{S \subseteq \{D, I, S\} \setminus \{S_i\}} \frac{1}{n \binom{n-1}{|S|}} (f_k(S \cup \{S_i\}) - f_k(S)). \quad (46)$$

where $n = 3$ for the choice $\{D, I, S\}$. The averaged path-specific excess loss of S_i and the averaged TV difference of S_i are equal to the Shapley values of S_i associated with functions f_1, f_2 , respectively,

$$\phi^1(S_i) = \text{APSEL}(S_i) \quad (47)$$

$$\phi^2(S_i) = \text{ATVD}(S_i), \quad (48)$$

with $\text{APSEL}(S_i)$ defined as

$$\frac{1}{3!} \sum_{\substack{\pi \in \mathcal{G}_{\text{PSEL}}: \\ \emptyset \text{ to } \{D, I, S\}}} \text{PSEL}(\pi^{<S_i} \rightarrow \pi^{<S_i} \cup S_i), \quad (49)$$

and $\text{ATVD}(S_i)$ defined as

$$\frac{1}{3!} \sum_{\substack{\pi \in \mathcal{G}_{\text{PSEL}}: \\ \emptyset \text{ to } \{D, I, S\}}} \text{TVD}(\pi^{<S_i} \rightarrow \pi^{<S_i} \cup S_i). \quad (50)$$

The above proposition illustrates how averaging the influence of removing a causal effect over all possible ways of reaching $\{D, I, S\}$ from \emptyset is equivalent to computing the Shapley values of an appropriate value function f . We remark that the attribution in Prop. 3 is not the only approach one could take. An alternative would be to average the contributions of each pathways across *all edges* of the graph $\mathcal{G}_{\text{PSEL}}$ (instead of focusing on all pathways from \emptyset to $\{D, I, S\}$). Such an attribution would not satisfy the Shapley axioms, however. We next introduce the notion of a causal fairness/utility ratios:

Definition 4 (Causal Fairness/Utility Ratio (CFUR)). *The causal fairness/utility ratio (CFUR) for a causal path S_i is defined as*

$$\text{CFUR}(S_i) = \frac{\text{ATVD}(S_i)}{\text{APSEL}(S_i)}. \quad (51)$$

Whenever $\text{APSEL}(S_i) = 0$, $\text{CFUR}(S_i)$ is equal to 0.

The CFUR quantity may be particularly useful for comparing different causal effects, and the connection of the CFUR with local TVD and PSEL values is described in Appendix D. The intuition behind the quantity is simple – for removing a causal effect S_i from our predictor \hat{Y} , we want to compute how much of a reduction in the disparity that results in (measured in terms of the ATVD measure) *per unit change in the incurred excess loss*. This quantity attempts to assign a single number to a causal path that succinctly summarizes how much fairness is gained vs. how much predictive power is lost by imposing such a causal constraint.

3 Causally-Fair Constrained Learning

In the preceding section, we developed an approach for quantifying the tension between fairness and accuracy from a causal viewpoint. The results were contingent on finding the optimal causally-fair predictors \hat{Y}^S following Def. 2. However, computing the predictors \hat{Y}^S is quite challenging in practice, due to several complex causal constraints in the optimization problem. We now develop a practical approach for solving this problem, by first introducing a Lagrangian form of the optimal causally-fair predictor:

Definition 5 (Causal Lagrange Predictor \hat{Y}^S). *The causally S -fair λ -optimal predictor $\hat{Y}^S(\lambda)$ with respect to pathways in S and the loss function L is the solution of the following:*

$$\arg \min_f \mathbb{E} L(Y, f(X, Z, W)) + \quad (52)$$

$$\lambda (\text{NDE}_{x_0, x_1}(f) - \text{NDE}_{x_0, x_1}(y) \cdot \mathbb{1}(\text{DE} \notin S))^2 + \quad (53)$$

$$\lambda (\text{NIE}_{x_1, x_0}(f) - \text{NIE}_{x_1, x_0}(y) \cdot \mathbb{1}(\text{IE} \notin S))^2 + \quad (54)$$

$$\lambda (\text{NSE}_{x_0}(f) - \text{NSE}_{x_0}(y) \cdot \mathbb{1}(\text{SE} \notin S))^2 + \quad (55)$$

$$\lambda (\text{NSE}_{x_1}(f) - \text{NSE}_{x_1}(y) \cdot \mathbb{1}(\text{SE} \notin S))^2 \quad (56)$$

The above definition reformulates the problem of finding \hat{Y}^S to a Lagrangian form. This makes the problem amenable to standard gradient descent methods, and we propose a procedure for finding a suitable predictor \hat{Y}^S in Alg. 2 (CFCL). In principle, many non-parametric learners such as boosting or neural networks could be used for fitting \hat{Y}^S , and here we describe a neural approach. The key challenge for constructing

Algorithm 2: Causally-Fair Constrained Learning (CFCL)

Input: training data \mathcal{D}_t , evaluation data \mathcal{D}_e , set S , precision ϵ

- 1 Set $\lambda_{\text{low}} = 0$, $\lambda_{\text{high}} = \text{large}$
 - 2 **while** $|\lambda_{\text{high}} - \lambda_{\text{low}}| > \epsilon \mathbf{do}$
 - 3 set $\lambda_{\text{mid}} = \frac{1}{2}(\lambda_{\text{low}} + \lambda_{\text{high}})$
 - 4 fit a neural network to solves the optimization problem in Eqs. 52-56 with $\lambda = \lambda_{\text{mid}}$ on \mathcal{D}_t to obtain the predictor $\hat{Y}^S(\lambda_{\text{mid}})$
 - 5 compute the causal measures of fairness NDE, NIE, NSE of $\hat{Y}^S(\lambda_{\text{mid}})$ on evaluation data \mathcal{D}_e
 - 6 test the hypothesis

$$H_0^{CE} : \text{NCE}(\hat{y}^S(\lambda_{\text{mid}})) = \text{NCE}(y) \cdot \mathbb{1}(\text{CE} \notin S) \quad (57)$$

where NCE ranges in
 $\text{NDE}_{x_0, x_1}, \text{NIE}_{x_1, x_0}, \text{NSE}_{x_0}, \text{NSE}_{x_1}$
 - 7 **if any of** $H_0^{DE}, H_0^{IE}, H_0^{SE_0}, H_0^{SE_1}$ **rejected then**
 $\lambda_{\text{low}} = \lambda_{\text{mid}}$ **else** $\lambda_{\text{high}} = \lambda_{\text{mid}}$
 - 8 **return predictor** $\hat{Y}^S(\lambda_{\text{mid}})$
-

\hat{Y}^S is finding the appropriate value of the tuning parameter λ . While the user may simply use a grid of values of λ , and inspect the loss function and the fairness measures, and then choose a λ value, we propose a data-driven approach. We note that if λ is too small, insufficient weight may be given to the fairness constraints, which therefore may be violated on new test data. If λ is too high, however, we may give insufficient weight to minimizing the loss L , which may lead to poor performance on test data. Therefore, we propose a binary search type of procedure that first splits the data into train and evaluation folds, \mathcal{D}_t and \mathcal{D}_e . CFCL starts with an interval $[\lambda_{\text{low}}, \lambda_{\text{high}}]$ and takes the midpoint λ_{mid} . For this parameter value, it computes the optimal predictor $\hat{Y}^S(\lambda_{\text{mid}})$ for the optimization problem in Eqs. 52-56 by fitting a feed-forward neural network with n_h hidden layers and n_v nodes in each layer. Then, for this fixed value of λ_{mid} , we test the hypotheses that

$$H_0^{CE} : \text{NCE}(\hat{y}^S(\lambda_{\text{mid}})) = \text{NCE}(y) \cdot \mathbb{1}(\text{CE} \notin S) \quad (58)$$

on evaluation data \mathcal{D}_e (done in Eq. 57), which essentially test if the fairness constraints hold on the evaluation set \mathcal{D}_e , i.e., out of sample. If any of the hypotheses are rejected, it means that this value of λ is too small to ensure that the fairness constraints are satisfied on new training data. Therefore, we want to find a larger λ , and the algorithm moves to the interval $[\lambda_{\text{mid}}, \lambda_{\text{high}}]$. If none of the hypotheses are rejected, it means that λ_{mid} is large enough to enforce the fairness constraints, and there may be an even smaller λ that achieves this, so the algorithm moves to the interval $[\lambda_{\text{low}}, \lambda_{\text{mid}}]$. In this way, CFCL leads to a data-driven way to select the tuning parameter λ . As the number of training and evaluation samples increases $|\mathcal{D}_t|, |\mathcal{D}_e| \rightarrow \infty$, the method is expected to perform increasingly well. An alternative approach would

be to use a framework that automatically allows the learning of the λ parameter (Fioretto et al. 2021).

4 Experiments

In this section, we perform the causal fairness-accuracy analysis described in Sec. 2 on the Census 2018 dataset (Ex. 2). Additional analyses of the COMPAS (Ex. 3) and UCI Credit (Ex. 4) datasets are reported in Appendix E. All code for reproducing the experiments can be found in our Github repository <https://github.com/dplecko/causal-acc-decomp>.

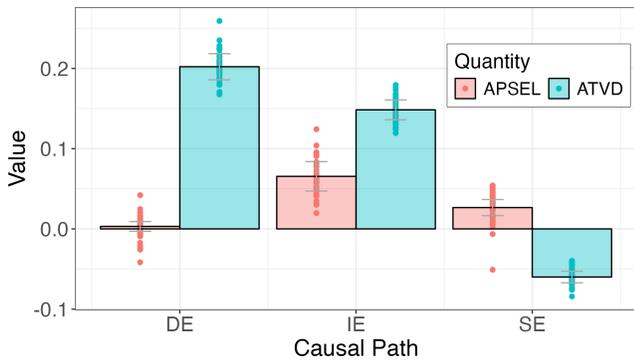
Example 2 (Salary Increase of Government Employees (Plečko and Bareinboim 2024)). The US government is building a tool for automated allocation of salaries for new employees. For developing the tool, they use the data collected by the United States Census Bureau in 2018, including

- confounders Z , consisting of demographic information (Z_1 for age, Z_2 for race, Z_3 for nationality),
- gender X (x_0 female, x_1 male),
- mediators W , including marital and family status M , education L , and work-related information R ,
- outcome Y , salary.

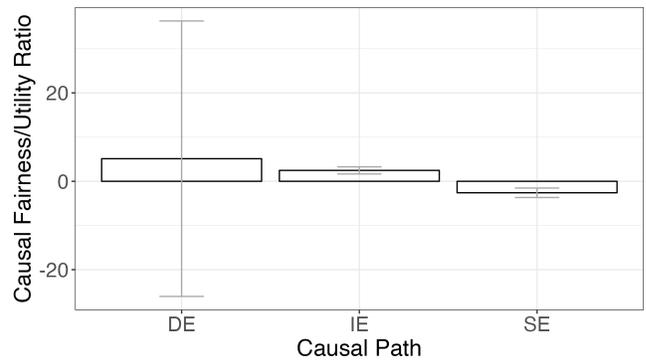
The government wants to predict the outcome Y , the yearly salary of the employees (transformed to a log-scale), in order to assign salaries for prospective employees. The standard fairness model (Fig. 2) is constructed as $\{X = X, Z = \{Z_1, Z_2, Z_3\}, W = \{M, L, R\}, Y = Y\}$.

The team developing the ML predictor is also concerned with the fairness of the allocated salaries. In particular, they wish to understand how the different causal effects from the protected attribute X to the predictor \hat{Y} affect the prediction, and how much the salary predictions would have to deviate from the optimal prediction to remove an effect along a specific pathway (in particular, they focus on the root mean squared error (RMSE) loss). For analyzing this, they utilize the tools from Alg. 1, and build causally fair predictors \hat{Y}^S (for different choices of S -sets) using Alg. 2.

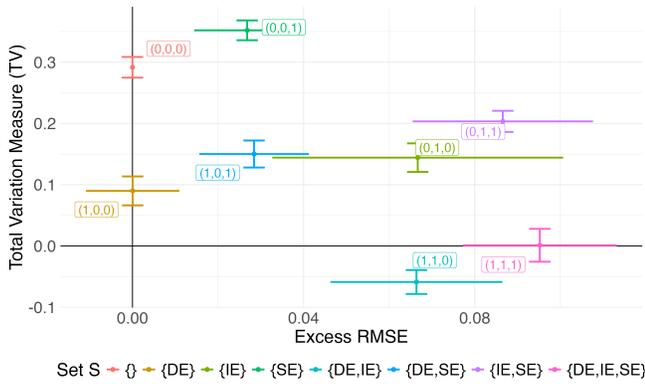
The analysis results are shown in Fig. 4, with uncertainty bars indicating standard deviations over 10 bootstrap repetitions. In the analysis of PSEL and TVD values (Fig. 4a), the team notices that imposing fairness constraints does not reduce RMSE for any of the effects. The largest excess loss is observed for the indirect effect, and smaller excess losses for direct and spurious effects. When looking at TVD values, they find that removing direct and indirect effects reduces the group differences substantially. In terms of causal fairness-utility ratios (Fig. 4b), the team finds that removing the direct effect has the best value in terms of reducing the disparity between groups vs. increasing the loss. The TV measure vs. excess loss dependence for different predictors \hat{Y}^S is shown in Fig. 4c (binary labels (D, I, S) in the figure indicate which effects were removed). The graph $\mathcal{G}_{\text{PSEL}}$ with the values of PSEL and TVD for each transition is shown in Fig. 4d. Based on the analysis, the team realizes that it is possible to substantially reduce discrimination with a small amount of excess RMSE loss. They decide to implement the predictor \hat{Y}^D with the direct effect removed. \square



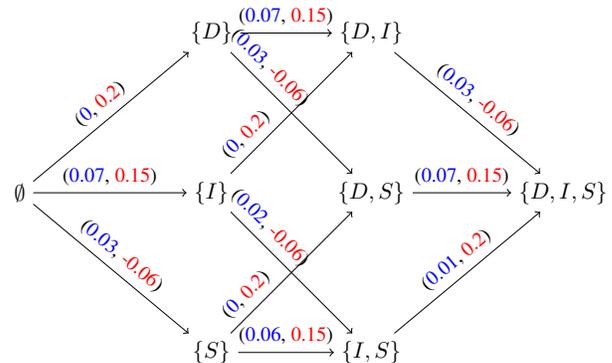
(a) (A)PSEL and (A)TVD values.



(b) CFUR values.



(c) Fairness/Utility of \hat{Y}^S predictors.



(d) $\mathcal{G}_{\text{PSEL}}$ with PSEL (blue) and TVD values (red).

Figure 4: Application of Alg. 1 on the Census 2018 dataset. (a) Estimated APSEL (Eq. 40) and ATVD (Eq. 41) values; (b) The causal fairness-utility ratios (Eq. 51); (c) The Pareto plot for trade-offs between fairness (TV measure on vertical axis) and utility (excess RMSE on horizontal axis) for different predictors. The vector (s_1, s_2, s_3) indicates which of the DE/IE/SE pathways are constrained to zero; (d) The $\mathcal{G}_{\text{PSEL}}$ graph from Fig. 5 populated with PSEL and TVD values.

5 Conclusion

The tension between fairness and accuracy is a fundamental concern in the modern applications of machine learning. The importance of this tension is also recognized in the legal frameworks of anti-discrimination, such as the disparate impact doctrine, which may allow for the usage of covariates correlated with the protected attribute if they are sufficiently important for the decision-maker’s utility – in legal texts, this concept is known as business necessity or job-relatedness. In this work, we developed tools for analyzing the fairness-accuracy trade-off from a causal standpoint. Our approach allows the system designer to quantify how much excess loss is incurred when removing a path-specific causal effect from an automated predictor (Def. 3). We also showed how the total excess loss, defined as the difference between the loss of the predictor fair along all causal pathways vs. an unconstrained predictor, can be decomposed into a sum of path-specific excess losses (Thm. 1). At the same time, enforcing fairness constraints may reduce the overall disparity between groups. Based on this, we developed an algorithm for attributing excess loss to different causal pathways (Alg. 1), and introduced the notion of a causal fairness-utility ratio

that captures the

$$\frac{\text{fairness gain}}{\text{excess loss}}$$

ratio and in this way summarizes the trade-off for each causal path. Since our approach requires access to causally-fair predictors (Def. 2), we introduced a new neural approach for constructing such predictors (Def. 5, Alg. 2). Finally, we analyzed several real-world datasets, in order to investigate if fairness and utility are in a trade-off in practice. Our findings are that, from causal perspective, fairness and utility are almost always in tension (see Exs. 2-4), contrary to some other works appearing in the fair ML literature.

Acknowledgements

This research was supported in part by the NSF, ONR, AFOSR, DoE, Amazon, JP Morgan, and The Alfred P. Sloan Foundation.

References

1993. *Elston v. Talladega County Bd. of Educ.* 997 F.2d 1394 (11th Cir. 1993). United States Court of Appeals for the Eleventh Circuit.

- Act, C. R. 1964. Civil rights act of 1964. *Title VII, Equal Employment Opportunities*.
- Adams-Prassl, J.; Binns, R.; and Kelly-Lyth, A. 2023. Directly discriminatory algorithms. *The Modern Law Review*, 86(1): 144–175.
- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2016. Machine bias: There’s software used across the country to predict future criminals. And it’s biased against blacks. *ProPublica*.
- Bareinboim, E.; Correa, J. D.; Ibeling, D.; and Icard, T. 2022. On Pearl’s Hierarchy and the Foundations of Causal Inference. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, 507–556. New York, NY, USA: Association for Computing Machinery, 1st edition.
- Barocas, S.; and Selbst, A. D. 2016. Big data’s disparate impact. *Calif. L. Rev.*, 104: 671.
- Berk, R.; Heidari, H.; Jabbari, S.; Kearns, M.; and Roth, A. 2021. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 50(1): 3–44.
- Blau, F. D.; and Kahn, L. M. 1992. The gender earnings gap: learning from international comparisons. *The American Economic Review*, 82(2): 533–538.
- Blau, F. D.; and Kahn, L. M. 2017. The gender wage gap: Extent, trends, and explanations. *Journal of economic literature*, 55(3): 789–865.
- Brennan, T.; Dieterich, W.; and Ehret, B. 2009. Evaluating the predictive validity of the COMPAS risk and needs assessment system. *Criminal Justice and Behavior*, 36(1): 21–40.
- Buolamwini, J.; and Gebru, T. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In Friedler, S. A.; and Wilson, C., eds., *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, 77–91. NY, USA.
- Chiappa, S. 2019. Path-specific counterfactual fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7801–7808.
- Chouldechova, A. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. Technical Report arXiv:1703.00056, arXiv.org.
- Corbett-Davies, S.; Pierson, E.; Feller, A.; Goel, S.; and Huq, A. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*, 797–806.
- Darlington, R. B. 1971. Another look at “cultural fairness” 1. *Journal of educational measurement*, 8(2): 71–82.
- Datta, A.; Tschantz, M. C.; and Datta, A. 2015. Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination. *Proceedings on Privacy Enhancing Technologies*, 2015(1): 92–112.
- Dutta, S.; Wei, D.; Yueksel, H.; Chen, P.-Y.; Liu, S.; and Varshney, K. 2020. Is there a trade-off between fairness and accuracy? a perspective using mismatched hypothesis testing. In *International conference on machine learning*, 2803–2813. PMLR.
- Fioretto, F.; Van Hentenryck, P.; Mak, T. W.; Tran, C.; Baldo, F.; and Lombardi, M. 2021. Lagrangian duality for constrained deep learning. In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part V*, 118–135. Springer.
- Fish, B.; Kun, J.; and Lelkes, Á. D. 2016. A confidence-based approach for balancing fairness and accuracy. In *Proceedings of the 2016 SIAM international conference on data mining*, 144–152. SIAM.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29: 3315–3323.
- Khandani, A. E.; Kim, A. J.; and Lo, A. W. 2010. Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11): 2767–2787.
- Kilbertus, N.; Rojas Carulla, M.; Parascandolo, G.; Hardt, M.; Janzing, D.; and Schölkopf, B. 2017. Avoiding discrimination through causal reasoning. *Advances in neural information processing systems*, 30.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*.
- Kusner, M. J.; Loftus, J.; Russell, C.; and Silva, R. 2017. Counterfactual fairness. *Advances in neural information processing systems*, 30.
- Larson, J.; Mattu, S.; Kirchner, L.; and Angwin, J. 2016. How we analyzed the COMPAS recidivism algorithm. *ProPublica* (5 2016), 9.
- Mahoney, J. F.; and Mohen, J. M. 2007. Method and system for loan origination and underwriting. US Patent 7,287,008.
- Maity, S.; Mukherjee, D.; Yurochkin, M.; and Sun, Y. 2020. There is no trade-off: enforcing fairness can improve accuracy.
- Mitchell, S.; Potash, E.; Barocas, S.; D’Amour, A.; and Lum, K. 2021. Algorithmic fairness: Choices, assumptions, and definitions. *Annual review of statistics and its application*, 8(1): 141–163.
- Nabi, R.; and Shpitser, I. 2018. Fair inference on outcomes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Nilforoshan, H.; Gaebler, J. D.; Shroff, R.; and Goel, S. 2022. Causal conceptions of fairness and their consequences. In *International Conference on Machine Learning*, 16848–16887. PMLR.
- Pager, D. 2003. The mark of a criminal record. *American journal of sociology*, 108(5): 937–975.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press. 2nd edition, 2009.
- Pearl, J. 2001. Direct and Indirect Effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, 411–420. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

Plečko, D.; and Bareinboim, E. 2024. Causal Fairness Analysis: A Causal Toolkit for Fair Machine Learning. *Foundations and Trends® in Machine Learning*, 17(3): 304–589.

Plečko, D.; and Bareinboim, E. 2024a. Causal Fairness for Outcome Control. *Advances in Neural Information Processing Systems*, 36.

Plečko, D.; and Bareinboim, E. 2024b. Reconciling predictive and statistical parity: A causal approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38 (13), 14625–14632.

Plečko, D.; and Meinshausen, N. 2020. Fair data adaptation with quantile preservation. *Journal of Machine Learning Research*, 21: 242.

Rodolfa, K. T.; Lamba, H.; and Ghani, R. 2021. Empirical observation of negligible fairness–accuracy trade-offs in machine learning for public policy. *Nature Machine Intelligence*, 3(10): 896–904.

Sanburn, J. 2015. Facebook Thinks Some Native American Names Are Inauthentic. *Time*.

Shapley, L. S.; et al. 1953. *A value for n-person games*. Princeton University Press Princeton.

Sweeney, L. 2013. Discrimination in Online Ad Delivery. Technical Report 2208240, SSRN.

Sweeney, L. T.; and Haney, C. 1992. The influence of race on sentencing: A meta-analytic review of experimental studies. *Behavioral Sciences & the Law*, 10(2): 179–195.

Taddeo, M.; and Floridi, L. 2018. How AI can be a force for good. *Science*, 361(6404): 751–752.

Topol, E. J. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nature medicine*, 25(1): 44–56.

Weerts, H.; Xenidis, R.; Tarissan, F.; Olsen, H. P.; and Pechenizkiy, M. 2023. Algorithmic unfairness through the lens of EU non-discrimination law: Or why the law is not a decision tree. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 805–816.

Wu, Y.; Zhang, L.; Wu, X.; and Tong, H. 2019. Pc-fairness: A unified framework for measuring causality-based fairness. *Advances in neural information processing systems*, 32.

Yeh, I.-C. 2016. Default of Credit Card Clients. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C55S3H>.

Zhang, J.; and Bareinboim, E. 2018a. Equality of Opportunity in Classification: A Causal Approach. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*, 3671–3681. Montreal, Canada: Curran Associates, Inc.

Zhang, J.; and Bareinboim, E. 2018b. Fairness in decision-making—the causal explanation formula. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Zhang, J.; Tian, J.; and Bareinboim, E. 2022. Partial Counterfactual Identification from Observational and Experimental Data. In *Proceedings of the 39th International Conference on Machine Learning*.