

Reconciling Predictive and Statistical Parity: A Causal Approach

Drago Plečko and Elias Bareinboim

Department of Computer Science, Columbia University, New York, NY 10027
dp3144@columbia.edu, eb@cs.columbia.edu

Abstract

Since the rise of fair machine learning as a critical field of inquiry, many different notions on how to quantify and measure discrimination have been proposed in the literature. Some of these notions, however, were shown to be mutually incompatible. Such findings make it appear that numerous different kinds of fairness exist, thereby making a consensus on the appropriate measure of fairness harder to reach, hindering the applications of these tools in practice. In this paper, we investigate one of these key impossibility results that relates the notions of statistical and predictive parity. Specifically, we derive a new causal decomposition formula for the fairness measures associated with predictive parity, and obtain a novel insight into how this criterion is related to statistical parity through the legal doctrines of disparate treatment, disparate impact, and the notion of business necessity. Our results show that through a more careful causal analysis, the notions of statistical and predictive parity are not really mutually exclusive, but complementary and spanning a spectrum of fairness notions through the concept of business necessity. Finally, we demonstrate the importance of our findings on a real-world example.

1 Introduction

As society increasingly relies on AI-based tools, an ever larger number of decisions that were once made by humans are now delegated to automated systems, and this trend is likely to only accelerate in the coming years. Such automated systems may exhibit discrimination based on gender, race, religion, or other sensitive attributes, as witnessed by various examples in criminal justice (Angwin et al. 2016), facial recognition (Harwell 2019; Buolamwini and Gebru 2018), targeted advertising (Detrixhe and Merrill 2019), and medical treatment allocation (Rajkomar et al. 2018), to name a few.

In light of these challenges, a large amount of effort has been invested in attempts to detect and quantify undesired discrimination based on society’s current ethical standards, and then design learning methods capable of removing possible unfairness from future predictions and decisions. During this process, many different notions on how to quantify discrimination have been proposed. In fact, the current literature is abundant with different fairness metrics, some of

This paper is part of the Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI 2024).

which are mutually incompatible (Corbett-Davies and Goel 2018). The incompatibility of these measures can create a serious obstacle for practitioners since choosing among them, even for the system designer, is usually non-trivial.

In the real world, issues of discrimination and unfairness are analyzed through two major legal doctrines. The first one is *disparate treatment*, which enforces the equality of treatment of different groups and prohibits the use of the protected attribute (e.g., race) during the decision process. One of the legal formulations for showing disparate treatment is that “a similarly situated person who is not a member of the protected class would not have suffered the same fate” (Barocas and Selbst 2016). Disparate treatment is commonly associated with the notion of direct effects in the causal literature. The second doctrine is known as *disparate impact* and focuses on *outcome fairness*, namely, the equality of outcomes among protected groups. Discrimination through disparate impact occurs if a facially neutral practice has an adverse impact on members of the protected group, including cases where discrimination is unintended or implicit. In practice, the law may not necessarily prohibit the usage of all characteristics correlated with the protected attribute due to their relevance to the business itself, which is legally known as “business necessity” (labeled BN from now on) or “job-relatedness”. Therefore, some of the variables may be used to distinguish between individuals, even if they are associated with the protected attribute (Kilbertus et al. 2017). From a causal perspective, disparate impact is realized through indirect forms of discrimination, and taking into account BN considerations is the essence of this doctrine (Barocas and Selbst 2016).

1.1 Relationship to Previous Literature

Consider now a set of causal pathways between the attribute X and the predictor \hat{Y} , labeled C_1, \dots, C_k . For example, these pathways could represent the direct, indirect, and spurious effects of X on \hat{Y} , or more generally any set of path-specific effects. Previous works on path-specific counterfactual fairness (Nabi and Shpitser 2018; Wu et al. 2019; Chippa 2019) constrain the effects along paths C_i that are considered discriminatory (i.e., *not* in the BN set), by requiring them to be equal to 0, written $C_i(X, \hat{Y}) = 0$. Other works demonstrate how such causal variations can be disentangled




Figure 1: BN specification (left) and the spectrum between statistical and predictive parity (right).

within statistical measures of fairness (Zhang and Bareinboim 2018b,a).

Importantly, however, these works do not specify anything about causal effects along pathways that *are in the BN set*, i.e., are not considered discriminatory. If C_i is in the BN set, the transmitted causal effect does not need to equal 0 in general, i.e., $C(X, \hat{Y}) \neq 0$ may be allowed. In this paper, we demonstrate that a natural way of constraining effects along BN paths exists, and the transmitted causal effect should not take an arbitrary value in this case. Rather, it should equal the *transmitted effect from X to the original outcome Y (observed in the real world) along the same pathway*, written as $C_i(X, \hat{Y}) = C_i(X, Y)$. In this way, we complement the existing literature on path-specific fairness.

As we will demonstrate, considerations of business necessity are also closely related to the impossibility result between statistical and predictive parity (Barocas, Hardt, and Narayanan 2017). When reasoning about fairness, the system designer needs to decide which of the causal pathways C_1, \dots, C_k are considered discriminatory, and which ones are not, as shown in Fig. 1 (left side). Therefore, considerations of BN can be summarized as a 0/1 vector of length k , where each pathway C_i is represented by an entry. As we demonstrate both intuitively and formally throughout the paper, the choice of the BN set being empty, written $(0, \dots, 0)$, will ensure the notion of statistical parity. The choice of the BN set that includes all pathways, written $(1, \dots, 1)$, will lead to predictive parity. Crucially, various intermediate fairness notions between the two ends of the spectrum are possible, depending on what is allowed or not; for an illustration, see Fig. 1 (right side).

The unification of the principles behind statistical and predictive parity through the concept of BN offers a new perspective, by providing an argument against the seemingly discouraging impossibility result between them. The practitioner is no longer faced with a false dichotomy of choosing between statistical or predictive parity but rather faces a spectrum of different fairness notions determined by the choice of the business necessity set, which is usually fixed through societal consensus and legal requirements.

Finally, we also mention the work on counterfactual predictive parity (Ctf-PP) (Coston et al. 2020) that is similar in name to our notion of causal predictive parity, but is in fact a very different notion. Ctf-PP deals with the setting of decision-making and considers counterfactuals of the outcome Y with respect to a treatment decision D that precedes it, while our work considers counterfactuals of the outcome Y and the predictor \hat{Y} with respect to the protected attribute X , in the context of fair predictions, and thus offers a rather

different line of reasoning.

1.2 Organization & Contributions

In Sec. 2, we introduce important preliminary notions in causal inference, and the formal tools for specifying causal pathways C_1, \dots, C_k described above. Further, we discuss the notions of statistical and predictive parity, together with the impossibility result that separates them. In Sec. 2.2 we discuss how different causal variations with the statistical parity measure can be disentangled using an additive decomposition based on the previous result of (Zhang and Bareinboim 2018b). In Sec. 3, we develop a novel decomposition of the predictive parity measure in terms of the underlying causal mechanisms and discuss how this notion is in fact complementary to statistical parity through the concept of business necessity. In Sec. 3.1, we unify the theoretical findings by introducing a formal procedure that shows how to assess the legal doctrines of discrimination by leveraging both the concepts of causal predictive parity and causal statistical parity. In Sec. 4, we apply our approach in the context of criminal justice using the COMPAS dataset (Angwin et al. 2016), and demonstrate empirically the trade-off between SP and PP. Our key formal contributions are the following:

- We develop the first non-parametric decomposition of the predictive parity measure in terms of the underlying causal mechanisms (Thm. 1).
- Building on the previous result, we define a natural notion of causal predictive parity (Def. 5). We then develop a procedure (Alg. 1) for evaluating if a classifier satisfies the desired notions of causal statistical parity and causal predictive parity. This provides a unified framework for incorporating considerations of business necessity and sheds light on the impossibility theorem between statistical and predictive parity.

2 Background

We use the language of structural causal models (SCMs) as our basic semantical framework (Pearl 2000). A structural causal model (SCM) is a tuple $\mathcal{M} := \langle V, U, \mathcal{F}, P(u) \rangle$, where V, U are sets of endogenous (observables) 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} (Pearl 2000) 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 $X = x$. Building on the notion of a potential response, one can further define the notions of counterfactual and factual contrasts, given by:

Definition 1 (Contrasts (Plečko and Bareinboim 2022)). *Given an SCM \mathcal{M} , a contrast \mathcal{C} is any quantity of the form*

$$\mathcal{C}(C_0, C_1, E_0, E_1) = \mathbb{E}[y_{C_1} | E_1] - \mathbb{E}[y_{C_0} | E_0], \quad (1)$$




Figure 2: Standard Fairness Model.

where E_0, E_1 are observed (factual) events and C_0, C_1 are counterfactual clauses to which the outcome Y responds. A clause C is an intervention of the form $V' = v'$ where V' is a subset of the observable V . Furthermore, whenever

- (a) $E_0 = E_1$, the contrast \mathcal{C} is said to be counterfactual;
- (b) $C_0 = C_1$, the contrast \mathcal{C} is said to be factual.

For instance, the contrast $(C_0 = \{x_0\}, C_1 = \{x_1\}, E_0 = \emptyset, E_1 = \emptyset)$ corresponds to the *average treatment effect (ATE)* $\mathbb{E}[y_{x_1} - y_{x_0}]$. Similarly, the contrast $(C_0 = \{x_0\}, C_1 = \{x_1\}, E_0 = \{x_0\}, E_1 = \{x_0\})$ corresponds to the *effect of treatment on the treated (ETT)* $\mathbb{E}[y_{x_1} - y_{x_0} | x_0]$. Many other important causal quantities can be represented as contrasts, as exemplified later on.

Throughout this manuscript, we assume a specific cluster causal diagram \mathcal{G}_{SFM} known as the standard fairness model (SFM) (Plečko and Bareinboim 2022) 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 also encodes the assumptions typically used in the causal inference literature about the lack of hidden confounding¹. We next introduce the key notions and results from the fair ML literature needed for our discussion.

2.1 Statistical & Predictive Parity Notions

The notions of statistical parity (labeled SP from now on) and predictive parity (labeled PP) are defined as follows:

Definition 2 (Statistical (Darlington 1971) and Predictive (Chouldechova 2017) Parity). *Let X be the protected attribute, Y the true outcome, and \hat{Y} the outcome predictor. The predictor \hat{Y} satisfies statistical parity (SP) with respect to X if $X \perp\!\!\!\perp \hat{Y}$, or equivalently if the statistical parity measure (SPM, for short) satisfies:*

$$SPM_{x_0, x_1}(\hat{y}) := P(\hat{y} | x_1) - P(\hat{y} | x_0) = 0. \quad (2)$$

Further, \hat{Y} satisfies predictive parity (PP) with respect to X, Y if $Y \perp\!\!\!\perp X | \hat{Y}$, or equivalently if $\forall \hat{y}$ we have

$$PPM_{x_0, x_1}(y | \hat{y}) := P(y | x_1, \hat{y}) - P(y | x_0, \hat{y}) = 0. \quad (3)$$

¹Partial identification (bounding) of effects can be used to relax these assumptions (Zhang, Tian, and Bareinboim 2022).




Figure 3: Standard Fairness Model with $Z = \emptyset$ from Thm. 1, extended with the predictor \hat{Y} .

Statistical parity ($\hat{Y} \perp\!\!\!\perp X$) requires that the predictor \hat{Y} contains no information about the protected attribute X . In contrast to this, the notion of predictive parity ($Y \perp\!\!\!\perp X | \hat{Y}$), requires that \hat{Y} should “exhaust”, or capture all the variations of X coming into the outcome Y in the current real world. Chouldechova (2017) introduced predictive parity for a binary predictor \hat{Y} , whereas the condition was termed calibration for a continuous classification score. In this paper, we are agnostic to this distinction, and almost all of the results are unaffected by this difference, as discussed in detail in Appendix D. Therefore, we simply use the term predictive parity throughout. An important known result on predictive parity is the following (see proof in Appendix C):

Proposition 1 (PP and Efficient Learning). *Let \mathcal{M} be an SCM compatible with the Standard Fairness Model (SFM). Suppose that the predictor \hat{Y} is based on the features X, Z, W . Suppose also that \hat{Y} is an efficient learner, i.e.:*

$$\hat{Y}(x, z, w) = P(y | x, z, w). \quad (4)$$

Then \hat{Y} satisfies predictive parity w.r.t. X and Y .

Prop. 1 that pertains to a continuous score-based predictor shows that PP is expected to hold whenever we learn the true conditional distribution $P(y | x, z, w)$. This is in stark contrast with the SP notion. For the SP to hold, the predictor \hat{Y} should not be associated at all with the attribute X , whereas for PP to hold, \hat{Y} should contain all variations coming from X . Perhaps unsurprisingly after this discussion, the following impossibility result holds:

Proposition 2 (SP and PP impossibility (Barocas, Hardt, and Narayanan 2017)). *Statistical parity ($\hat{Y} \perp\!\!\!\perp X$) and predictive parity ($Y \perp\!\!\!\perp X | \hat{Y}$) are mutually exclusive except in degenerate cases when $Y \perp\!\!\!\perp X$.*

The above theorem might lead the reader to believe that SP and PP criteria come from two different realities, and bear no relation to each other. After all, the theorem states that it is not possible for the predictor \hat{Y} to include *all variations* of X coming into Y and simultaneously include *no variations* of X coming into Y . This realization is the starting point of our discussion in the rest of the manuscript.

2.2 Statistical Parity Decomposition

We continue by analyzing the decomposition of the statistical parity measure (also known as parity gap, or total variation), which is commonly used to determine if statistical parity is satisfied. The causal decomposition of the SPM requires the usage of causal measures known as counterfactual direct, indirect, and spurious effects, which are defined as:

Definition 3 (Counterfactual Causal Measures). *The counterfactual- $\{direct, indirect, spurious\}$ effects of X on \hat{Y} are defined as follows*

$$Ctf-DE_{x_0, x_1}(\hat{y} | x) = P(\hat{y}_{x_1, W_{x_0}} | x) - P(\hat{y}_{x_0} | x) \quad (5)$$

$$Ctf-IE_{x_1, x_0}(\hat{y} | x) = P(\hat{y}_{x_1, W_{x_0}} | x) - P(\hat{y}_{x_1} | x) \quad (6)$$

$$Ctf-SE_{x_1, x_0}(\hat{y}) = P(\hat{y}_{x_1} | x_0) - P(\hat{y}_{x_1} | x_1). \quad (7)$$

The measures capture the variations from X to \hat{Y} going through: (i) the direct mechanism $X \rightarrow \hat{Y}$; (ii) the indirect mechanism $X \rightarrow W \rightarrow \hat{Y}$; (iii) the spurious mechanisms $X \leftarrow Z \rightarrow \hat{Y}$, and $X \leftarrow Z \rightarrow W \rightarrow \hat{Y}$, respectively (see Fig. 2). Based on the defined measures, $SPM_{x_0, x_1}(\hat{y})$ admits an additive decomposition, first obtained in (Zhang and Bareinboim 2018b), given in the following theorem²:

Proposition 3 (Causal Decomposition of Statistical Parity (Zhang and Bareinboim 2018b)). *The statistical parity measure admits a decomposition into its direct, indirect, and spurious variations:*

$$SPM_{x_0, x_1}(\hat{y}) = Ctf-DE_{x_0, x_1}(\hat{y} | x_0) - Ctf-IE_{x_1, x_0}(\hat{y} | x_0) - Ctf-SE_{x_1, x_0}(\hat{y}). \quad (8)$$

This result shows how we can disentangle direct, indirect, and spurious variations within the SPM. We emphasize the importance of this result in the context of assessing the legal doctrines of discrimination. If a causal pathway (direct, indirect, or spurious) does not lie in the business necessity set, then the corresponding counterfactual measure (Ctf-DE, IE, or SE) needs to equal 0. This notion follows directly from the literature on path-specific fairness (Nabi and Shpitser 2018; Chiappa 2019; Wu et al. 2019), and we refer to it as *causal statistical parity*:

Definition 4 (Causal Statistical Parity). *We say that \hat{Y} satisfies causal statistical parity with respect to the protected attribute X if*

$$Ctf-DE_{x_0, x_1}(\hat{y} | x_0) = Ctf-IE_{x_1, x_0}(\hat{y} | x_0) \quad (9)$$

$$= Ctf-SE_{x_1, x_0}(\hat{y}) = 0. \quad (10)$$

In practice, causal statistical parity can be a strong requirement, but the notion can be easily relaxed to include only a subset of the Ctf- $\{DE, IE, \text{ or } SE\}$ measures, under BN requirements.

3 Predictive Parity Decomposition

After discussing the decomposition of the SPM, our aim is to obtain a causal understanding of the predictive parity criterion, $Y \perp\!\!\!\perp X | \hat{Y}$. To do so, we derive a formal decomposition result of the PP measure that involves both Y and \hat{Y} , shown in the following theorem:

Theorem 1 (Causal Decomposition of Predictive Parity). *Let \mathcal{M} be an SCM compatible with the causal graph in*

²We note that the same decomposition can also be applied for the true outcome Y instead of the predictor \hat{Y} , as will be relevant in later sections.

Fig. 3 (i.e., SFM with $Z = \emptyset$). Then, it follows that the $PPM_{x_0, x_1}(y | \hat{y}) = P(y | x_1, \hat{y}) - P(y | x_0, \hat{y})$ can be decomposed into its causal and spurious anti-causal variations as:

$$PPM_{x_0, x_1}(y | \hat{y}) = P(y_{x_1} | x_1, \hat{y}) - P(y_{x_0} | x_1, \hat{y}) + P(y_{x_0} | \hat{y}_{x_1}) - P(y_{x_0} | \hat{y}_{x_0}). \quad (11)$$

Thm. 1 offers a non-parametric decomposition result of the predictive parity measure that can be applied to any SCM compatible with the graph in Fig. 3. In Appendix A.1 we provide a proof of the theorem (together with the proof of Cor. 1 stated below), and in Appendix A.2 we perform an empirical study to verify the decomposition result of the theorem (for supplements see full paper version at <https://arxiv.org/pdf/2306.05059.pdf>). For the additional special case of linear SCMs, the terms appearing in the decomposition in Eq. 11 can be computed explicitly:

Corollary 1 (Causal Decomposition of Predictive Parity in the Linear Case). *Under the additional assumptions that (i) the SCM \mathcal{M} is linear and Y is continuous; (ii) the learner \hat{Y} is efficient, then*

$$\mathbb{E}(y_{x_1} | x_1, \hat{y}) - \mathbb{E}(y_{x_0} | x_1, \hat{y}) = \alpha_{XW}\alpha_{WY} + \alpha_{XY}, \quad (12)$$

$$\mathbb{E}(y_{x_0} | x_1, \hat{y}_{x_1}) - \mathbb{E}(y_{x_0} | x_1, \hat{y}_{x_0}) = -\alpha_{XW}\alpha_{WY} - \alpha_{XY}, \quad (13)$$

where $\alpha_{V_i V_j}$ is the linear coefficient between V_i and V_j .

We now carefully unpack the key insight from Thm. 1. In particular, we showed that in the case of an SFM with $Z = \emptyset$ ³ the predictive parity measure can be written as:

$$PPM = \underbrace{P(y_{x_1} | x_1, \hat{y}) - P(y_{x_0} | x_1, \hat{y})}_{\text{Term (I) causal}} + \underbrace{P(y_{x_0} | \hat{y}_{x_1}) - P(y_{x_0} | \hat{y}_{x_0})}_{\text{Term (II) reverse-causal spurious}}. \quad (14)$$

Term (I) capturing causal variations can be expanded as

$$\sum_u \left[\underbrace{y_{x_1}(u) - y_{x_0}(u)}_{\text{unit-level difference}} \right] \underbrace{P(u | x_1, \hat{y})}_{\text{posterior}}. \quad (15)$$

The expansion shows us that the term is a weighted average of unit-level differences $y_{x_1}(u) - y_{x_0}(u)$. Each unit-level difference measures the causal effect of a transition $x_0 \rightarrow x_1$ on Y . The associated weights are given by the posterior $P(u | x_1, \hat{y})$, which determines the probability mass corresponding to a unit u within the set of all units compatible with x_1, \hat{y} . Therefore, the term represents an average causal effect of X on Y for a specific group of units chosen by the x_1, \hat{y} conditioning. Interestingly, after selecting the units by x_1, \hat{y} , the effect in Term (I) *no longer depends on the constructed predictor \hat{Y}* , but only on the true mechanism of Y that already exists in the real world, i.e., it is not under the

³We remark that the essence of the argument is unchanged in the case with $Z \neq \emptyset$, but handling this case limits the clarity of presentation.

Algorithm 1: Business Necessity Cookbook

```

1: Input: data  $\mathcal{D}$ , BN-Set  $\subseteq \{\text{DE}, \text{IE}, \text{SE}\}$ , predictor  $\hat{Y}$ 
2: for  $\text{CE} \in \{\text{DE}, \text{IE}, \text{SE}\}$  do
3:   if  $\text{CE} \in \text{BN}$  then
4:     Compute the effects  $\text{Ctf-CE}(y)$ ,  $\text{Ctf-CE}(\hat{y})$ 
5:     Assert that  $\text{Ctf-CE}(y) = \text{Ctf-CE}(\hat{y})$ , otherwise
       FAIL
6:   else
7:     Compute the effect  $\text{Ctf-CE}(\hat{y})$ 
8:     Assert that  $\text{Ctf-CE}(\hat{y}) = 0$ , otherwise FAIL
9:   end if
10: end for
11: if not FAIL then
12:   return SUCCESS
13: end if
14: Output: SUCCESS or FAIL of ensuring that disparate
       impact and treatment hold under BN.

```

control of the predictor designer. The additional linear result in Cor. 1 may also help the reader ground this idea, since it shows that Term (I) indeed captures the causal effect, which, in the linear case, can be obtained using the path-analysis of (Wright 1934).

Thus, to achieve the criterion $\text{PPM} = 0$, the second term needs to be exactly the reverse of the causal effect, captured by the spurious variations induced by changing $\hat{y}_{x_0} \rightarrow \hat{y}_{x_1}$ in the selection of units. The second term, which is in the control of the predictor \hat{Y} designer, needs to cancel out the causal effect measured by the first term for PPM to vanish. Therefore, we see that achieving predictive parity is about constructing \hat{Y} in a way that reflects the causal effect of X on Y , across various groups of units. This key observation motivates a novel definition of *causal predictive parity*:

Definition 5 (Causal Predictive Parity). *Let \hat{Y} be a predictor of the outcome Y , and let X be the protected attribute. Then, \hat{Y} is said to satisfy causal predictive parity (CPP) with respect to a counterfactual contrast (C_0, C_1, E, E) if*

$$\mathbb{E}[y_{C_1} | E] - \mathbb{E}[y_{C_0} | E] = \mathbb{E}[\hat{y}_{C_1} | E] - \mathbb{E}[\hat{y}_{C_0} | E]. \quad (16)$$

Furthermore, \hat{Y} is said to satisfy CPP with respect to a factual contrast (C, C, E_0, E_1) if

$$\mathbb{E}[y_C | E_1] - \mathbb{E}[y_C | E_0] = \mathbb{E}[\hat{y}_C | E_1] - \mathbb{E}[\hat{y}_C | E_0]. \quad (17)$$

The intuition behind the notion of causal predictive parity captures the intuition behind predictive parity. If a contrast \mathcal{C} describes some amount of variation in the outcome Y , then it should describe the same amount of variation in the predicted outcome \hat{Y} . For any of the contrasts $\text{Ctf-}\{\text{DE}, \text{IE}, \text{SE}\}$ corresponding to a causal pathway, causal predictive parity would require that $\mathcal{C}(X, \hat{Y}) = \mathcal{C}(X, Y)$.

3.1 Reconciling Statistical and Predictive Parity

We now tie the notions of statistical and predictive parity through the concept of *business necessity*. In particular, if

a contrast \mathcal{C} is associated with variations that are not in the business necessity set, then the value of this contrast should be $\mathcal{C}(X, \hat{Y}) = 0$, following the intuition of causal statistical parity from Def. 4. However, if the variations associated with the contrast *are* in the business necessity set, then the value of that contrast should be equal for the predictor to the value for the true outcome

$$\mathcal{C}(X, \hat{Y}) = \mathcal{C}(X, Y), \quad (18)$$

following the intuition of causal predictive parity. Combining these two notions through business necessity results in Alg. 1. The algorithm requires the user to compute the measures

$$\text{Ctf-}\{\text{DE}, \text{IE}, \text{SE}\}(y), \text{Ctf-}\{\text{DE}, \text{IE}, \text{SE}\}(\hat{y}). \quad (19)$$

Importantly, under the SFM, these measures are *identifiable* from observational data:

Proposition 4. *Under the assumptions of the standard fairness model in Fig. 2, the causal measures of fairness $\text{Ctf-}\{\text{DE}, \text{IE}, \text{SE}\}(y)$, $\text{Ctf-}\{\text{DE}, \text{IE}, \text{SE}\}(\hat{y})$ are identifiable from observational data, that is, they can be computed uniquely from the observational distribution $P(V)$.*

The explicit identification expression for each of the measures is given in Appendix B. The above result guarantees that the procedure of Alg. 1 is applicable to practical data analysis, in a fully non-parametric nature.

Further, for each of the DE, IE, and SE effects, the user needs to determine whether the causal effect (CE) in question falls into the business necessity set. If yes, then the algorithm asserts that

$$\text{Ctf-CE}(y) = \text{Ctf-CE}(\hat{y}). \quad (20)$$

In the other case, when the causal effect is not in the business necessity set, the algorithm asserts that

$$\text{Ctf-CE}(\hat{y}) = 0. \quad (21)$$

Alg. 1 is the key result of this paper, and provides a practical approach to incorporating ideas of business necessity through a causal lens. We also remark that it is written in its population level version, in which the causal effects are estimated perfectly with no uncertainty. In the finite sample case, one needs to perform hypothesis testing to see whether the effects differ. If one is also interested in constructing a new fair predictor before using Alg. 1 (instead of testing an existing one), one may use tools for causally removing discrimination, such as (Chiappa 2019) or (Plečko and Meinshausen 2020; Plečko, Bennett, and Meinshausen 2021). In Appendix E we show formally that a predictor \hat{Y} satisfying the conditions of Alg. 1 can be constructed for any choice of the BN-set.

Finally, we also remark that Alg. 1 generalizes to arbitrary causal diagrams \mathcal{G} . For the SFM, we can only distinguish three causally different pathways (direct, indirect, spurious), whereas for arbitrary graphs we may consider various collections of path-specific effects. If the contrast \mathcal{C}_π captures variations along a pathway π , then Alg. 1 would simply test the equality $\mathcal{C}_\pi(X, \hat{Y}) = \mathcal{C}_\pi(X, Y)$ if π is in the BN set. For larger graphs, however, the number of BN paths may scale unfavorably, representing a computational challenge.




Figure 4: SFM for the COMPAS dataset.

4 Experiments

We now apply Alg. 1 to the COMPAS dataset (Angwin et al. 2016), as described in the following example. Additionally, in Appendix F, we also analyze the Census 2018 dataset.

Courts in Broward County, Florida use machine learning algorithms, developed by Northpointe, to predict whether individuals released on parole are at high risk of re-offending within 2 years (Y). The algorithm is based on the demographic information Z (Z_1 for gender, Z_2 for age), race X (x_0 denoting White, x_1 Non-White), juvenile offense counts J , prior offense count P , and degree of charge D .

We construct the standard fairness model (SFM) for this example, which is shown in Fig. 4. The bidirected arrow between X and $\{Z_1, Z_2\}$ indicates possible co-variations of race with age and sex, which may not be causal in nature⁴. Furthermore, $\{Z_1, Z_2\}$ are the confounders, not causally affected by race X . The set of mediators $\{J, P, D\}$, however, may be affected by race, due to an existing societal bias in policing and criminal justice. Finally, all of the above-mentioned variables may influence the outcome Y .

Having access to data from Broward County, and equipped with Alg. 1, we want to prove that the recidivism predictions produced by Northpointe (labeled \hat{Y}^{NP}) violate legal doctrines of anti-discrimination. Suppose that in an initial hearing, the Broward County district court determines that the direct and indirect effects are not in the business necessity set, while the spurious effect is. In words, gender (Z_1) and age (Z_2) are allowed to be used to distinguish between the minority and majority groups when predicting recidivism, while other variables are not. If Northpointe’s predictions are found to be discriminatory, we are required by the court to produce better, non-discriminatory predictions.

In light of this information, we proceed as follows (see <https://github.com/dplecko/sp-to-pp/blob/main/sp-pp-compas.R> for source code). We first obtain a causally fair predictor \hat{Y}^{FP} using the fairadapt package, which performs optimal transport in order to remove the direct and indirect effects of the protected attribute from the prediction. Additionally, we also add predictors \hat{Y}^{SP} (obtained using group-specific thresholds), \hat{Y}^{PP} (the unconstrained predictor using random forests) that satisfy statistical and predictive parity, respectively. Then, we compute the counterfactual

⁴The causal model is non-committal regarding the complex historical/social processes that lead to such co-variations.

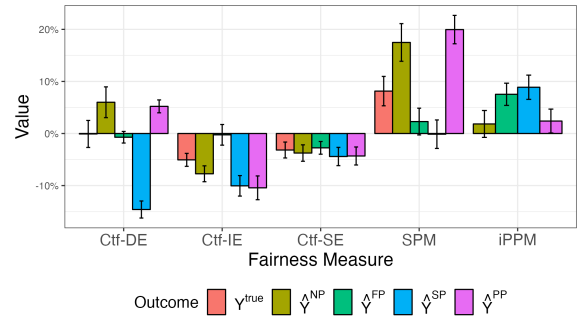


Figure 5: Fairness measures on COMPAS.

tual causal measures of fairness for the true outcome Y , Northpointe’s predictions \hat{Y}^{NP} , fair predictions \hat{Y}^{FP} , and \hat{Y}^{SP} , \hat{Y}^{PP} (see Fig. 5). For the direct effect, we have:

$$\text{Ctf-DE}_{x_0, x_1}(y | x_0) = -0.08\% \pm 2.59\%, \quad (22)$$

$$\text{Ctf-DE}_{x_0, x_1}(\hat{y}^{NP} | x_0) = 6\% \pm 2.96\%, \quad (23)$$

$$\text{Ctf-DE}_{x_0, x_1}(\hat{y}^{FP} | x_0) = -0.72\% \pm 1.11\%. \quad (24)$$

The indicated 95% confidence intervals are computed using repeated bootstrap repetitions of the dataset. Since the direct effect is not in the business necessity set, Northpointe’s predictions clearly violate the disparate treatment doctrine (yellow bar for the Ctf-DE measure in Fig. 5). Our predictions, however, do not exhibit a statistically significant direct effect of race on the outcome, so they do not violate the criterion (green bar). Next, for the indirect effects, we obtain:

$$\text{Ctf-IE}_{x_1, x_0}(y | x_0) = -5.06\% \pm 1.24\%, \quad (25)$$

$$\text{Ctf-IE}_{x_1, x_0}(\hat{y}^{NP} | x_0) = -7.73\% \pm 1.53\%, \quad (26)$$

$$\text{Ctf-IE}_{x_1, x_0}(\hat{y}^{FP} | x_0) = -0.25\% \pm 1.98\%. \quad (27)$$

Once again, the indirect effect, which is not in the business necessity set, is different from 0 for the Northpointe’s predictions (violating disparate impact, see yellow bar for Ctf-IE in Fig. 5), but not statistically different from 0 for our predictions (green bar). Interestingly, the indirect effect is different from 0 for the true outcome (red bar), indicating a bias in the current real world. Finally, for spurious effects we obtain

$$\text{Ctf-SE}_{x_1, x_0}(y) = -3.17\% \pm 1.53\%, \quad (28)$$

$$\text{Ctf-SE}_{x_1, x_0}(\hat{y}^{NP}) = -3.75\% \pm 1.58\%, \quad (29)$$

$$\text{Ctf-SE}_{x_1, x_0}(\hat{y}^{FP}) = -2.75\% \pm 1.22\%. \quad (30)$$

Since the spurious effect is in the business necessity set, and the hypotheses $\text{SE}(y) = \text{SE}(\hat{y}^{NP})$, $\text{SE}(y) = \text{SE}(\hat{y}^{FP})$ are not rejected (p-values 0.30, 0.32, respectively), no violations with respect to spurious effects are found.

We conclude that Northpointe’s predictions \hat{Y}^{NP} violate the legal doctrines of fairness, while our predictions \hat{Y}^{FP} do not. Importantly, tying back to the original discussion that motivated our approach, Northpointe’s predictions \hat{Y}^{NP} are further away from statistical parity than our predictions \hat{Y}^{FP} according to the SPM (see SPM column in

Fig. 5), while at the same time better calibrated according to the integrated PP measure (iPPM) that averages the PP measures from Eq. 3 across different values of \hat{y} (see iPPM column in the figure). This demonstrates the trade-off between statistical and predictive parity through business necessity.

The choice of the business necessity set in the example above, $BN = \{W\}$, was arbitrary. In general, the BN set can be any of $\emptyset, Z, W, \{Z, W\}$, which we explore next.

Statistical vs. Predictive Parity Pareto Frontier. We now investigate other possible choices of the BN set, namely BN sets $\emptyset, Z, W, \{Z, W\}$. Based on the theoretical analysis of Sec. 3, we expect that different choices of the BN set will lead to different trade-offs between SP and PP.

In particular, with the BN set being empty, any variation from X to Y would be considered discriminatory, so one would expect to be close to statistical parity. In contrast, if the BN set is $\{Z, W\}$, then all variations (apart from the direct effect) are considered as non-discriminatory, so one would expect to be closer to satisfying predictive parity. The remaining options, $BN = W$, and $BN = Z$ should interpolate between these two extremes.

Based on the above intuition, we proceed as follows. For each choice of the business necessity set,

$$BN \in \{\emptyset, Z, W, \{Z, W\}\}, \quad (31)$$

we compute the adjusted version of the data again using the `fairadapt` package, with the causal effects in the complement of the BN set being removed. That is, if $BN = W$, then our procedure removes the spurious effect, but keeps the indirect effect intact (other choices of the BN set are similar). Therefore, for each BN set, we obtain an appropriately adjusted predictor \hat{Y}_{BN}^{FP} , and in particular compute the predictors $\hat{Y}_{\emptyset}^{FP}, \hat{Y}_Z^{FP}, \hat{Y}_W^{FP}$, and $\hat{Y}_{\{Z,W\}}^{FP}$. We note that the fair predictor \hat{Y}^{FP} from the first part of the example is the predictor with the BN set equal to Z , i.e., it corresponds to \hat{Y}_Z^{FP} . For each \hat{Y}_{BN}^{FP} we compute the SPM, iPPM, and area under receiver operator characteristic (AUC):

$$SPM_{x_0, x_1}(\hat{y}_{BN}^{FP}), iPPM_{x_0, x_1}(\hat{y}_{BN}^{FP}), AUC(\hat{y}_{BN}^{FP}), \quad (32)$$

across 10 different repetitions of the adjustment procedure that yields \hat{Y}_{BN}^{FP} . For each business necessity set, this allows us to compute the SPM (measuring statistical parity), and iPPM (measuring predictive parity).

The results of the experiment are shown in Fig. 6, where the error bars indicate the standard deviation over different repetitions. As predicted by our theoretical analysis, the choice of $BN = \emptyset$ yields the lowest SPM, but the largest iPPM. Conversely, $BN = \{Z, W\}$ yields the lowest iPPM, but the largest SPM. The BN sets Z, W interpolate between the two notions, but the data indicates the spurious effect explained by Z does not have a major contribution. Fig. 6, therefore, shows a trade-off between statistical parity on the one hand, and predictive parity and accuracy on the other. The trade-off is captured through different business necessity options, and gives an empirical validation of the hypothesized spectrum of fairness notions in Fig. 1.

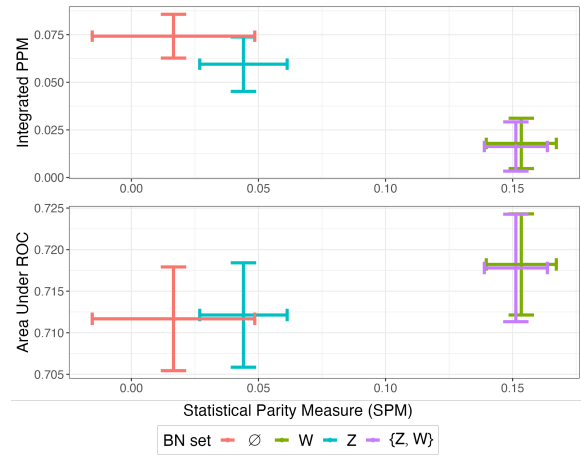


Figure 6: SP vs. PP and AUC Pareto frontiers on COMPAS.

5 Conclusions

The literature in fair ML is abundant with fairness measures (Corbett-Davies and Goel 2018), many of which are mutually incompatible. Nonetheless, it is doubtful that each of these measures corresponds to a fundamentally different ethical conception of fairness. The multitude of possible approaches to quantifying discrimination makes the consensus on an appropriate notion of fairness unattainable. Further, the impossibility results between different measures may be discouraging to data scientists who wish to quantify and remove discrimination, but are immediately faced with a choice of which measure they wish to subscribe to. Our work falls under the rubric of analyzing impossibility results through a causal lens (Makar and D’Amour 2022).

In particular, we focus on the impossibility of simultaneously achieving SP and PP. As our discussion shows, the guiding idea behind SP is that variations transmitted along causal pathways from the protected attribute to the predictor should equal 0, i.e., the decision should not depend on the protected attribute through the causal pathway in question (Def. 4). Complementary to this notion, and based on Thm. 1, the guiding principle behind PP is that variations transmitted along a causal pathway should be equal for the predictor as they are for the outcome *in the real world* (Def. 5). PP will therefore be satisfied when the BN set includes all variations coming from X to Y , while SP will be satisfied when the BN set is empty. The choice of the BN set interpolates between SP and PP forming a spectrum of fairness notions (see Fig. 1), in a way that can be formally assessed based on Alg. 1.

Therefore, our work complements the previous literature by reconciling the impossibility result between SP and PP (Barocas, Hardt, and Narayanan 2017). Furthermore, it complements the existing literature on path-specific notions of fairness (Kilbertus et al. 2017; Wu et al. 2019; Chiappa 2019), which does not consider the true outcome Y and the predictor \hat{Y} simultaneously, and does not explicitly specify which levels of discrimination are deemed acceptable along causal pathways in the BN set, as described in Alg. 1.

Acknowledgements

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

References

- 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.; Hardt, M.; and Narayanan, A. 2017. Fairness in machine learning. *Nips tutorial*, 1: 2017.
- Barocas, S.; and Selbst, A. D. 2016. Big data’s disparate impact. *Calif. L. Rev.*, 104: 671.
- 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.
- Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.
- 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.; and Goel, S. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*.
- Coston, A.; Mishler, A.; Kennedy, E. H.; and Chouldechova, A. 2020. Counterfactual risk assessments, evaluation, and fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 582–593.
- Darlington, R. B. 1971. Another Look At “Cultural Fairness”. *Journal of Educational Measurement*, 8(2): 71–82.
- Detrixhe, J.; and Merrill, J. B. 2019. The fight against financial advertisers using Facebook for digital redlining.
- Harwell, D. 2019. Federal study confirms racial bias of many facial-recognition systems, casts doubt on their expanding use. <https://www.washingtonpost.com/technology/2019/12/19/federal-study-confirms-racial-bias-many-facial-recognition-systems-casts-doubt-their-expanding-use/>.
- Kilbertus, N.; Rojas-Carulla, M.; Parascandolo, G.; Hardt, M.; Janzing, D.; and Schölkopf, B. 2017. Avoiding discrimination through causal reasoning. *arXiv preprint arXiv:1706.02744*.
- Makar, M.; and D’Amour, A. 2022. Fairness and robustness in anti-causal prediction. *arXiv preprint arXiv:2209.09423*.
- Nabi, R.; and Shpitser, I. 2018. Fair inference on outcomes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press. 2nd edition, 2009.
- Plečko, D.; and Bareinboim, E. 2022. Causal Fairness Analysis. *arXiv preprint arXiv:2207.11385*. (To appear in *Foundations and Trends in Machine Learning*).
- Plečko, D.; Bennett, N.; and Meinshausen, N. 2021. fairadapt: Causal Reasoning for Fair Data Pre-processing. *arXiv preprint arXiv:2110.10200*.
- Plečko, D.; and Meinshausen, N. 2020. Fair data adaptation with quantile preservation. *Journal of Machine Learning Research*, 21: 242.
- Rajkumar, A.; Hardt, M.; Howell, M. D.; Corrado, G.; and Chin, M. H. 2018. Ensuring fairness in machine learning to advance health equity. *Annals of internal medicine*, 169(12): 866–872.
- Tian, J.; and Pearl, J. 2000. Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1): 287–313.
- U.S. Census Bureau. 2018. Public Use Microdata Sample (PUMS), 2018 American Community Survey. <https://www.census.gov/programs-surveys/acs/microdata/documentation.html>. Accessed: 18-07-2021.
- Wright, S. 1934. The method of path coefficients. *The annals of mathematical statistics*, 5(3): 161–215.
- 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.
- 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*.