



HAL
open science

On the Impossibility of non-Trivial Accuracy in Presence of Fairness Constraints

Carlos Pinzón, Catuscia Palamidessi, Pablo Piantanida, Frank Valencia

► **To cite this version:**

Carlos Pinzón, Catuscia Palamidessi, Pablo Piantanida, Frank Valencia. On the Impossibility of non-Trivial Accuracy in Presence of Fairness Constraints. Proceedings of the AAAI 36th Conference on Artificial Intelligence, Feb 2022, Vancouver / Virtual, Canada. pp.7993-8000, 10.1609/aaai.v36i7.20770 . hal-03452324

HAL Id: hal-03452324

<https://hal.science/hal-03452324>

Submitted on 26 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

On the impossibility of non-trivial accuracy under fairness constraints

Carlos Pinzón^{1,4}, Catuscia Palamidessi^{1,4}, Pablo Piantanida^{2,3}, and Frank Valencia^{2,4}

¹Inria, France

²CNRS, France

³Laboratoire des Signaux et Systèmes (L2S), CentraleSupélec, Université Paris-Saclay

⁴LIX, École Polytechnique, Institut Polytechnique de Paris

Abstract—One of the main concerns about fairness in machine learning (ML) is that, in order to achieve it, one may have to trade off some accuracy. To overcome this issue, Hardt et al. proposed the notion of equality of opportunity (EO), which is compatible with maximal accuracy when the target label is deterministic with respect to the input features.

In the probabilistic case, however, the issue is more complicated: It has been shown that under differential privacy constraints, there are data sources for which EO can only be achieved at the total detriment of accuracy, in the sense that a classifier that satisfies EO cannot be more accurate than a trivial (i.e., constant) classifier. In our paper we strengthen this result by removing the privacy constraint. Namely, we show that for certain data sources, the most accurate classifier that satisfies EO is a trivial classifier. Furthermore, we study the trade-off between accuracy and EO loss (opportunity difference), and provide a sufficient condition on the data source under which EO and non-trivial accuracy are compatible.

I. INTRODUCTION

During the last decade, the intersection between machine learning and social discrimination has gained considerable attention from the academia, the industry and the public in general. A similar trend occurred before between machine learning and privacy, and even the three fields have been studied together recently Pujol et al. [2020], Cummings et al. [2019], Kearns and Roth [2019], Agarwal [2020].

Fairness, has proven to be harder to conceptualize than privacy, for which differential privacy has become the de-facto definition. Fairness is subjective and laws vary between countries. Even in academia, depending on the application, the words fairness and bias have different meanings Crawford [2017]. The current general consensus is that fairness can not be summarized into a unique universal definition; and for the most popular definitions, several trade-offs, implementation difficulties and impossibility theorems have been found Kleinberg et al. [2016], Chouldechova [2017]. One such definition of fairness is equal-opportunity Hardt et al. [2016].

To contrast equal-opportunity (EO) with accuracy, we borrow the notion of trivial accuracy from the work by Cumming et al. Cummings et al. [2019]. A *non-trivial* classifier is one that has higher accuracy than any constant classifier. Since constant classifiers are independent of the input, trivial accuracy determines a very low minimum performance level that any correctly trained classifier should overcome. Yet as

shown in related works Cummings et al. [2019], Agarwal [2020], under the simultaneous constraints of differential privacy and equal-opportunity, it is impossible to have non-trivially accurate classifiers.

In this paper, we complement these existing theorems by showing that even without the assumption of differential-privacy, there are distributions for which equal-opportunity implies trivial accuracy. This is only possible, however, under a probabilistic data source, i.e. when the correct label for a given input is not necessarily deterministic.

Probability plays two different roles in this paper. On the one hand, we allow classifiers to be probabilistic. This is important because sometimes randomness is the only fair way to distribute an indivisible limited resource. Indeed, equal-opportunity and equal-odds demand probabilistic predictors in some scenarios. On the other hand, we consider the possibility that the data source may be probabilistic. This decision is motivated by two different reasons:

- 1) It enables a more realistic and general representation of reality: one in which the information in the input may be insufficient to conclude definitely the yes-no decision, or in which real-life constraints force the decision to be different for identical inputs.
- 2) It provides a more general, yet simple, perspective for understanding the trade-off between fairness and accuracy. Also, it can take into account that in practice, input datasets are a noisy (thus probabilistic) approximation of reality.

Our contributions are the following.

- 1) We prove that for certain probabilistic distributions, no predictor can achieve EO and non-trivial accuracy simultaneously.
- 2) We provide a sufficient condition that guarantees compatibility between non-trivial accuracy and EO.
- 3) We explain how to modify existing results that assume deterministic data sources to the probabilistic case:
 - a) We prove that for certain distributions, the Bayes classifier does not satisfy EO. As a consequence, in these cases, EO can only be achieved by trading-off some accuracy.
 - b) We give sufficient and necessary conditions for non-trivially accurate predictors to exist.

- 4) We prove and depict several algebraic and geometric properties about the feasible region in the plane of opportunity-difference versus error.

II. RELATED WORKS

Our paper is strongly related to the following two works that consider a randomized learning algorithm guaranteeing (exact) equal-opportunity and also satisfying differential privacy: Cummings et al. [2019] shows that, for certain distributions, these constraints imply trivial accuracy. Agarwal [2020] proves the same claim for any arbitrary distribution and for non-exact equal-opportunity, i.e. bounded opportunity-difference. It also highlights that, although there appears to be an error in the proof of Cummings et al. [2019], the statement is still correct. In contrast, in this paper, we prove the existence of particular distributions in which trivial accuracy is implied directly from the (exact) equal-opportunity constraint, without any differential privacy assumption.

Another work about differential-privacy and fairness is Pujol et al. [2020]. Unlike our work and that of Agarwal [2020], which are limited to equal-opportunity (and equal-odds), they use a different notion of fairness, and this allows them to provide examples in which privacy and fairness are not necessarily in a trade-off.

There are also several works that focus on the feasibility of fairness constraints. In Kleinberg et al. [2016], it is shown that several different fairness notions can not hold simultaneously, except for exceptional cases. In Lipton et al. [2018], it is shown that the two main legal notions of discrimination are in conflict for some scenarios. In particular when impact parity and treatment parity are imposed, the learned model seems to decide based on irrelevant attributes. These works differ from our work in that they reveal contradictions that arise when different notions of fairness are imposed together. In contrast, we fix a single fairness notion, equal-opportunity, and reveal a contradiction that may arise with non-trivial accuracy.

Lastly, our geometric plots differ from those in the seminal paper of equal-opportunity Hardt et al. [2016]. Graphically, their analysis is carried out over ROC curves while we plot directly the two metrics of interest. In this sense, we provide a complementary geometric perspective for analyzing equal-opportunity and accuracy together.

III. PRELIMINARIES

The notation described in this section is summarized in Table I.

We consider the problem of binary classification with a binary protected feature. *Protected features*, also called sensible attributes or sensible features, are input features that represent race, gender, religion, nationality, age, or any other variable that could be used to discriminate against a group of people. A feature may be considered as a protected feature in some contexts and not in others, depending on whether the classification task should ideally consider that feature or not. For our purposes, we assume the simple and fundamental case in which there is a single protected attribute that can only take two values, e.g. man or woman, or, religious or non-religious.

(X, A, Y)	Data source
X	Non-protected feature vector in \mathbb{R}^d
A	Protected feature in $\{0, 1\}$
Y	Target label in $\{0, 1\}$
Q, q	Soft target label $Q \stackrel{\text{def}}{=} \mathbb{E}[Y X, A]$
π	Distribution of (X, A)
(π, q)	Distribution of (X, A, Y)
\hat{Q}, \hat{q}	Predictor $\hat{Q} = \hat{q}(X, A) = \mathbb{E}[\hat{Y} X, A]$
\hat{Y}	Predicted label in $\{0, 1\}$
\mathcal{Q}	Set of all predictors
$\text{acc}(\hat{Q})$	Accuracy of \hat{Q} : $\mathbb{P}[\hat{Y}=Y]$
$\text{oppDiff}(\hat{Q})$	Opportunity difference of \hat{Q} : $\mathbb{E}[\hat{Q} Y=1, A=1] - \mathbb{E}[\hat{Q} Y=1, A=0]$

Table I. Notation used in the paper.

Data source

We consider an observable underlying statistical model consisting of three random variables over a probability space $(\Omega, \mathcal{E}, \mathbb{P})$: the *protected feature* $A : \Omega \rightarrow \{0, 1\}$, the *non-protected feature vector* $X : \Omega \rightarrow \mathbb{R}^d$ for some positive integer d , and the *target label* $Y : \Omega \rightarrow \{0, 1\}$. We refer to this statistical model as the *data source*.

The distribution of (X, A) is denoted by the measure π that computes for each $((X, A)$ -measurable) event $E \subseteq \mathbb{R}^d \times \{0, 1\}$, the probability $\pi(E) \stackrel{\text{def}}{=} \mathbb{P}[(X, A) \in E]$. To reduce the verbosity of the discrete case, we denote the probability mass function as $\pi(x, a) \stackrel{\text{def}}{=} \pi(\{(x, a)\})$, i.e. $\pi(x, a) = \mathbb{P}[X=x, A=a]$.

The expectation of Y conditioned on (X, A) is denoted both as the function $q(x, a) \stackrel{\text{def}}{=} \mathbb{E}[Y | X=x, A=a]$ (conditional expectation, see the supplementary material for more details) and the random variable $Q \stackrel{\text{def}}{=} \mathbb{E}[Y | X, A] = q(X, A)$.

The random variable Q plays the role of a soft target label because, since $q(x, a) = \mathbb{P}[Y=1 | X=x, A=a]$, then Y can be modeled as a Bernoulli random variable with success probability Q .

The distribution of (X, A, Y) is completely characterized by the pair (π, q) , hence we refer to this pair as the distribution of the data source. And we distinguish two cases: the data source is *probabilistic* in general, but if $Q \in \{0, 1\}$ (with probability 1), then it is said to be *deterministic*. This distinction is crucial, because several statements hold exclusively in one of the two cases.

Classifiers and predictors

Analogously to the data source, we model the estimation \hat{Y} as a Bernoulli random variable with success probability $\hat{Q} = \hat{q}(X, A)$ for some $((X, A)$ -measurable) function \hat{q} . We refer to \hat{Y} as a (hard) *classifier*, and to \hat{Q} or \hat{q} as a (soft) *predictor*. Notice that \hat{Y} is deterministic when $\hat{Q} \in \{0, 1\}$ (with probability 1), in which case, $\hat{Y} = \hat{Q}$ (w.p. 1). Hence all deterministic classifiers are also predictors.

The set of all soft predictors is denoted as \mathcal{Q} . We highlight the following predictors in \mathcal{Q} :

- 1) the two constant classifiers, $\hat{0}$ and $\hat{1}$, given by $\hat{0}(x, a) \stackrel{\text{def}}{=} 0$ and $\hat{1}(x, a) \stackrel{\text{def}}{=} 1$,
- 2) for each $\hat{Q} \in \mathcal{Q}$, the $1/2$ -threshold classifier given by $\hat{Q}_{1/2} \stackrel{\text{def}}{=} \mathbf{1}_{\hat{Q} > 1/2}$,

- 3) the data source soft target Q , and
- 4) the Bayes classifier $Q_{1/2} = \mathbf{1}_{Q>1/2}$.

It is well known¹ that the Bayes classifier $Q_{1/2}$ has minimal error among all predictors in \mathcal{Q} , regardless of whether the data source is deterministic or not.

Evaluation metrics

To refer to *equal-opportunity* [Hardt et al. \[2016\]](#), we introduce a continuous metric called the *opportunity-difference*. The opportunity-difference of a predictor $\hat{Q} \in \mathcal{Q}$ is defined as

$$\text{oppDiff}(\hat{Q}) \stackrel{\text{def}}{=} +\mathbb{P}[\hat{Y}=1 \mid A=1, Y=1] - \mathbb{P}[\hat{Y}=1 \mid A=0, Y=1]$$

and a predictor $\hat{Q} \in \mathcal{Q}$ is said to satisfy equal-opportunity whenever $\text{oppDiff}(\hat{Q}) = 0$.

The *error* and the *accuracy* of a predictor $\hat{Q} \in \mathcal{Q}$ are defined as

$$\begin{aligned} \text{err}(\hat{Q}) &\stackrel{\text{def}}{=} \mathbb{P}[\hat{Y} \neq Y] \\ \text{acc}(\hat{Q}) &\stackrel{\text{def}}{=} 1 - \text{err}(\hat{Q}) \end{aligned}$$

Additionally, we consider a minimal reference level of accuracy that should be outperformed intuitively by any well-trained predictor. The *trivial accuracy* [Cummings et al. \[2019\]](#) is defined as $\tau \stackrel{\text{def}}{=} \max \{ \text{acc}(\hat{Q}) : \hat{Q} \in \text{Triv} \}$, where Triv is the set of (trivial) predictors whose output does not depend on X and A at all, and as a consequence is independent of Y as well. In other words, Triv consists of all constant soft predictors $\text{Triv} \stackrel{\text{def}}{=} \{ ((x, a) \mapsto c) : c \in [0, 1] \}$. According to the Neyman-Pearson Lemma, the most accurate trivial predictor is always hard, i.e. must be either $\hat{0}$ or $\hat{1}$. Thus τ is well defined and can be computed as

$$\tau = \max \{ \mathbb{P}[Y=0], \mathbb{P}[Y=1] \}.$$

A predictor $\hat{Q} \in \mathcal{Q}$ is said to be *trivially accurate* if $\text{acc}(\hat{Q}) \leq \tau$, and *non-trivially accurate*, or *non-trivial* otherwise. Notice that for a degenerated data source in which the decision Y is independent of X and A , all predictors are forcibly trivially accurate.

IV. THE ERROR VS OPPORTUNITY-DIFFERENCE REGION

In this section, we remark several properties of the region $M \subseteq [0, 1] \times [-1, +1]$ given by

$$M \stackrel{\text{def}}{=} \{ (\text{err}(\hat{Q}), \text{oppDiff}(\hat{Q})) : \hat{Q} \in \mathcal{Q} \}$$

which represents the feasible combinations of the evaluation metrics (error and opportunity-difference) that can be obtained for a given source distribution (π, q) .

M determines the tension between error and opportunity difference. Figure 1 shows an example of this region.

Theorem 1. *Assuming a discrete data source with finitely many possible outcomes, the region M of feasible combinations of error versus opportunity-difference satisfies the following claims.*

- 1) M is a convex polygon.

¹See for instance Chapter 3 of [Fukunaga \[2013\]](#).

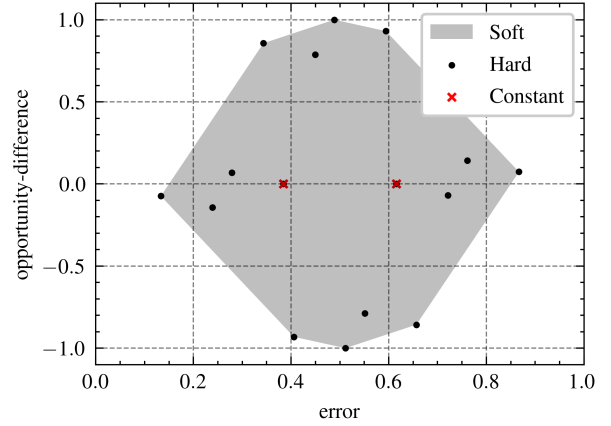


Fig. 1. Region M for an arbitrary source distribution. The data source is not deterministic because the Bayes classifier does not satisfy equal-opportunity. The Python code for generating this figure is in the supplementary material.

- 2) The vertices of the polygon M correspond to some deterministic predictors.

- 3) M is symmetric with respect to the point $(1/2, 0)$.

Proof. The proof is in the supplementary material. It uses the fact that affine transformations map polytopes into polytopes (See Chapter 3 of [Grünbaum \[2013\]](#)). \square

The reader is invited to visualize the aforementioned properties of M in Figure 1, which depicts the region M for a particular instance² of \vec{P} and \vec{Q} .

V. STRONG IMPOSSIBILITY RESULT

Contrasting with Figure 1 in the previous section, Figure 2 shows a data source for which the constant classifiers are vertices of the polygon. This figure illustrates the strong incompatibility that may occur (especially in highly probabilistic distributions). Namely, among the predictors satisfying equal-opportunity (those in the X-axis), the minimal error is achieved by a constant classifier.

In other words, there are data sources for which no predictor can achieve equal-opportunity and non-trivial accuracy simultaneously. This is Theorem 3.

Since Theorem 3 is our strongest result, we also show how to generalize it to non-finite domains. For this purpose, and focusing on formality, we state in Definition 2 very precisely, for which kind of domains it applies.

Definition 2. The *essential range* of a random variable $S : \Omega \rightarrow \mathbb{R}^k$ is the set

$$\{ \vec{s} \in \mathbb{R}^k : (\forall \epsilon > 0) \mathbb{P}[\|S - \vec{s}\| < \epsilon] > 0 \}$$

We call a set $D \subseteq \mathbb{R}^k$ an *essential domain* if it is the essential range of any random variable.

Definition 2 excludes pathological domains such as non-measurable sets, the Cantor set or the irrationals. But it allows for isolated points, convex and closed sets, finite unions of

²Namely $\mathbb{P}=[0.267 \ 0.344 \ 0.141 \ 0.248]$, $\mathbb{Q}=[0.893 \ 0.896 \ 0.126 \ 0.207]$ and $\mathbb{A}=[0 \ 1 \ 0 \ 1]$.




Fig. 2. Randomly generated example using Algorithm 1 (Python code in supplementary material). The constant classifiers are vertices of the polygon, thus the constraints of equal-opportunity and non-trivial accuracy can not be satisfied simultaneously.

them and countable unions of them as long as the resulting set is closed. This includes typical domains, such as products of closed intervals $\prod_{i=1}^n [l_i, r_i]$, or the whole space \mathbb{R}^n .

Theorem 3. *For any essential domain $\mathcal{X} \subseteq \mathbb{R}^d$ with $|\mathcal{X}| \geq 2$ there exists a data source (X, A, Y) whose essential range is $\mathcal{X} \times \{0, 1\}^2$ and such that the accuracy $\text{acc}(\hat{Q})$ of any predictor $\hat{Q} \in \mathcal{Q}$ that satisfies equal opportunity is at most the trivial accuracy τ .*

Proof. The complete proof is contained in the supplementary material. Here we highlight only the sketch, the intuition and some relevant details.

Partition the non-protected input space \mathcal{X} into two non-empty sets $\mathcal{X}_1, \mathcal{X}_2$, and the input space $\mathcal{X} \times \{0, 1\}$ into three regions R_j :

$$\begin{aligned} R_1 &= \mathcal{X}_1 \times \{0\} \\ R_2 &= \mathcal{X}_2 \times \{0\} \\ R_3 &= \mathcal{X} \times \{1\} \end{aligned}$$

For any distribution (π, q) for which these 3 regions have positive probabilities, denote $\vec{P}_j \stackrel{\text{def}}{=} \mathbb{P}[(X, A) \in R_j] > 0$ and $\vec{Q}_j \stackrel{\text{def}}{=} \mathbb{P}[Y=1 | (X, A) \in R_j]$ for $j \in \{1, 2, 3\}$. We search for constraints over \vec{P} and \vec{Q} that are feasible and cause $\text{acc}(\hat{Q}) \leq \tau$ for any predictor $\hat{Q} \in \mathcal{Q}$ satisfying EO. As shown in the supplementary material, the following constraints suffice:

- C1) $\vec{P} \in (0, 1)^3$ and, for probabilism, also $\vec{Q} \in (0, 1)^3$.
- C2) The accuracy of $\hat{Q} = \hat{1}$ is higher than that of $\hat{Q} = \hat{0}$ (for fixing an orientation).
- C3) $\vec{Q}_1 < 1/2$ and $\vec{Q}_2, \vec{Q}_3 > 1/2$.
- C4) $\vec{Q}_3 + \vec{Q}_1 \geq 1$, and
- C5) $\vec{P}_1 \vec{Q}_1 + \vec{P}_2 \vec{Q}_2 < \vec{P}_3 \vec{Q}_1$.

The last three constraints are not straightforward, but their main consequence can be explained graphically. Let us characterize each predictor \hat{Q} , with a vector \vec{F} given by $\vec{F}_j \stackrel{\text{def}}{=} \mathbb{P}[\hat{Y}=1, (X, A) \in R_j]$, so that $\hat{Q} = \hat{1}$ corresponds to $\vec{F} = \vec{P}$. Figure 3 depicts the set of all predictors (box), and those that satisfy EO (plane). This plane can be characterized by the two vectors \vec{Z} and \vec{P} .

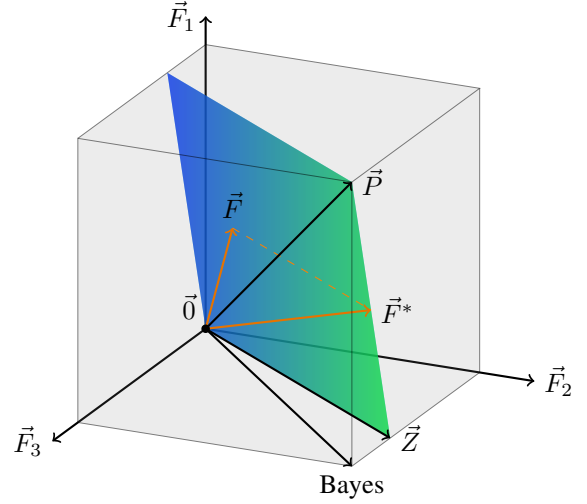


Fig. 3. In vectorial form, the predictors that satisfy equal-opportunity form a plane inside the rectangular box of all predictors.

Constraint C3 simply fixes the location of the Bayes classifier at $(0, \vec{P}_2, \vec{P}_3)$. Constraints C4 and C5 force the gradient of the accuracy along the plane to be non-decreasing in the directions \vec{Z} and $\vec{P} - \vec{Z}$, so that for each predictor \vec{F} there is a pivot \vec{F}^* with higher accuracy than \vec{F} and lower accuracy than \vec{P} . As a consequence, when the constraints are satisfied, $\hat{1}$ has maximal accuracy in \mathcal{Q} .

In order to satisfy the constraints, we propose a randomized algorithm (Algorithm 1) that generates random vectors \vec{P}, \vec{Q} satisfying the five constraints, regardless of the seed and the random sampling function, e.g. uniform. The proof is presented in the supplementary material. To corroborate, Figure 2 shows a particular output of the algorithm³.

The proof concludes by showing that given \vec{P} and \vec{Q} (produced by the algorithm), it is always possible to split $\mathcal{X} \times \{0, 1\}$ into the three regions R_j (for $j \in \{1, 2, 3\}$) that satisfy $\mathbb{P}[(X, A) \in R_j] = \vec{P}_j$ and $\mathbb{P}[Y=1 | (X, A) \in R_j] = \vec{Q}_j$. \square

Algorithm 1 Random generator for Theorem 3.

- 1: **procedure** VECTORGENERATOR(seed)
 - 2: Initialize random sampler with the seed
 - 3: $\vec{Q}_1 \leftarrow$ random in $(0, 1/2)$
 - 4: $\vec{Q}_2 \leftarrow$ random in $(1/2, 1)$
 - 5: $\vec{Q}_3 \leftarrow$ random in $(1 - \vec{Q}_1, 1)$
 - 6: $\vec{P}_3 \leftarrow$ random in $(1/2, 1)$
 - 7: $a \leftarrow \max\{(1 - \vec{P}_3)\vec{Q}_1, 1/2 - \vec{P}_3\vec{Q}_3\}$
 - 8: $b \leftarrow \min\{(1 - \vec{P}_3)\vec{Q}_2, \vec{P}_3\vec{Q}_1\}$
 - 9: $c \leftarrow$ random in (a, b)
 - 10: $\vec{P}_2 \leftarrow (c - \vec{Q}_1(1 - \vec{P}_3)) / (\vec{Q}_2 - \vec{Q}_1)$
 - 11: $\vec{P}_1 \leftarrow 1 - \vec{P}_3 - \vec{P}_2$
 - 12: **return** \vec{P}, \vec{Q}
-

Finally, to conclude this section we present Example 1. It shows that there are many other scenarios, not necessarily

³The output is $\mathbb{P}=[0.131 \ 0.096 \ 0.772]$ and $\mathbb{Q}=[0.274 \ 0.858 \ 0.891]$. Also, $\mathbb{A}=[0 \ 0 \ 1]$ from the partition $\{R_1, R_2, R_3\}$.




Fig. 4. Example 1. One of the constant classifiers is Pareto-optimal.

those of Theorem 3, in which EO and non-trivial accuracy are incompatible.

Example 1. Consider a data source (X, A, Y) over $\{0, 1\}^3$ whose distribution is given by

x	a	$\pi(x, a)$	$q(x, a)$
0	0	3/8	9/20
0	1	2/8	15/20
1	0	1/8	15/20
1	1	2/8	16/20

Then, (i) there are predictors satisfying equal-opportunity, (ii) there are predictors with non-trivial accuracy, but (iii) there are no predictors satisfying both. (End)

Indeed, Figure 4 depicts the region M for Example 1. On the one hand, the set of non-trivially accurate predictors corresponds to the area with error strictly smaller than the left constant classifier. On the other hand, the set of equal-opportunity predictors is (for this particular example) the closed segment between the two constant classifiers. As claimed in Example 1 (and depicted in Figure 4), these two sets are non-empty and do not intersect each other.

VI. PROBABILISTIC VS DETERMINISTIC SOURCES

In this section we compare the tension between error and opportunity-difference when the data source is deterministic and probabilistic. Particularly, we show that some known properties that apply for the discrete case may fail to hold for the probabilistic one, and under what conditions this happens.

As commented in the introduction, analyzing the more general probabilistic case is motivated by two different reasons. On the one hand, scientifically, it provides a more general perspective. On the other hand, ethically, it enables a more realistic representation of a fair environment: one in which the information in (X, A) may be insufficient to conclude the correct decision Y ; in which (X, A) are prone to errors; or in which real-life constraints force the correct decision to be different for identical inputs (x, a) , e.g. two identical candidates for a single job position.

A. Deterministic sources

Under the assumption that the data source is deterministic, there are some important existing results showing the compatibility between equal-opportunity and high accuracy:

Fact 4. Assuming a deterministic data source, if $\tau < 1$, then there is always a non-trivial predictor, for instance the Bayes classifier $Q_{1/2}$. Otherwise (degenerated case) all predictors are trivially accurate.

Fact 5. Assuming a deterministic data source, the Bayes classifier $Q_{1/2}$ satisfies equal-opportunity necessarily.

As a consequence, EO and maximal accuracy (thus also non-trivial accuracy) are always compatible provided $\tau < 1$, because the Bayes classifier satisfies both. This is a celebrated fact and it was part of the motivations of [Hardt et al. \[2016\]](#) for defining equal-opportunity, because other notions of fairness, including statistical parity, are incompatible with accuracy.

B. Probabilistic sources

If we allow the data source to be probabilistic, the results of the deterministic case change. In particular, Fact 4 is generalized by Proposition 6 and Fact 5 is affected by Proposition 7 and Example 1.

Analogous to τ for deterministic sources, we define a secondary reference value $\tau^* \in [0, 1]$. We let

$$\tau^* \stackrel{\text{def}}{=} \max \{ \mathbb{P}[Q \geq 1/2], \mathbb{P}[Q \leq 1/2] \},$$

highlighting that (i) $Q = q(X, A)$ is a random variable varying in $[0, 1]$, (ii) τ and τ^* are equal when the data source is deterministic, and (iii) the condition $\tau = 1$ implies $\tau^* = 1$, but not necessarily the opposite.

As shown in Proposition 6, the equation $\tau^* = 1$ characterizes the necessary and sufficient conditions on the data source for non-trivially accurate predictors to exist.

Particularly, in the deterministic case, we have $\tau^* = \tau$, and Proposition 6 resembles Fact 4.

Proposition 6. (Characterization of the impossibility of non-trivial accuracy)

For any arbitrary source distribution (π, q) , non-trivial predictors exist if and only if $\tau^* < 1$.

Proof. The proof is in the supplementary material. Intuitively, if $\mathbb{P}[Q \geq 1/2] = 1$, then predicting 1 for any input is optimal, and vice versa. □

Finally, in Proposition 7 and its proof, we show a simple family of probabilistic examples for which equal-opportunity and optimal accuracy (obtained by the Bayes classifier) are not compatible. This issue does not merely arise from the fact that the Bayes classifier is hard while the data distribution is soft. Adding randomness to the classifier does not solve the issue. To justify this, and also for completeness, we considered the soft predictor Q and showed that it also fails to satisfy equal-opportunity.

Proposition 7. There are data sources for which neither the Bayes classifier $Q_{1/2}$ nor the predictor Q satisfy equal-opportunity.

Proof. Fix any data source with $\mathbb{P}[A=a, Y=1] > 0$ for each $a \in \{0, 1\}$, pick an arbitrary $((X, A)$ -measurable) function $c: \mathbb{R}^d \rightarrow (0, 1/2)$ and let

$$q(x, a) \stackrel{\text{def}}{=} \begin{cases} 1/2 - c(x) & \text{if } a = 0 \\ 1/2 + c(x) & \text{if } a = 1 \end{cases}$$

for each $(x, a) \in \mathbb{R}^d \times \{0, 1\}$.

Since we know that $Q_{1/2}(x, a) = a$, then the term $\mathbb{E}[Q_{1/2}(X, A) | A=a, Y=1]$ can be reduced more simply into $\mathbb{E}[A | A=a, Y=1] = a$. Therefore, the Bayes classifier satisfies $\text{oppDiff}(Q_{1/2}) = 1 - 0 > 0$.

Regarding Q , we have $\mathbb{E}[Q | A=1, Y=1] = 1/2 + \mathbb{E}[c(X) | A=1, Y=1]$ and $\mathbb{E}[Q | A=0, Y=1] = 1/2 + \mathbb{E}[c(X) | A=0, Y=1]$. Notice from the range of c , that $\mathbb{E}[Q | A=1, Y=1] \in (1/2, 1)$ and $\mathbb{E}[Q | A=0, Y=1] \in (0, 1/2)$. Hence $\text{oppDiff}(Q) > 0$.

Therefore neither $Q_{1/2}$ nor Q satisfy equal-opportunity. \square

As a remark, notice that the data sources proposed in the proof of Proposition 7, contrast the extreme case $Y = A$ because they allow some mutual information between X and Y after A is known, as one would expect in a real-life distribution. Nevertheless, there is an evident inherent demographic disparity in these distributions, and this can be the reason why equal-opportunity hinders optimal accuracy for these examples.

VII. SUFFICIENCY CONDITION

In this section, we provide a simple sufficient (but not necessary) condition (Theorem 9) that guarantees that equal-opportunity and non-triviality are compatible. It is not very restrictive and it is valid for discrete, continuous and mixed data sources. Therefore, it may be used as a minimal assumption for any research work on equal-opportunity dealing with probabilistic data sources. It can also be used to verify whether a data source (X, A, Y) of a particular application is pathogenic for equal-opportunity or not.

Figure 5 summarizes the sufficiency condition in simple manner. The proof consists of showing that when the 4 events highlighted in Figure 5 have positive probabilities, then it is possible to use one of them to improve the performance of the best constant classifier and another one to compensate for equal opportunity.

$Q < 1/2$ $A = 1$	$Q = 1/2$ $A = 1$	$Q > 1/2$ $A = 1$
$Q < 1/2$ $A = 0$	$Q = 1/2$ $A = 0$	$Q > 1/2$ $A = 0$

Fig. 5. Sufficiency condition: If the 4 blue events have positive probability, then equal-opportunity and non-triviality are compatible.

Lemma 8. Assume, for EO to be well defined, that $\mathbb{P}[Y=1, A=a] > 0$ for each $a \in \{0, 1\}$. For any predictor \hat{Q} , we have

$$\mathbb{P}[\hat{Y}=1 | Y=1, A=a] = \frac{\mathbb{E}[\hat{Q}Q | A=a]}{\mathbb{E}[Q | A=a]}$$

Proof. Proved in the supplementary material. \square

Theorem 9. (Sufficiency condition, Figure 5)

For any given data source (X, A, Y) , not-necessarily discrete, if

$$\mathbb{P}[Q > 1/2, A=a], \mathbb{P}[Q < 1/2, A=a] > 0$$

for each $a \in \{0, 1\}$, then equal-opportunity and non-triviality are compatible.

Proof. We begin by noticing that $\mathbb{P}[Q > 1/2, A=a] > 0$ implies $\mathbb{P}[Y=1, A=a] > 0$ for each $a \in \{0, 1\}$, thus equal-opportunity is well-defined.

The proof is divided into two cases depending on which constant classifier is optimal (either $\hat{0}$ or $\hat{1}$). The distinction is needed because equal-opportunity treats $Y = 1$ and $Y = 0$ differently.

Case 1. Assume $\text{err}(\hat{0}) \leq \text{err}(\hat{1})$.

We will show that there are constants $\hat{q}_0, \hat{q}_1 \in [0, 1]$ such that the following predictor satisfies equal-opportunity and non-triviality.

$$\hat{Q} \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } Q \leq 1/2 \\ \hat{q}_0 & \text{if } Q > 1/2, A = 0 \\ \hat{q}_1 & \text{if } Q > 1/2, A = 1 \end{cases}$$

For \hat{Q} to satisfy non-triviality, it suffices to have smaller error than $\hat{0}$. This holds whenever \hat{q}_0 or \hat{q}_1 are positive, because for each $a \in \{0, 1\}$, the optimal classification over the region $\{Q > 1/2, A = a\}$ is 1, thus any positive value \hat{q}_a improves the misclassification of $\hat{0}$. For this reason, we restrict $\hat{q}_0, \hat{q}_1 \in (0, 1]$.

Regarding equal-opportunity, recall from Lemma 8 that

$$\mathbb{P}[\hat{Y}=1 | Y=1, A=a] = \frac{\mathbb{E}[\hat{Q}Q | A=a]}{\mathbb{E}[Q | A=a]}$$

Let us call $\alpha_a \stackrel{\text{def}}{=} \mathbb{E}[\hat{Q}Q | A=a]$ to the numerator. Since $\hat{Q} = 0$ for $Q \leq 1/2$, then α_a may be computed as

$$\alpha_a = \hat{q}_a \mathbb{E}[Q | A=a, Q > 1/2] \mathbb{P}[Q > 1/2 | A=a]$$

and it is positive.

Hence, equal-opportunity may be stated as

$$\frac{\hat{q}_1}{\hat{q}_0} = \frac{\alpha_0 \mathbb{E}[Q | A=1]}{\alpha_1 \mathbb{E}[Q | A=0]}$$

The right hand side term is always well-defined, and it is a positive real number. Since \hat{q}_0 and \hat{q}_1 can be made arbitrarily small, there are always solutions to this equation in the range $\hat{q}_0, \hat{q}_1 \in (0, 1]$.

Case 2. Assume $\text{err}(\hat{1}) < \text{err}(\hat{0})$.

Analogously, for $\hat{q}_0, \hat{q}_1 \in [0, 1]$, let \hat{Q} be given by

$$\hat{Q} \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } Q \geq 1/2 \\ \hat{q}_0 & \text{if } Q < 1/2, A = 0 \\ \hat{q}_1 & \text{if } Q < 1/2, A = 1 \end{cases}$$

If $\hat{q}_0, \hat{q}_1 < 1$, then \hat{Q} satisfies non-triviality, because it has less error than $\hat{1}$. Hence, we restrict $\hat{q}_0 \hat{q}_1 < 1$.

Equal-opportunity, may be equivalently stated in terms of $\mathbb{P}[\hat{Y}=0 | Y=1, A=a]$ because it is the complement of $\mathbb{P}[\hat{Y}=1 | Y=1, A=a]$. Recall from Lemma 8 that

$$\mathbb{P}[\hat{Y}=0 | Y=1, A=a] = \frac{\mathbb{E}[(1-\hat{Q})Q | A=a]}{\mathbb{E}[Q | A=a]}$$

Let us call $\alpha_a \stackrel{\text{def}}{=} \mathbb{E}[(1-\hat{Q})Q | A=a]$ to the numerator. Since $(1-\hat{Q})=0$ for $Q \geq 1/2$, then α_a may be computed as

$$\alpha_a = (1-\hat{q}_a)\mathbb{E}[Q | A=a, Q < 1/2]\mathbb{P}[Q < 1/2 | A=a]$$

and it is non-negative.

Hence, equal-opportunity may be stated as

$$(1-\hat{q}_1)\alpha_1\mathbb{E}[Q | A=0] = (1-\hat{q}_0)\alpha_0\mathbb{E}[Q | A=1]$$

If $\alpha_1\mathbb{E}[Q | A=0] = 0$, we let $\hat{q}_0 = 1$ and $\hat{q}_1 = 1/2$. If $\alpha_0\mathbb{E}[Q | A=1] = 0$, we let $\hat{q}_1 = 1$ and $\hat{q}_0 = 1/2$. And if none of the two is zero, we use the same argument as in the first case: since $1-\hat{q}_0$ and $1-\hat{q}_1$ can be made arbitrarily small, there are always solutions to this equation in the range $\hat{q}_0, \hat{q}_1 \in [0, 1]$. \square

VIII. CONCLUSION AND FUTURE WORK

Our work extends existing results about equal-opportunity and accuracy from a deterministic data source to a probabilistic one. The main result, Theorem 3, states that for certain probabilistic data sources, no predictor can achieve equal-opportunity and non-trivial accuracy simultaneously. We also provided a sufficient condition on the data source under which EO and non-trivial accuracy are guaranteed to be compatible.

Our method focuses on the fairness notion of equal-opportunity, which seeks for equal true positive rates. A symmetric analysis can be carried out for equal false positive rates using the same ideas we provided. Since the notion of equal-odds seeks for both equal true positive rates and equal false positive rates, our methodology and results can be used for equal-odds. However, our method is not compatible with statistical parity or with fairness notions based on individual fairness.

An interesting question left for future work is whether the scenarios in which equal-opportunity and non-trivial accuracy are incompatible require the data source to be unfair on its own in some sense. If this is true, it would provide additional justification for equal-opportunity as a fairness notion.

Other future lines of research include characterizing completely the conditions under which equal-opportunity and non-trivial accuracy are compatible, studying the trade-off and Pareto-optimality between accuracy and opportunity-difference, and bounding the opportunity-difference by taking into account the learning process and the statistical sampling.

In addition, we plan to measure the accuracy gap between the Bayes classifier and the most accurate predictor that satisfies equal-opportunity.

REFERENCES

- Sushant Agarwal. Trade-offs between fairness, interpretability, and privacy in machine learning. Master's thesis, University of Waterloo, 2020.
- Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163, 2017.
- Kate Crawford. The trouble with bias. In *Conference on Neural Information Processing Systems (NIPS)*, 2017.
- Rachel Cummings, Varun Gupta, Dhamma Kimpara, and Jamie Morgenstern. On the compatibility of privacy and fairness. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*, pages 309–315, 2019.
- Keinosuke Fukunaga. *Introduction to statistical pattern recognition*. Elsevier, 2013.
- Branko Grünbaum. *Convex polytopes*, volume 221. Springer Science & Business Media, 2013.
- Moritz Hardt, Eric Price, and Nathan Srebro. Equality of opportunity in supervised learning. *arXiv preprint arXiv:1610.02413*, 2016.
- Michael Kearns and Aaron Roth. *The ethical algorithm: The science of socially aware algorithm design*. Oxford University Press, 2019.
- Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.
- Zachary C Lipton, Alexandra Chouldechova, and Julian McAuley. Does mitigating ml's impact disparity require treatment disparity? In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 8136–8146, 2018.
- David Pujol, Ryan McKenna, Satya Kuppam, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. Fair decision making using privacy-protected data. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 189–199, 2020.

SUPPLEMENTARY MATERIAL

A. Remark on conditional expectation

The notation $\mathbb{E}[Y | X=x, A=a]$ for defining $q(x, a)$ is not an expectation conditioned on the possibly null event $(X = x, A = a)$. Instead, it is syntactic sugar for the conditional expectation function.

Formally speaking, the function q is not necessarily unique. It is uniquely defined at the points (x, a) where $\mathbb{P}[X=x, A=a] > 0$. In the other points, it is defined almost everywhere uniquely in a conditional expectation sense, so that for any other valid conditional expectation function q' , we have $q(X, A) = q'(X, A)$ almost surely. Throughout the paper, we prioritize studying the discrete case to avoid this extreme level of formalism without losing rigor.

B. The error vs opportunity-difference region

In this section, we analyze the region $M \subseteq [0, 1] \times [-1, +1]$ given by

$$M \stackrel{\text{def}}{=} \{(\text{err}(\hat{Q}), \text{oppDiff}(\hat{Q})) : \hat{Q} \in \mathcal{Q}\}$$

which represents the feasible combinations of the evaluation metrics (error and opportunity-difference) that can be obtained for a given source distribution (π, q) . This region determines the tension between error and opportunity difference. Figure 6 shows an example of the region M .

The results presented in this section assume that the data source is discrete and its range is finite. We will use the following vectorial notation to represent both the distribution (π, q) and any arbitrary predictor $\hat{Q} \in \mathcal{Q}$.

Definition 1. Suppose (X, A) can only take a finite number of outcomes $\{(x_i, a_i)\}_{i=1}^n$ (each with positive probability) for some integer $n > 0$. In order to represent π, q and any $\hat{Q} \in \mathcal{Q}$ respectively, let $\vec{P}, \vec{Q}, \vec{F} \in \mathbb{R}^n$ be the vectors given by

$$\begin{aligned} \vec{P}_i &\stackrel{\text{def}}{=} \mathbb{P}[X=x_i, A=a_i] \\ \vec{Q}_i &\stackrel{\text{def}}{=} \mathbb{P}[Y=1 | X=x_i, A=a_i] \\ \vec{F}_i &\stackrel{\text{def}}{=} \mathbb{P}[\hat{Y}=1, X=x_i, A=a_i] \end{aligned}$$

For notation purposes, let also $\vec{Q}^{(0)}, \vec{Q}^{(1)} \in \mathbb{R}^n$ be given by $\vec{Q}_i^{(a)} \stackrel{\text{def}}{=} \vec{Q}_i \cdot \mathbf{1}_{a_i=a}$, and, following the definition of $\text{err}(\hat{Q})$ and $\text{oppDiff}(\hat{Q})$, let

$$\begin{aligned} \text{err}(\vec{F}) &\stackrel{\text{def}}{=} \langle \vec{P}, \vec{Q} \rangle + \langle \vec{F}, 1 - 2\vec{Q} \rangle \\ \text{oppDiff}(\vec{F}) &\stackrel{\text{def}}{=} \frac{\langle \vec{F}, \vec{Q}^{(1)} \rangle}{\langle \vec{P}, \vec{Q}^{(1)} \rangle} - \frac{\langle \vec{F}, \vec{Q}^{(0)} \rangle}{\langle \vec{P}, \vec{Q}^{(0)} \rangle} \end{aligned}$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product:

$$\langle u, v \rangle \stackrel{\text{def}}{=} u_1 v_1 + \dots + u_n v_n. \quad (\text{End})$$

Regarding Definition 1, we highlight three important remarks:

- 1) $\vec{Q} \in [0, 1]^n$, $\vec{P} \in (0, 1]^n$, $\|\vec{P}\|_1 = 1$ and \vec{F} lies in the rectangular n -dimensional box given by

$$0 \preceq \vec{F} \preceq \vec{P},$$

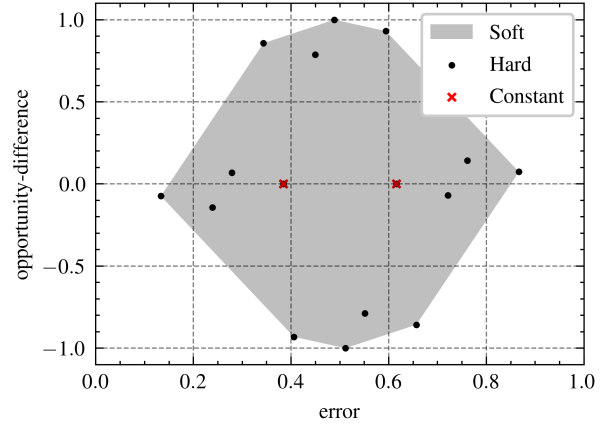


Fig. 6. Region M for an arbitrary source distribution. The data source is not deterministic because the Bayes classifier does not satisfy equal-opportunity.

where \preceq denotes the componentwise order in \mathbb{R}^n , i.e. $0 \preceq \vec{F}_i \preceq \vec{P}_i$ for each $i \in \{1, \dots, n\}$. Moreover, from the definition of \vec{P} and \vec{F} , the vertices of this rectangular box correspond precisely with the deterministic predictors.

- 2) The vectorial definitions of error and opportunity-difference correspond to those of the non-vectorial case as shown in Lemma 14. Moreover, their gradients are constant.
- 3) There is a one-to-one correspondence between the predictors $\hat{q} \in \mathcal{Q}$ and the vectors \vec{F} that satisfy $0 \preceq \vec{F} \preceq \vec{P}$. Indeed, each predictor is uniquely given by its pointwise values $\hat{q}(x_i, a_i) = \frac{\vec{F}_i}{\vec{P}_i}$ and each vector by its pointwise coordinates $\vec{F}_i = \vec{P}_i \hat{q}(x_i, a_i)$. Therefore

$$M = \{(\text{err}(\vec{F}), \text{oppDiff}(\vec{F})) : 0 \preceq \vec{F} \preceq \vec{P}\}$$

We now make use of results from a different research area in mathematics, geometry, to conclude the main properties of the region M .

Theorem 10. Assuming a discrete data source with finitely many possible outcomes, the region M of feasible combinations of error versus opportunity-difference satisfies the following claims.

- 1) M is a convex polygon.
- 2) The vertices of the polygon M correspond to some deterministic predictors.
- 3) M is symmetric with respect to the point $(1/2, 0)$.

Proof. Assume the notation of Definition 1.

Part 1. In geometrical terms, M is the result of applying an affine transformation, i.e. a linear transformation and a translation, to the n -dimensional polytope given by $0 \preceq \vec{F} \preceq \vec{P}$.

Affine transformations are known to map polytopes into polytopes (See Chapter 3 of Grünbaum [2013]), therefore M must be a 2-dimensional polytope, i.e. the region M is a convex polygon. In theory, this region may also be a 1-dimensional segment, but this can only occur in the extreme case that $Q = 1/2$ (with probability 1).

Part 2. The vertices of a polytope, also called extremal points, are the points in the polytope that are not in the segment between any two other points in the polytope. It is known from geometry theory that affine mappings preserve collinearity, i.e. they map segments into segments, thus they map non-vertices into non-vertices. As a consequence, the vertices of the polygon M correspond to some vertices of the polytope $0 \preceq \vec{F} \preceq \vec{P}$, that is, to some deterministic classifiers.

Part 3. Notice (Lemma 15) that

$$\begin{aligned} \text{err}(\vec{P} - \vec{F}) &= 1 - \text{err}(\vec{F}) \\ \text{oppDiff}(\vec{P} - \vec{F}) &= -\text{oppDiff}(\vec{F}) \end{aligned}$$

This implies that for each point $(\text{err}(\vec{F}), \text{oppDiff}(\vec{F})) \in M$, there is another one, namely $(\text{err}(\vec{P} - \vec{F}), \text{oppDiff}(\vec{P} - \vec{F})) \in M$ that is symmetrical to the former w.r.t the point $(1/2, 0)$. Geometrically, this means that the polygon M is symmetric with respect to the point $(1/2, 0)$. \square

The reader is invited to visualize the aforementioned properties of M in Figure 6, which depicts the region M for a particular instance⁴ of \vec{P} and \vec{Q} .

C. Impossibility theorem

Theorem 11. For any essential domain $\mathcal{X} \subseteq \mathbb{R}^d$ with $|\mathcal{X}| \geq 2$ there exists a data source (X, A, Y) whose essential range is $\mathcal{X} \times \{0, 1\}^2$ and such that the accuracy $\text{acc}(\hat{Q})$ of any predictor $\hat{Q} \in \mathcal{Q}$ that satisfies equal opportunity is at most the trivial accuracy τ .

Proof. The proof is divided in four parts. We will (i) reduce the problem into an algebraic one; (ii) find the linear constraints that solve the algebraic problem when satisfied; (iii) provide an algorithm that generates vectors that satisfy the linear constraints; and finally, (iv) convert the vectorial solution back into a distribution (π, q) for the given domain.

Part 1. Reduction to an algebraic problem.

Partition the non-protected input space \mathcal{X} into two non-empty sets $\mathcal{X}_1, \mathcal{X}_2$, and the input space $\mathcal{X} \times \{0, 1\}$ into three regions R_j :

$$\begin{aligned} R_1 &= \mathcal{X}_1 \times \{0\} \\ R_2 &= \mathcal{X}_2 \times \{0\} \\ R_3 &= \mathcal{X} \times \{1\} \end{aligned}$$

For any distribution (π, q) for which these 3 regions have positive probabilities, denote $\vec{P}_j \stackrel{\text{def}}{=} \mathbb{P}[(X, A) \in R_j] > 0$ and $\vec{Q}_j \stackrel{\text{def}}{=} \mathbb{P}[Y=1 | (X, A) \in R_j]$ for $j \in \{1, 2, 3\}$. We search for constraints over \vec{P} and \vec{Q} that are feasible and cause $\text{acc}(\hat{Q}) \leq \tau$ for any fair predictor $\hat{Q} \in \mathcal{Q}$ satisfying EO. The first such constraint is

$$C6) \quad \vec{P}, \vec{Q} \in (0, 1)^3.$$

That is, we require \vec{P}_j to be positive, and Y to have at least some degree of randomness in each region.

Given a reference predictor \hat{Q} , let $\vec{F} \in [0, 1]^3$ be the vector given by $\vec{F}_j \stackrel{\text{def}}{=} \mathbb{P}[\hat{Y}=1, (X, A) \in R_j]$. Lemma 14 shows that

⁴Namely $\mathbb{P}=[0.267 \ 0.344 \ 0.141 \ 0.248]$, $\mathbb{Q}=[0.893 \ 0.896 \ 0.126 \ 0.207]$ and $\mathbb{A}=[0 \ 1 \ 0 \ 1]$

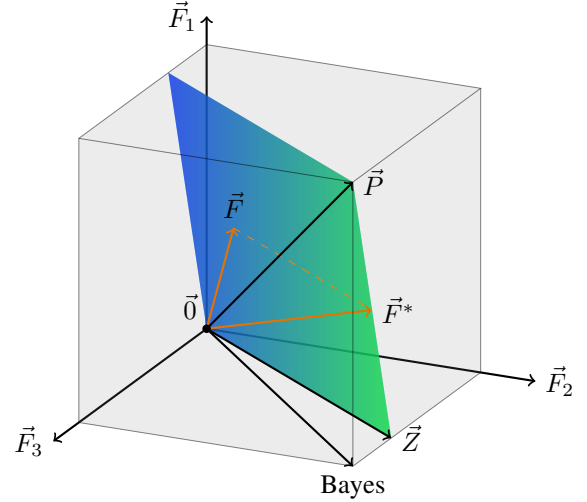


Fig. 7. In vectorial form, the predictors that satisfy equal-opportunity form a plane inside the rectangular box of all predictors.

the accuracy and the opportunity-difference of any predictor \hat{Q} can be computed from \vec{P} , \vec{Q} and \vec{F} as

$$\begin{aligned} \text{acc}(\hat{Q}) &= \langle \vec{F}, 2\vec{Q} - 1 \rangle + C_{\vec{Q}} \\ \text{oppDiff}(\hat{Q}) &= \frac{\vec{F}_3}{\vec{P}_3} - \frac{\vec{F}_1\vec{Q}_1 + \vec{F}_2\vec{Q}_2}{\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2} \end{aligned}$$

where $C_{\vec{Q}} \stackrel{\text{def}}{=} 1 - \langle \vec{P}, \vec{Q} \rangle$ is a constant and the operator $\langle \cdot, \cdot \rangle$ denotes the inner product explained in Definition 1. Since we are interested in relative accuracies with respect to the trivial predictors, the constant $C_{\vec{Q}}$ is mostly irrelevant. For this reason, we let $L(\vec{F}) \in [-1, 1]$ denote the non-constant component of the accuracy $L(\vec{F}) \stackrel{\text{def}}{=} \langle \vec{F}, 2\vec{Q} - 1 \rangle$.

Both accuracy and opportunity-difference are completely determined for any predictor by the vectors \vec{P} , \vec{Q} and \vec{F} as shown above. Moreover, both quantities are linear with respect to \vec{F} .

Regarding equal-opportunity, the constraint $\text{oppDiff}(\hat{Q}) = 0$ forms a plane in \mathbb{R}^3 , depicted in Figure 7. This plane passes through the origin, is determined by \vec{P} and \vec{Q} , and contains all vectors \vec{F} (restricted to $0 \leq \vec{F}_j \leq \vec{P}_j$) that satisfy

$$\vec{F}_3(\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2) - \vec{P}_3(\vec{F}_1\vec{Q}_1 + \vec{F}_2\vec{Q}_2) = 0$$

or equivalently, all vectors F that are normal to the vector $(-\vec{P}_3\vec{Q}_1, -\vec{P}_3\vec{Q}_2, \vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2)$.

Regarding accuracy, the two constant predictors correspond to $\vec{F} = \vec{0}$ and $\vec{F} = \vec{P}$, thus $\tau = C_{\vec{Q}} + \max\{L(\vec{0}), L(\vec{P})\}$. Importantly, both of them lie on the equal-opportunity plane.

The problem is now reduced to finding vectors \vec{P} and \vec{Q} such that all vectors \vec{F} in the equal-opportunity plane satisfy $L(\vec{F}) \leq \max\{L(\vec{0}), L(\vec{P})\}$.

Part 2. Constraints for the algebraic solution.

To fix an orientation, let us impose these constraints:

$$C7) \quad \text{Among the constant predictors, the accuracy of } \vec{F} = \vec{P} \text{ is higher than that of } \vec{F} = \vec{0}. \text{ This is } L(\vec{P}) > 0 = L(\vec{0}).$$

C8) The Bayes classifier is located at $(0, \vec{P}_2, \vec{P}_3)$ as in Figure 7. Algebraically this means $\vec{Q}_1 < 1/2$ and $\vec{Q}_2, \vec{Q}_3 > 1/2$.

In order to derive the constraints that make the scalar field L maximal at \vec{P} over the plane, consider the vector \vec{Z} that lies on the plane and has minimal \vec{Z}_1 and maximal \vec{Z}_2 , i.e.

$$\vec{Z} \stackrel{\text{def}}{=} (0, \vec{P}_2, \vec{P}_3 \frac{\vec{P}_2 \vec{Q}_2}{\vec{P}_1 \vec{Q}_1 + \vec{P}_2 \vec{Q}_2})$$

Since the gradient of L is given by $2\vec{Q} - 1$ and has signs $(-, +, +)$, then for any vector \vec{F} in the plane, there is \vec{F}^* in the segment between \vec{P} and \vec{Z} such that $\vec{F}_1 = \vec{F}_1^*$ and $L(\vec{F}^*) \geq L(\vec{F})$ (refer to Figure 7). This implies that the L attains its maximal value on the segment between \vec{P} and \vec{Z} . Hence, for L to be maximal at \vec{P} , it would suffice to have $L(\vec{P}) > L(\vec{Z})$. As shown in Lemma 20, this can be achieved by imposing, in addition,

- C9) $\vec{Q}_3 + \vec{Q}_1 \geq 1$, and
 C10) $\vec{P}_1 \vec{Q}_1 + \vec{P}_2 \vec{Q}_2 < \vec{P}_3 \vec{Q}_1$.

Part 3. Solution to the constraints.

Algorithm 2 is a randomized algorithm that generates random vectors. We will prove that the output vectors \vec{P} and \vec{Q} satisfy the constraints of the previous parts of this proof, regardless of the seed and the random sampling function, e.g. uniform.

Algorithm 2 Random generator for Theorem 11.

```

1: procedure VECTORGENERATOR(seed)
2:   Initialize random sampler with the seed
3:    $\vec{Q}_1 \leftarrow$  random in  $(0, 1/2)$ 
4:    $\vec{Q}_2 \leftarrow$  random in  $(1/2, 1)$ 
5:    $\vec{Q}_3 \leftarrow$  random in  $(1 - \vec{Q}_1, 1)$ 
6:    $\vec{P}_3 \leftarrow$  random in  $(1/2, 1)$ 
7:    $a \leftarrow \max\{(1 - \vec{P}_3)\vec{Q}_1, 1/2 - \vec{P}_3\vec{Q}_3\}$ 
8:    $b \leftarrow \min\{(1 - \vec{P}_3)\vec{Q}_2, \vec{P}_3\vec{Q}_1\}$ 
9:    $c \leftarrow$  random in  $(a, b)$ 
10:   $\vec{P}_2 \leftarrow (c - \vec{Q}_1(1 - \vec{P}_3)) / \vec{Q}_2 - \vec{Q}_1$ 
11:   $\vec{P}_1 \leftarrow 1 - \vec{P}_3 - \vec{P}_2$ 
12:  return  $\vec{P}, \vec{Q}$ 

```

Two immediate observations about Algorithm 2 are that the construction of \vec{Q} implies that constraints C8 and C9 are satisfied, and the construction of \vec{P} implies $\vec{P}_1 + \vec{P}_2 + \vec{P}_3 = 1$. To prove the correctness of the algorithm, it remains to prove that (i) $a < b$ (otherwise the algorithm would not be well-defined), that (ii) $\vec{P}_2 \in (0, 1)$ for constraint C6, and also that (iii) constraints C7 and C10 are satisfied. For better readability, the algebraic proof of these claims are moved to Lemma 21.

Part 4. Construction of the distribution.

Generate a pair of vectors \vec{P} and \vec{Q} using the algorithm of the previous part (Part 3). The first goal is to partition \mathcal{X} into \mathcal{X}_1 and \mathcal{X}_2 to generate the regions R_1, R_2 and R_3 . The second goal is to define π in such a way that $\mathbb{P}[(X, A) \in R_j] = \vec{P}_j$ for each $j \in \{1, 2, 3\}$. The third and last goal is to define q so that $\mathbb{E}[Q | (X, A) \in R_j] = \vec{Q}_j$ for each j . This can be done immediately by letting $q(x, a) \stackrel{\text{def}}{=} \vec{Q}_j$ for all $(x, a) \in R_j$. Thus only the first two goals remain.

For the first goal, since $|\mathcal{X}| \geq 2$, we may create a simple Voronoi clustering diagram by choosing two different arbitrary points $s_1, s_2 \in \mathcal{X}$, and letting $\mathcal{X}_1 \stackrel{\text{def}}{=} \{s \in \mathcal{X} : \|s - s_1\| \leq \|s - s_2\|\}$ and $\mathcal{X}_2 \stackrel{\text{def}}{=} \mathcal{X} \setminus \mathcal{X}_1$.

For the second goal, since \mathcal{X} is an essential domain, there exists a random variable S whose essential range is \mathcal{X} . Notice that $\mathbb{P}[S \in \mathcal{X}_j] \geq \mathbb{P}[\|S - s_j\| < \|s_1 - s_2\|/2] > 0$ for each $j \in \{1, 2\}$. For each $((X, A)$ -measurable) event E , let $E_a \stackrel{\text{def}}{=} \{x : (x, a) \in E\}$, and define $\pi(E)$ as

$$\begin{aligned} \mathbb{P}[(X, A) \in E] &\stackrel{\text{def}}{=} \sum_{a=0,1} \mathbb{P}[X \in E_a, A=a] \\ \mathbb{P}[X \in E_0, A=0] &\stackrel{\text{def}}{=} \sum_{j=1,2} \mathbb{P}[S \in E_0 | S \in \mathcal{X}_j] \vec{P}_j \\ \mathbb{P}[X \in E_1, A=1] &\stackrel{\text{def}}{=} \mathbb{P}[S \in E_1] \vec{P}_3 \end{aligned}$$

This forces $\mathbb{P}[(X, A) \in R_j] = \vec{P}_j$ for each $j \in \{1, 2, 3\}$ as desired. \square

D. Characterization of the impossibility of non-trivial accuracy

Proposition 12. (Characterization of the impossibility of non-trivial accuracy)

For any arbitrary source distribution (π, q) , non-trivial predictors exist if and only if $\tau^* < 1$.

Proof. We will prove that all predictors are trivially accurate if and only if $\tau^* = 1$.

(\Leftarrow) Suppose $\tau^* = 1$, i.e. $\mathbb{P}[Q \leq 1/2] = 1$ or $\mathbb{P}[Q \geq 1/2] = 1$.

In the former case, the Bayes classifier $Q_{1/2}$ is the constant predictor $(x, a) \mapsto 0$, thus $\text{acc}(Q_{1/2}) \leq \tau$ necessarily. In the latter case, the alternative Bayes classifier $Q_{1/2}^*$ (defined in Lemma 18) is the constant predictor $(x, a) \mapsto 1$, thus $\text{acc}(Q_{1/2}^*) \leq \tau$. According to Lemma 18, $\text{acc}(Q_{1/2}) = \text{acc}(Q_{1/2}^*)$, thus we may conclude $\text{acc}(Q_{1/2}) \leq \tau$ as well.

It follows that $\text{acc}(\hat{Q}) \leq \text{acc}(Q_{1/2}) \leq \tau$ for all $\hat{Q} \in \mathcal{Q}$ because $Q_{1/2}$ has maximal accuracy in \mathcal{Q} .

(\Rightarrow) Suppose $\tau^* < 1$. We will suppose that the Bayes classifier $Q_{1/2}$ is not trivially accurate and find a contradiction.

Suppose $\text{acc}(Q_{1/2}) - \tau = 0$. According to Lemmas 16 and 19 we may rewrite this as $\mathbb{E}[|Q - 1/2| - |\mathbb{E}[Y] - 1/2|] = 0$. Using the reverse triangle inequality, we conclude $\mathbb{E}[|Q - \mathbb{E}[Y]|] = 0$, thus $Q = \mathbb{E}[Y]$ is constant.

If $\mathbb{E}[Y] \leq 1/2$, then $\mathbb{P}[Q \leq 1/2] = 1$. If $\mathbb{E}[Y] \geq 1/2$, then $\mathbb{P}[Q \geq 1/2] = 1$. In any case, we have $\tau^* = 1$ which contradicts the initial supposition. \square

E. Lemmas

F. Some helper lemmas

Lemma 13. For every $\hat{Q} \in \mathcal{Q}$,

$$\text{err}(\hat{Q}) = \mathbb{E}[|\hat{Q} - Y|]$$

Proof. Notice that $\mathbb{P}[\hat{Y} \neq Y | Y=1] = \mathbb{E}[1 - \hat{Q} | Y=1]$ and $\mathbb{P}[\hat{Y} \neq Y | Y=0] = \mathbb{E}[\hat{Q} | Y=0]$. In both cases, we may write $\mathbb{P}[\hat{Y} \neq Y | Y=y] = \mathbb{E}[|Y - \hat{Q}| | Y=y]$.

Hence, marginalizing over Y we conclude $\mathbb{P}[\hat{Y} \neq Y] = \mathbb{E}[|Y - \hat{Q}|]$. \square

Lemma 14. (Vectorial metrics) Using the notation of Definition 1, we have

$$\begin{aligned} \text{err}(\hat{Q}) &= \text{err}(\vec{F}) \\ \text{oppDiff}(\hat{Q}) &= \text{oppDiff}(\vec{F}) \end{aligned}$$

Proof. For the error, we marginalize over (X, A) . Notice

$$\begin{aligned} &\mathbb{P}[Y \neq \hat{Y} | X = x_i, A = a_i] \\ &= + (1 - q(x_i, a_i))\hat{q}(x_i, a_i) \\ &\quad + q(x_i, a_i)(1 - \hat{q}(x_i, a_i)) \\ &= (1 - \vec{Q}_i) \frac{\vec{F}_i}{\vec{P}_i} + \vec{Q}_i \frac{\vec{P}_i - \vec{F}_i}{\vec{P}_i} \\ &= \frac{\vec{Q}_i \vec{P}_i + \vec{F}_i(1 - 2\vec{Q}_i)}{\vec{P}_i} \end{aligned}$$

Thus $\mathbb{P}[Y \neq \hat{Y}, X = x_i, A = a_i] = \vec{F}_i \vec{P}_i + \vec{F}_i(1 - 2\vec{Q}_i)$, and

$$\begin{aligned} \text{err}(f) &= \mathbb{P}[\hat{Y} \neq Y] \\ &= \sum_{i=1}^n \mathbb{P}[\hat{Y} \neq Y, X = x_i, A = a_i] \\ &= \langle \vec{P}, \vec{Q} \rangle + \langle \vec{F}, 1 - 2\vec{Q} \rangle \\ &= \text{err}(\vec{F}) \end{aligned}$$

For opportunity-difference, we also marginalize over (X, A) . Notice that

$$\begin{aligned} &\mathbb{P}[\hat{Y} = 1, Y = 1, X = x_i, A = a_i] \\ &= \vec{P}_i \mathbb{E}[\hat{Q}Q | X = x_i, A = a_i] \\ &= \vec{P}_i \frac{\vec{F}_i}{\vec{P}_i} \vec{Q}_i = \vec{F}_i \vec{Q}_i, \end{aligned}$$

hence $\mathbb{P}[\hat{Y} = 1, Y = 1, X = x_i, A = a_i] = \vec{F}_i \vec{Q}_i^{(a)}$. In addition, $\mathbb{P}[Y = 1, X = x_i, A = a] = \vec{P}_i \vec{Q}_i^{(a)}$ and

$$\begin{aligned} &\mathbb{P}[\hat{Y} = 1 | Y = 1, A = a] \\ &= \frac{\sum_{i=1}^n \mathbb{P}[\hat{Y} = 1, Y = 1, X = x_i, A = a]}{\sum_{i=1}^n \mathbb{P}[Y = 1, X = x_i, A = a]} \\ &= \frac{\langle \vec{F}, \vec{Q}^{(a)} \rangle}{\langle \vec{P}, \vec{Q}^{(a)} \rangle} \end{aligned}$$

Therefore

$$\begin{aligned} \text{oppDiff}(\hat{Q}) &= + \mathbb{P}[\hat{Y} = 1 | Y = 1, A = 1] \\ &\quad - \mathbb{P}[\hat{Y} = 1 | Y = 1, A = 0] \\ &= \frac{\langle \vec{F}, \vec{Q}^{(1)} \rangle}{\langle \vec{P}, \vec{Q}^{(1)} \rangle} - \frac{\langle \vec{F}, \vec{Q}^{(0)} \rangle}{\langle \vec{P}, \vec{Q}^{(0)} \rangle} \\ &= \text{oppDiff}(\vec{F}) \end{aligned}$$

Lemma 15. (Metrics symmetry) Using the notation of Definition 1, we have

$$\begin{aligned} \text{err}(\vec{P} - \vec{F}) &= 1 - \text{err}(\vec{F}) \\ \text{oppDiff}(\vec{P} - \vec{F}) &= -\text{oppDiff}(\vec{F}) \end{aligned}$$

Proof. According to Lemma 14, opportunity-difference is a linear transformation. Since linear transformations preserve scalar multiplication and vector addition, it follows that $\text{oppDiff}(\vec{P} - \vec{F}) = \text{oppDiff}(\vec{P}) - \text{oppDiff}(\vec{F})$. Moreover, since $\text{oppDiff}(\vec{P}) = 1 - 1 = 0$, then $\text{oppDiff}(\vec{P} - \vec{F}) = -\text{oppDiff}(\vec{F})$.

According to the same lemma, the error is an affine transformation with offset $\langle \vec{P}, \vec{Q} \rangle$. Hence

$$\begin{aligned} \text{err}(\vec{P} - \vec{F}) &= \text{err}(\vec{P}) - \text{err}(\vec{F}) + \langle \vec{P}, \vec{Q} \rangle \\ &= 2\langle \vec{P}, \vec{Q} \rangle - \langle \vec{P}, 1 - 2\vec{Q} \rangle - \text{err}(\vec{F}) \\ &= \langle \vec{P}, 1 \rangle - \text{err}(\vec{F}) \\ &= 1 - \text{err}(\vec{F}) \end{aligned}$$

because $\sum_{i=1}^n \vec{P}_i = 1$. \square

Lemma 16. (Bayes accuracy)

$$\text{acc}(Q_{1/2}) = 1/2 + \mathbb{E}[|Q - 1/2|]$$

Proof. Out of Lemma 13, we know $\text{err}(Q_{1/2}) = 1 - \mathbb{E}[\epsilon]$ where $\epsilon \stackrel{\text{def}}{=} |Q_{1/2} - Y|$. Let us condition on $Q < 1/2$ and $Q \geq 1/2$ separately (whenever these events have possible probabilities).

For $Q < 1/2$, we have $\mathbb{E}[\epsilon | Q < 1/2] = \mathbb{E}[Y | Q < 1/2] = \mathbb{E}[Q | Q < 1/2]$ and $Q = 1/2 - (1/2 - Q)$. For $Q \geq 1/2$, we have $\mathbb{E}[\epsilon | Q \geq 1/2] = \mathbb{E}[1 - Y | Q \geq 1/2] = \mathbb{E}[1 - Q | Q \geq 1/2]$ and $1 - Q = 1/2 - (Q - 1/2)$.

These cases partition Ω and in both cases we have $\mathbb{E}[\epsilon] = 1/2 - \mathbb{E}[1/2 - Q]$. It follows that $\text{err}(Q_{1/2}) = 1/2 - \mathbb{E}[|Q - 1/2|]$. \square

Lemma 17. (Uniform case) Let $\hat{Q} \in \mathcal{Q}$ and $\epsilon \stackrel{\text{def}}{=} |\hat{Q} - Y|$ be the random variable of the error of \hat{Q} (according to Lemma 13). If $\mathbb{P}[Q = 1/2] > 0$, then

$$\mathbb{E}[\epsilon | Q = 1/2] = 1/2$$

Proof. Define $r \stackrel{\text{def}}{=} \mathbb{E}[\hat{Q}]$. Let us condition on $Y = 0$ and $Y = 1$ separately. For $Y = 0$, we have $\mathbb{E}[\epsilon | Q = 1/2, Y = 0] = r$, and for $Y = 1$, we have $\mathbb{E}[\epsilon | Q = 1/2, Y = 1] = 1 - r$.

Since $\mathbb{P}[Y = y | Q = 1/2] = 1/2$, we can compute the marginal as

$$\mathbb{E}[\epsilon | Q = 1/2] = (1/2)(r + 1 - r) = 1/2$$

\square

Lemma 18. (Alternative Bayes) The alternative Bayes classifier $Q_{1/2}$ given by $\mathbf{1}_{q(x,a) \geq 1/2}$ (\geq instead of $>$) has also maximal accuracy.

Proof. We will prove that $\text{err}(Q_{1/2}) = \text{err}(Q_{1/2}^*)$. Following Lemma 13, let $\epsilon \stackrel{\text{def}}{=} |Q_{1/2} - Y|$ and $\epsilon^* \stackrel{\text{def}}{=} |Q_{1/2}^* - Y|$.

Conditioned to $Q \neq 1/2$ we have $Q_{1/2} = Q_{1/2}^*$ from their definitions, and thus also $\mathbb{E}[\epsilon - \epsilon^* | Q \neq 1/2] = 0$. It suffices to check the complement event $Q = 1/2$. Suppose $\mathbb{P}[Q = 1/2] > 0$. Conditioned to $Q = 1/2$, Lemma 17 implies that $\mathbb{E}[\epsilon - \epsilon^* | Q = 1/2] = 1/2 - 1/2 = 0$.

Hence $\mathbb{E}[\epsilon] = \mathbb{E}[\epsilon^*]$, i.e. $\text{err}(q_{1/2}) = \text{err}(q_{1/2}^*)$. \square

Lemma 19. (Trivial error as an expectation)

$$\tau = 1/2 + |\mathbb{E}[Y] - 1/2|$$

Proof. The constant 0 predictor ($\hat{0}$) has error $\mathbb{E}[Y]$, while the constant 1 predictor ($\hat{1}$) has error $1 - \mathbb{E}[Y]$. We can rewrite these quantities respectively as $1/2 - (1/2 - \mathbb{E}[Y])$ and $1/2 + (1/2 - \mathbb{E}[Y])$, whose maximum is $\tau = 1/2 + |1/2 - \mathbb{E}[Y]|$. \square

Lemma 20. Let $\vec{P}, \vec{Q} \in (0, 1)^3$, with $\vec{Q}_1 < 1/2$ and $\vec{Q}_2 - \vec{Q}_1 > 0$ (as in Theorem 11).

If \vec{P} and \vec{Q} satisfy also $\vec{Q}_3 + \vec{Q}_1 \geq 1$ and $\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2 < \vec{P}_3\vec{Q}_1$, then $\langle \vec{P}, 2\vec{Q} - 1 \rangle > \langle \vec{Z}, 2\vec{Q} - 1 \rangle$.

Proof. We have the following equivalences and at the end an implication.

$$\begin{aligned} & \langle 2\vec{Q} - 1, \vec{P} \rangle - \langle 2\vec{Q} - 1, \vec{Z} \rangle > 0 \\ \Leftrightarrow & \langle 2\vec{Q} - 1, \vec{P} - \vec{Z} \rangle > 0 \\ \Leftrightarrow & (2\vec{Q}_1 - 1)\vec{P}_1 + (2\vec{Q}_3 - 1)\vec{P}_3 \frac{\vec{P}_1\vec{Q}_1}{\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2} > 0 \\ \Leftrightarrow & (2\vec{Q}_1 - 1)(\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2) + (2\vec{Q}_3 - 1)\vec{P}_3\vec{Q}_1 > 0 \\ \Leftrightarrow & (1 - 2\vec{Q}_1)(\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2) < (2\vec{Q}_3 - 1)\vec{P}_3\vec{Q}_1 \\ \Leftrightarrow & (1 - 2\vec{Q}_1 \leq 2\vec{Q}_3 - 1) \wedge (\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2 < \vec{P}_3\vec{Q}_1) \end{aligned}$$

It is given that $\vec{Q}_3 + \vec{Q}_1 \geq 1$ and $\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2 < \vec{P}_3\vec{Q}_1$, which are equivalent to the last two inequalities. Thus, they imply that $\langle 2\vec{Q} - 1, \vec{P} \rangle - \langle 2\vec{Q} - 1, \vec{Z} \rangle > 0$. \square

Lemma 21. (Complementary part of Theorem 11) Algorithm 2 is correct.

Proof. We will prove $a < b$, $\vec{P}_2 \in (0, 1)$ and the fulfillment of constraints C7 and C10.

Part 1. Proof that $a < b$.

Recall $a = \max\{(1 - \vec{P}_3)\vec{Q}_1, 1/2 - \vec{P}_3\vec{Q}_3\}$ and $b = \min\{(1 - \vec{P}_3)\vec{Q}_2, \vec{P}_3\vec{Q}_1\}$.

- 1) Since $\vec{Q}_1 < 1/2 < \vec{Q}_2$ and $\vec{P}_3 \in (0, 1)$, then $(1 - \vec{P}_3)\vec{Q}_1 < (1 - \vec{P}_3)\vec{Q}_2$.
- 2) Since $\vec{P}_3 \in (1/2, 1)$, then $(1 - \vec{P}_3)\vec{Q}_1 < \vec{P}_3\vec{Q}_1$.
- 3) Since $\vec{P}_3 \in (0, 1)$ and $\vec{Q}_3 \in (1/2, 1)$, then $\vec{P}_3(\vec{Q}_2 - \vec{Q}_3) < 1 \cdot (\vec{Q}_2 - 1/2)$, or equivalently, $1/2 - \vec{P}_3\vec{Q}_3 < (1 - \vec{P}_3)\vec{Q}_2$.
- 4) Since $\vec{P}_3 > \frac{1}{2(\vec{Q}_1 + \vec{Q}_3)}$ then $1/2 - \vec{P}_3\vec{Q}_3 < \vec{P}_3\vec{Q}_1$.

Since the inequalities hold for all available choices for a and b , then, in general, $a < b$ holds.

Part 2. Proof that $\vec{P}_2 \in (0, 1)$.

We know $c > \vec{Q}_1(1 - \vec{P}_3)$ and $c < \vec{Q}_2(1 - \vec{P}_3)$. These inequalities imply that $c - \vec{Q}_1(1 - \vec{P}_3) \in (0, \vec{Q}_2 - \vec{Q}_1)$, hence also that $\vec{P}_2 \in (0, 1)$.

Part 3. Constraint C7 is satisfied.

Since $\vec{P}_1 + \vec{P}_2 = 1 - \vec{P}_3$ and $\vec{Q}_2 > \vec{Q}_1$, then the term $\vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2$ is minimal when $\vec{P}_1 = 1 - \vec{P}_3$ and $\vec{P}_2 = 0$. Thus,

$$\begin{aligned} & \vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2 + \vec{P}_3\vec{Q}_3 \\ & \geq (1 - \vec{P}_3)\vec{Q}_1 + \vec{P}_3\vec{Q}_3 \\ & = \vec{Q}_1 + \vec{P}_3(\vec{Q}_3 - \vec{Q}_1) \\ & > \vec{Q}_1 + \frac{\vec{Q}_3 - \vec{Q}_1}{2} \\ & = \frac{\vec{Q}_3 + \vec{Q}_1}{2} \geq 1/2. \end{aligned}$$

Part 4. Constraint C10 is satisfied.

Since $b \leq \vec{P}_3\vec{Q}_1$, then $\vec{P}_2(\vec{Q}_2 - \vec{Q}_1) < \vec{P}_3\vec{Q}_1 - \vec{Q}_1(1 - \vec{P}_3)$. From this inequality, we may derive constraint C10 as follows.

$$\begin{aligned} & \vec{P}_2(\vec{Q}_2 - \vec{Q}_1) < \vec{P}_3\vec{Q}_1 - \vec{Q}_1(1 - \vec{P}_3) \\ & \vec{P}_2\vec{Q}_2 < (2\vec{P}_3 - 1 + \vec{P}_2)\vec{Q}_1 \\ & \vec{P}_2\vec{Q}_2 < \vec{P}_3\vec{Q}_1 - \vec{P}_1\vec{Q}_1 \\ & \vec{P}_1\vec{Q}_1 + \vec{P}_2\vec{Q}_2 < \vec{P}_3\vec{Q}_1. \end{aligned}$$

\square

Lemma 22. Assume $\mathbb{P}[Y=1, A=a] > 0$ for each $a \in \{0, 1\}$. For any predictor \hat{Q} , we have

$$\mathbb{P}[\hat{Y}=1 | Y=1, A=a] = \frac{\mathbb{E}[\hat{Q}Q | A=a]}{\mathbb{E}[Q | A=a]}$$

hence also

$$\text{oppDiff}(\hat{Q}) = \frac{\mathbb{E}[\hat{Q}Q | A=1]}{\mathbb{E}[Q | A=1]} - \frac{\mathbb{E}[\hat{Q}Q | A=0]}{\mathbb{E}[Q | A=0]}$$

As an additional consequence, by considering the symmetric predictor $1 - \hat{Q}$, it is also true that

$$\mathbb{P}[\hat{Y}=0 | Y=1, A=a] = \frac{\mathbb{E}[(1 - \hat{Q})Q | A=a]}{\mathbb{E}[Q | A=a]}$$

Proof. Indeed, applying repetitively the Bayes rule, we get

$$\begin{aligned} & \mathbb{P}[\hat{Y}=1 | Y=1, A=a] \\ & = \frac{\mathbb{P}[\hat{Y}=1, Y=1, A=a]}{\mathbb{P}[Y=1, A=a]} \\ & = \frac{\mathbb{P}[A=a]\mathbb{E}[\hat{Q}Q | A=a]}{\mathbb{P}[Y=1, A=a]} \\ & = \frac{\mathbb{E}[\hat{Q}Q | A=a]}{\mathbb{P}[Y=1 | A=a]} \\ & = \frac{\mathbb{E}[\hat{Q}Q | A=a]}{\mathbb{E}[Q | A=a]} \end{aligned}$$

The second equality holds because $(Y, \hat{Y}) \perp A | (Q, \hat{Q})$ and $Y \perp \hat{Y} | (Q, \hat{Q})$. \square

G. Python code

Algorithm 2 and figure generation for reproducibility:

```
#!/usr/bin/env python
# coding: utf-8

## Algorithm 1 and Figures
#
# Requirements:
# ...
# pip3 install numpy scipy matplotlib SciencePlots
# ...

import itertools
import numpy as np
import matplotlib.pyplot as plt
from matplotlib.patches import Polygon
from matplotlib.collections import PatchCollection
from scipy.spatial import ConvexHull, convex_hull_plot_2d

plt.style.use('science') # pip3 install SciencePlots

class RandInterval(np.random.RandomState):
    '''Uniform sampling from (a, b)'''

    def open(self, a, b):
        assert a < b
        c = a
        while c == a:
            c = a + self.random() * (b - a)
        return c

def vectorGenerator(seed=None):
    '''Algorithm 1 of the paper. Asserts the 5 conditions'''
    R = RandInterval(seed)
    Q1 = R.open(0, 1 / 2)
    Q2 = R.open(1 / 2, 1)
    Q3 = R.open(1 - Q1, 1)
    P3 = R.open(1 / 2, 1)
    a = max((1 - P3) * Q1, 1 / 2 - P3 * Q3)
    b = min((1 - P3) * Q2, P3 * Q1)
    om = R.open(a, b)
    P2 = (om - Q1 * (1 - P3)) / (Q2 - Q1)
    P1 = 1 - P3 - P2
    P = np.array([P1, P2, P3])
    Q = np.array([Q1, Q2, Q3])
    try:
        # Verify the 5 conditions explicitly
        assert 0 < P1 < 1 and 0 < P2 < 1 and 0 < P3 < 1
        assert 0 < Q1 < 1 and 0 < Q2 < 1 and 0 < Q3 < 1
        assert P1 + P2 + P3 == 1
        assert P1 * Q1 + P2 * Q2 + P3 * Q3 > 1 / 2
        assert Q1 < 1 / 2 and Q2 > 1 / 2 and Q3 > 1 / 2
        assert Q3 + Q1 > 1
        assert P1 * Q1 + P2 * Q2 - P3 * Q1 < 0
    except AssertionError as e:
        e.args = (f'\n{P}={P.round(3)}\n{Q}={Q.round(3)}',)
        raise
    return P, Q

def arbitraryGenerator(seed=None, n=4):
    '''Arbitrary generator for explanation purposes'''
    R = np.random.RandomState(seed)
    P = R.random(n)
    P /= P.sum()
    Q = R.random(n)
    return P, Q

def test_generator(many_times):
    '''Test (run) the vectorGenerator many times'''
    assert many_times > 0
    for _ in range(many_times):
        vectorGenerator(seed=None)
    print(f'{many_times}_tests_passed')
    return

class DataSource:

    def __init__(self, A, P, Q):
        assert len(A) == len(P) == len(Q)
        self.n = len(A)
        self.A = A
        self.P = P
        self.Q = Q

    def print(self):
        print(f'A={self.A.round(3)}')
        print(f'P={self.P.round(3)}')
        print(f'Q={self.Q.round(3)}')

    def err(self, F):
        P = self.P
        Q = self.Q
        return np.sum(F * (1 - 2 * Q)) + np.sum(P * Q)

    def oppDiff(self, F):
        P = self.P
        Q = self.Q
        A = self.A
        f = lambda A: np.sum(F * Q * A) / np.sum(P * Q * A)
        return f(A) - f(1 - A)

    def iter_soft(self, k=11):
        dim = np.linspace(0, 1, k)
        for f in itertools.product(dim, repeat=self.n):
            F = np.array(f) * self.P
            yield F
```

```
def iter_const(self):
    yield self.P * 0
    yield self.P

def iter_hard(self):
    for f in itertools.product([0, 1], repeat=self.n):
        F = np.array(f) * self.P
        yield F

def metrics(self, F_iter):
    l = [(self.err(F), self.oppDiff(F)) for F in F_iter]
    return np.array(l)

def the_plot(self, styles=(), r=None):
    self.print()

    #soft = cls.metrics(self.iter_soft())
    hard = self.metrics(self.iter_hard())
    const = self.metrics(self.iter_const())
    hull = hard[ConvexHull(hard).vertices]

    with plt.style.context(['grid', 'scatter', *styles]):
        #plt.scatter(soft[:,0], soft[:,1], label='Soft', alpha=0.3)
        poly = Polygon(hull, label='Soft', alpha=0.25)
        plt.gca().add_patch(poly)
        plt.scatter(hard[:, 0], hard[:, 1], marker='.', label='Hard')
        plt.scatter(const[:, 0], const[:, 1], marker='x', label='Constant')
        plt.xlabel('error')
        plt.ylabel('opportunity-difference')
        plt.xlim((0, 1))
        if r is not None and r > 0:
            x0 = const[:, 0].min()
            y0 = const[const[:, 0] == x0, 1][0]
            if y0 != 0:
                print(y0)
            else:
                plt.xlim((x0 - r, x0 + r))
                plt.ylim((y0 - r, y0 + r))
        plt.legend()
        plt.show()
    return

test_generator(many_times=10**4)

# Arbitrary non-example for Section "Strong impossibility result"
A = np.array([0, 1, 0, 1])
P, Q = arbitraryGenerator(seed=3, n=4)
DataSource(A, P, Q).the_plot(styles=['ieee'])

# Generated example for Theorem "Impossibility result"
A = np.array([0, 0, 1])
P, Q = vectorGenerator(seed=0)
DataSource(A, P, Q).the_plot(styles=['ieee'])

# First example we found. Appears as "Example 1"
A = np.array([0, 1, 0, 1])
P = np.array([3, 2, 1, 2]) / 8.
Q = np.array([9, 15, 15, 16]) / 20.
DataSource(A, P, Q).the_plot(styles=['ieee'])

# Other arbitrary (not shown in paper)
A = np.array([0, 1, 0, 1, 0, 1, 0, 1])
P = np.array([1, 1, 1, 1, 1, 1, 1, 1]) / 8.
Q = np.array([1, 1, 1, 1, 1, 3, 3]) / 4
DataSource(A, P, Q).the_plot(styles=['ieee'])
```