

A Possibility in Algorithmic Fairness: Can Calibration and Equal Error Rates Be Reconciled?

Claire Lazar Reich ✉

MIT Statistics Center and Department of Economics, Cambridge, MA, USA

Suhas Vijaykumar ✉

MIT Statistics Center and Department of Economics, Cambridge, MA, USA

Abstract

Decision makers increasingly rely on algorithmic risk scores to determine access to binary treatments including bail, loans, and medical interventions. In these settings, we reconcile two fairness criteria that were previously shown to be in conflict: calibration and error rate equality. In particular, we derive necessary and sufficient conditions for the existence of calibrated scores that yield classifications achieving equal error rates at any given group-blind threshold. We then present an algorithm that searches for the most accurate score subject to both calibration and minimal error rate disparity. Applied to the COMPAS criminal risk assessment tool, we show that our method can eliminate error disparities while maintaining calibration. In a separate application to credit lending, we compare our procedure to the omission of sensitive features and show that it raises both profit and the probability that creditworthy individuals receive loans.

2012 ACM Subject Classification Mathematics of computing → Probability and statistics; Social and professional topics → Computing / technology policy; Computing methodologies → Supervised learning

Keywords and phrases fair prediction, impossibility results, screening decisions, classification, calibration, equalized odds, optimal transport, risk scores

Digital Object Identifier 10.4230/LIPIcs.FORC.2021.4

Related Version *Full Version*: <https://arxiv.org/abs/2002.07676>

Funding *Claire Lazar Reich*: Supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374.

Acknowledgements Many thanks to Anna Mikusheva, Iván Werning, and David Autor for their valuable advice. We're also deeply grateful for the support of Ben Deaner, Lou Crandall, Pari Sastry, Tom Brennan, Jim Poterba, Rachael Meager, and Frank Schilbach with whom we have had energizing and productive conversations. Thank you to Deborah Plana, Pooya Molavi, Adam Fisch, and Yonadav Shavit for commenting on the manuscript at its advanced stages.

1 Introduction

Today's algorithms reach deep into decisions that guide our lives, from loan approvals to medical treatments to foster care placements. Making these high-impact decisions fairly is an effort undergoing public scrutiny. In one investigation, *ProPublica* showed that an algorithm operating in the U.S. criminal justice system, COMPAS, discriminated against black defendants by misclassifying them as high-risk at significantly higher rates than white defendants [2]. On the other hand, it was later revealed that the same algorithm did satisfy a different form of fairness: calibration of scores for both black and white defendants [10]. This meant that on average, a defendant's score reflected the same risk level regardless of race.

Researchers have sought to explain how a screening algorithm like COMPAS can satisfy one natural notion of fairness but not another, spurring a research agenda to characterize how definitions of algorithmic fairness relate to one another. Multiple studies in this literature



© Claire Lazar Reich and Suhas Vijaykumar;
licensed under Creative Commons License CC-BY 4.0
2nd Symposium on Foundations of Responsible Computing (FORC 2021).

Editors: Katrina Ligett and Swati Gupta; Article No. 4; pp. 4:1–4:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

proved that algorithms face inevitable tradeoffs whenever they predict on groups that have different average outcomes [15, 5, 4, 7, 14]. These influential “impossibility results” have underscored the need for practitioners to target certain fairness criteria at the expense of others.

We show that it is in fact possible to reconcile the two notions of fairness that gained influence following the COMPAS investigation: calibration and equal error rates. In important previous work, these two criteria were proven to be mutually incompatible when both are applied to a *risk score* [15, 22] and when both are applied to a *classifier* [5]. Naturally these findings were interpreted as evidence that calibration and equal error rates are incompatible altogether [1]. It was therefore speculated that COMPAS’s enforcement of score calibration made its error rate imbalances inevitable [5].

In contrast, we show that both calibration and equal error rates can be reconciled in COMPAS and in many other real-world settings where protected groups have different mean outcomes. We relax the mathematical tension between these two fairness criteria by separately enforcing *calibration on the score* and *equal error rates on the corresponding classifier*. In particular, we prove that it is possible to design calibrated scores that yield equal error rate classifications at group-blind cutoffs, and we provide a method to do so with maximal accuracy. Furthermore, we develop practical extensions of the method, such as showing how to enforce weaker notions of the equal error rate criterion (like the “equality of opportunity” criterion of Hardt et al. [12]) and how to accommodate multiple protected subgroups.

Our framework and method can be applied to two settings. In the first, we consider the problem of providing risk scores to a profit-maximizing third-party agent, such as a lender, who then uses them to assign binary treatments, such as loan approvals and denials. We illustrate how to construct calibrated scores that lead this profit-maximizer to make classifications satisfying equal error rates. In the second setting, we consider risk assessments like COMPAS that output both scores and classification recommendations, and show that the scores can be made to satisfy calibration while the classification recommendations can be made to satisfy equal error rates.

This paper supports growing evidence on the complementary relationship between data quality and fairness objectives [11, 8, 9, 6, 13, 14]. In particular, we show that access to sufficiently informative features is required to satisfy our fairness criteria, and that the feasible set of solutions grows with the informativeness of the data. In an empirical credit lending example, we compare our method to a commonly practiced strategy of data omission. It yields higher lender profit while also improving access to loans for creditworthy applicants in all groups.

The results proceed as follows. In Section 2, we prove that it is possible to construct calibrated scores that lead to equal error rate classifications and we precisely characterize when such scores exist. In Section 3, we propose an algorithm that produces the most accurate possible score satisfying the fairness criteria and minimizing the decision-maker’s errors. We apply our method in Section 4 to two empirical settings. We first assess its performance in helping a lender screen loan applicants of various educational backgrounds. We also apply the method to the COMPAS criminal risk assessment tool, where we show that our procedure can eliminate error rate imbalances in risk classifications while preserving calibration of scores.

1.1 Related Work

Our paper belongs to a body of work that studies the mathematical relationships between various individual and group measures of fairness. Calibration and equal error rates have been formalized and extensively studied in prior work [22, 12, 15, 5]. In particular, Kleinberg et al. [15] and Pleiss et al. [22] show that these criteria are incompatible when applied to a risk score and Chouldechova [5] shows the corresponding result for binary classifiers. We consider a natural variation of the problem where we ask whether a calibrated score can, upon being supplied to a rational third-party, lead to equal-error predictions. Surprisingly, we find that the answer is yes.

Our work also contributes to a recent strand of the literature which studies how algorithmic prediction can interact with self-interested decision makers, bridging the classical problem of prediction with the traditionally economic problem of information design [20, 24]. From this perspective, we study the existence of scores that lead to desirable equilibria: those in which the final decision rule is group-blind due to calibration, and the resulting decisions satisfy equal error rates.

Finally, we believe it is important to emphasize that the two fairness criteria we study do not encompass all notions of fairness. Tradeoffs remain between these criteria and others. For example, enforcing equal error rates requires that the classifications' positive and negative predictive values will be unequal across groups, meaning that one groups' scores would carry greater signal to the decision-maker than the others' [5]. In addition, equal error rate classifications will generically require changes to the Bayes' optimal classifications, and enforcing calibration does not diminish this requirement [6].

Decisions for how to prioritize fairness conditions are likely to vary by application going forward. We hope that by clarifying the precise relationship between two influential criteria, we can facilitate these decisions, and that in settings where calibration and equal error rates are considered essential, our algorithm can help yield accurate predictions and fairer outcomes.

2 Theoretical Results

2.1 Formal Setting

Let us consider a triple (Y, X, A) on a common probability space \mathbb{P} , where $Y \in \{0, 1\}$ is an outcome variable, $X \in \mathbb{R}^d$ is a vector of features, and $A \in \{H, L\}$ is a protected attribute differentiating two groups with unequal base rates $\mu_A = \mathbb{E}[Y|A]$ of the outcome,

$$\mu_L < \mu_H. \tag{1}$$

Our goal is to estimate a score function $\hat{p} \equiv \hat{p}(X, A) \in [0, 1]$ that predicts Y with maximum accuracy subject to the constraints of calibration and equal error rates. Specifically, we hand \hat{p} to a decision-maker tasked with selecting classifications $\hat{y} \in \{0, 1\}$ that minimize their loss function

$$\ell(\hat{y}, y) = \begin{cases} 0 & y = \hat{y} \\ 1 & y > \hat{y} \\ k & y < \hat{y}, \end{cases} \tag{2}$$

where $k > 0$ is the relative cost of false positive classifications. Note that any loss function that is minimized when $y = \hat{y}$ is equivalent to ℓ after an affine transformation.

4:4 Calibrated Scores and Equal Error Classifiers

Let us suppose the decision-maker might be able to observe group affiliation A in addition to \hat{p} . To ensure that classifications are based only on \hat{p} and not on A , we constrain \hat{p} to satisfy *calibration within groups*,

$$\mathbb{E}[Y|A, \hat{p}] = \mathbb{E}[Y|\hat{p}] = \hat{p}. \quad (3)$$

If (3) holds, the decision-maker's expected loss given \hat{p} and A becomes

$$\mathbb{E}[\ell(Y, \hat{y})|\hat{p}, A] = \hat{p}(1 - \hat{y}) + k(1 - \hat{p})\hat{y}. \quad (4)$$

This expected loss is minimized with a cutoff decision rule that is independent of group affiliation A ,

$$\hat{y} = \mathbb{1}\{\hat{p} \geq \bar{p}\}, \quad (5)$$

where the cutoff $\bar{p} = k/(k+1)$ is fixed by the decision-maker's loss function.

Our second condition constrains \hat{y} to satisfy *equal error rates*, ensuring that the classification only depends on the group through the target variable. Following the decision rule (5), we may write this as

$$(\mathbb{1}\{\hat{p} \geq \bar{p}\} \perp\!\!\!\perp A) | Y. \quad (6)$$

Our calibration and equal error rate conditions are summarized by (3) and (6), respectively.

2.2 Relation to Impossibility Results


We first introduce a general impossibility result, relate it to previous work, and show where our assumptions diverge to make our proposed criteria satisfiable. The following theorem proves that a *single* algorithmic output Z cannot generally satisfy notions of both calibration and equal error rates.

► **Theorem 1.** *Let Y, A , and Z be random variables satisfying the following three conditions.*

- (i) $(Y \perp\!\!\!\perp A) | Z$,
- (ii) $(Z \perp\!\!\!\perp A) | Y$,
- (iii) $\mathbb{P}(A = H|Z), \mathbb{P}(Y = 1|A, Z) \in (0, 1)$.

Then A and (Z, Y) must be independent.

Proof. Suppose that (Y, A, Z) satisfy (i) (ii) and (iii). Assumption (iii) implies that the law of (A, Y, Z) is strictly positive. By the Hammersley-Clifford theorem (see e.g. [18]), the conditional independence relations are summarized by a graph on $\{Y, A, Z\}$ where every path from Y to A travels through Z , and every path from A to Z travels through Y . There are only two graphs with this property:



In neither of these graphs does there exist a path from A to (Y, Z) , so we conclude that A and (Y, Z) must be independent for (i) (ii) and (iii) to simultaneously hold. ◀

Note that when A denotes group affiliation and Y denotes outcomes, (i) is a form of calibration and (ii) is a form of the equal error rate condition. Assumption (iii) is a strong form of predictive uncertainty that is generalized in the appendix. Thus the theorem shows that

when there is predictive uncertainty and Y depends on A (i.e. when the base rates are unequal), it is impossible for a single Z to satisfy both calibration and equal error rates. For example, letting Z be a classifier recovers the result of Chouldechova that (i) equal positive and negative predictive values are unachievable alongside (ii) equal error rates [5]. Meanwhile, letting Z be a risk score shows that (i) calibration is unachievable alongside (ii) a condition that implies balance in the positive and negative class, similar to the result of Kleinberg et al. [15].

Our own setting bypasses the mathematical impossibility described in Theorem 1 by imposing constraints on *two* separate algorithmic outputs rather than one. We require (i) calibration from the scores \hat{p} and (ii) equal error rates from the resulting classifications $\hat{y} = \mathbb{1}\{\hat{p} \geq \bar{p}\}$.

2.3 Necessary and Sufficient Conditions

In this section we characterize exactly when there exists a calibrated \hat{p} that leads to equal error rate classifications \hat{y} at the cutoff \bar{p} . Our conditions can be easily checked in a given setting, and they are shown to depend on the informativeness of the features X .

The graphical framework in this section builds on methods developed by Hardt et al. [12]. All the necessary and sufficient conditions will be illustrated in \mathbb{R}^2 , with true positive rates on the vertical axis and false positive rates on the horizontal. The *feasible region* will be the set in \mathbb{R}^2 corresponding to error rates achievable by an equal error rate classifier $\hat{y} = \mathbb{1}\{\hat{p} \geq \bar{p}\}$ where \hat{p} is calibrated.

We first study the entire set corresponding to equal error rate classifiers, without regard to calibration or the decision-maker's cutoff \bar{p} . Then we study the entire set corresponding to classifiers that can be based on the cutoff \bar{p} applied to calibrated scores, without regard to the equal error rate condition. Finally, we prove that the intersection of these two sets determines feasibility of enforcing both conditions, and we characterize when the intersection is nonempty.

2.3.1 Classifiers Satisfying Equal Error Rates


We wish to identify the entire set of error rates in \mathbb{R}^2 achievable by classifiers with equal error rates. Hardt et al. [12] succeeded in doing so, and we review and adapt their results in this subsection. To lay the groundwork for the geometric reasoning to follow, we first denote the group A false positive rate and true positive rate associated with a given classifier \hat{y} as a point in \mathbb{R}^2 ,

$$\alpha(\hat{y}, A) = \left(\mathbb{P}(\hat{y} = 1 | Y = 0, A), \mathbb{P}(\hat{y} = 1 | Y = 1, A) \right).$$

We may now define the set of achievable error rates in \mathbb{R}^2 . Let \mathcal{H} be the set of all possibly random classifiers $h(X, A)$. The set of achievable error rates for group A is

$$S(A) = \{\alpha(\hat{y}, A) \mid \hat{y} = h(X, A), h \in \mathcal{H}\} \subseteq \mathbb{R}^2, \quad (7)$$

and the set of achievable rates for all classifiers satisfying equal error rates is given by $S(L) \cap S(H)$. To better understand this intersection, we characterize $S(A)$ in terms of Receiver Operator Characteristic (ROC) curves following Hardt et al. [12]. By definition, an ROC curve of a given score p traces the true and false positive rates associated with each possible cutoff rule $\mathbb{1}\{p \geq c\}$ for $c \in [0, 1]$. Therefore it contains all points $\alpha(\mathbb{1}\{p \geq c\}, A)$. With these tools in hand, we are ready to characterize the feasible set of rates $S(A)$ for group A .



■ **Figure 1** Achievable equal error rates (shaded). ■ **Figure 2** Achievable equal error rates from calibrated score at cutoff \bar{p} (shaded). The restrictions (11) correspond to half-spaces above the red dashed lines.

► Proposition 2. Let $p^* = p^*(X, A)$ be the Bayes optimal score satisfying $p^* = \mathbb{E}[Y|X, A]$, i.e., the best score given our data. Then the set of achievable rates $S(A)$ is exactly the convex hull of the union of the group- A ROC curve of the best score p^* and the group- A ROC curve of the worst score $1 - p^*$, i.e. the convex hull of

$$\begin{aligned}
 & \left\{ \alpha(\mathbb{1}\{p^* \geq c\}, A) \mid 0 \leq c \leq 1 \right\} \\
 & \cup \left\{ (1, 1) - \alpha(\mathbb{1}\{p^* \geq c\}, A) \mid 0 \leq c \leq 1 \right\}.
 \end{aligned}$$

Figure 1 illustrates typical examples of $S(L)$, $S(H)$, and the intersection $S(L) \cap S(H)$ which represents the rates achievable by equal error rate classifiers.

2.3.2 Classifiers Compatible with Calibration

We now put aside the equal error rate constraint and concentrate on identifying the entire set of classifiers that are implementable with the cutoff \bar{p} applied to some calibrated scores \hat{p} . The set is characterized by the following proposition.

► Proposition 3. A classifier \hat{y} can be written as $\hat{y} = \mathbb{1}\{\hat{p} \geq \bar{p}\}$ for some calibrated \hat{p} if and only if its group-specific positive predictive values exceed \bar{p} , and its group-specific negative predictive values exceed $1 - \bar{p}$. In particular, for $A \in \{L, H\}$,

$$\mathbb{P}(Y = 1|\hat{y} = 1, A) \geq \bar{p}, \quad \mathbb{P}(Y = 0|\hat{y} = 0, A) > 1 - \bar{p}. \tag{8}$$

Proof. Suppose that $\hat{y} = \mathbb{1}\{\hat{p} \geq \bar{p}\}$ where \hat{p} is calibrated. Then \hat{y} must satisfy the inequalities

$$\begin{aligned}
 \mathbb{P}(Y = 1|\hat{y} = 1, A) &= \mathbb{E}[Y|\hat{p} \geq \bar{p}, A] \\
 &= \mathbb{E}[\hat{p}|\hat{p} \geq \bar{p}, A] \geq \bar{p},
 \end{aligned} \tag{9}$$

$$\mathbb{P}(Y = 1|\hat{y} = 0, A) = \mathbb{E}[\hat{p}|\hat{p} < \bar{p}, A] < \bar{p}. \tag{10}$$

Therefore, if \hat{y} is based on a calibrated score \hat{p} at cutoff \bar{p} , then it is necessary for the group-specific positive and negative predictive values to exceed \bar{p} and $(1 - \bar{p})$, respectively.

Conversely, given *any* classifier \hat{y} that satisfies the inequalities (9) and (10), we can always put

$$\hat{p}(\hat{y}, A) = \mathbb{P}(Y = 1 | \hat{y}, A)$$

to obtain a calibrated score that takes just two possible values per group with the cutoff \bar{p} guaranteed to be between them. This choice of \hat{p} thus satisfies $\hat{y} = \mathbb{1}\{\hat{p} \geq \bar{p}\}$ by construction. \blacktriangleleft

As we will see in the following subsection, this result lays the foundation for the necessary and sufficient conditions for the satisfiability of our fairness criteria.

2.3.3 The Feasibility Region

Proposition 3 demonstrates that the following are equivalent:

- (i) There exists a calibrated score \hat{p} such that $\hat{y} = \mathbb{1}\{\hat{p} \geq \bar{p}\}$ satisfies equal error rates.
- (ii) There exists a classifier \hat{y} satisfying equal error rates and (8).

In practice, we propose checking (ii) to identify whether (i) holds. To do so, we use Bayes' rule to write (8) as group-specific restrictions on true and false positive rates so that we can consider them in the same space as the equal error rate constraints given by Hardt et al. [12]. The following theorem and the accompanying Figure 2 indicate that each restriction (8) corresponds to a half-space in \mathbb{R}^2 , and that the feasible region corresponds to the intersection of those half-spaces with each other and with the equal error rates region $S(L) \cap S(H)$.

► **Theorem 4.** *Let $\beta_A = \mu_A / (1 - \mu_A)$ denote the group-specific odds ratios, with $\beta_L < \beta_H$. Then our fairness criteria are simultaneously satisfiable at cutoff \bar{p} if and only if there exists $(\alpha_1, \alpha_1) \in S(L) \cap S(H)$ satisfying the two inequalities*

$$\frac{\alpha_2}{\alpha_1} \geq \frac{\bar{p}}{\beta_L(1 - \bar{p})}, \quad \frac{(1 - \alpha_1)}{(1 - \alpha_2)} > \frac{\beta_H(1 - \bar{p})}{\bar{p}}. \quad (11)$$

We next provide easily checkable necessary and sufficient conditions for when the feasible region is nonempty.

► **Corollary 5.** *Let $(\check{\alpha}_1, \check{\alpha}_2)$ denote the point at which the inequalities (11) hold with equality. Our fairness criteria are simultaneously satisfiable at cutoff \bar{p} if and only if any of the following holds: $\check{\alpha}_1 \leq 0$, $\check{\alpha}_1 \geq 1$, or both groups' ROC curves corresponding to p^* lie above $(\check{\alpha}_1, \check{\alpha}_2)$. Note that $(\check{\alpha}_1, \check{\alpha}_2)$ are fixed by the group base rates and decision-maker's cutoff \bar{p} ,*

$$\check{\alpha}_1 = \frac{\beta_L}{(\beta_H - \beta_L)} \left(\frac{\beta_H - (1 + \beta_H)\bar{p}}{\bar{p}} \right), \quad \check{\alpha}_2 = \frac{1}{(\beta_H - \beta_L)} \left(\frac{\beta_H(1 - \bar{p}) - \bar{p}}{1 - \bar{p}} \right). \quad (12)$$

We note that the feasible region depends on the decision-maker's cutoff \bar{p} , which in turn depends on their relative valuation of false positive and false negative classifications, k . In particular, when k is either very large or close to 0, the set of feasible error rates shrinks to include only those corresponding to no positive classifications or no negative classifications.

Data quality also contributes to the feasibility of enforcing both fairness criteria, as illustrated by Theorem 4 and Figure 2. Note that the intersection of the half-spaces defined in (11) are fixed by given parameters: β_L , β_H , and \bar{p} . Beyond these, what determines the size of the feasible region is the height of the ROC curves.

Higher ROC curves correspond to more accurate predictions, which can be achieved by including more informative features X . This expands the region $S(H) \cap S(L)$ and thus always weakens the constraints dictating whether equal error rates and calibration are compatible in a given setting. Therefore, increasing the quality of data that an algorithm can access promotes our notions of fairness, whereas removing data compromises them.

3 A Loss-Minimizing Algorithm

After checking that our fairness criteria are feasible in a given setting, a natural next step is to search for the constrained optimal solution, i.e. to identify the most accurate score \hat{p} that minimizes the decision-maker's loss subject to our fairness constraints. Our strategy is to first estimate the most accurate score $p^* = \mathbb{E}[Y|X, A]$ without regard to fairness, and then to transform the estimate in two separate stages. First, we identify the error rates that minimize loss subject to the fairness conditions (Section 3.1). Second, we identify the MSE-minimizing calibrated scores \hat{p} that gives rise to those error rates at the decision-maker's cutoff \bar{p} (Section 3.2). Lastly, we lay out extensions of the algorithm that can accommodate practical use cases (Section 3.3).

3.1 Stage 1: Error Rate Optimization

The first stage of the algorithm identifies feasible error rates that minimize the decision maker's loss.

Let R denote the set of points (α_1, α_2) in the feasible region, i.e. the pairs of error rates in $S(H) \cap S(L)$ that satisfy (11). Note that R is necessarily convex, as it is the intersection of four convex regions: $S(H)$, $S(L)$, and the half-spaces defined in (11). Moreover, according to the decision-maker's loss function, a classifier corresponding to error rates (α_1, α_2) obtains expected loss

$$\ell(\alpha_1, \alpha_2) \equiv k\alpha_1(1 - \mathbb{E}[Y]) + (1 - \alpha_2)\mathbb{E}[Y]. \quad (13)$$

Thus, straightforward convex optimization will identify the error rates that minimize the linear function ℓ over $(\alpha_1, \alpha_2) \in R$. The optimal error rates identified, $z^* = (\alpha_1^*, \alpha_2^*)$, will be on the upper-left boundary of the feasible region in Figure 2, with the precise point determined by the decision-maker's relative preference k over false positive and false negative classifications.

■ **Algorithm Stage 1** Find loss-minimizing feasible error rates.

Input: Raw scores $\{p_i^*\}$, labels $\{Y_i\}$, group identities $\{A_i\}$, base rates μ_A , cutoff \bar{p} , loss parameter k .

Step 1: Define convex feasible region R by taking intersection of rates (α_1, α_2) in $S(L) \cap S(H)$ that satisfy (11). To compute $S(L) \cap S(H)$, use $\{p_i^*\}$ to determine each group's ROC curves.

if R is empty **then**

Output: No feasible solution.

end if

Step 2: Minimize loss function (13) over $(\alpha_1, \alpha_2) \in R$.

Output: Optimal target rates (α_1^*, α_2^*) from Step 2.

► **Remark 6.** The sets $S(L)$ and $S(H)$ correspond to the Bayes optimal score $p^* = \mathbb{E}[Y|X, A]$, which needs to be estimated in practice. Given an estimated score p , we propose using a holdout sample to first calibrate p and then perform our algorithm. The resulting scores will satisfy the fairness criteria approximately by a law-of-large-numbers argument, where the fidelity is determined solely by the size of the holdout sample (see e.g. [27]).

3.2 Stage 2: Risk Score Optimization

Once a feasible set of error rates is chosen, the decision-maker's expected loss is determined. However, multiple choices of calibrated scores may achieve those target rates at the cutoff \bar{p} , and we expect that in practice, decision-makers would prefer more accurate scores. This section thus describes a method to recover the MSE-minimizing score \hat{p} that implements the target rates z^* by solving a constrained optimal transport problem [21].

We base the method on the finding that the best \hat{p} satisfying the fairness criteria is recoverable through post-processing the Bayes optimal score $p^* = \mathbb{E}[Y|X, A]$. We include a proof for this in the appendix, following a similar argument of Hardt et al. [12]. In the appendix we also discuss how our procedure can be thought of as finding the smallest *mean-preserving contraction* of p^* that yields the targeted error rates. Readers will note that the post-processing procedure requires some randomization of input scores. We explore the effects of the randomization empirically in our online appendix [23], and meanwhile highlight that our algorithm's accuracy objective limits the extent to which scores p^* change.

Our method defines one linear program per group A and seeks the most accurate \hat{p}_A that yields error rates at the cutoff \bar{p} given by

$$\alpha(\mathbb{1}\{\hat{p}_A \geq \bar{p}\}, A) = z^* = (\alpha_1^*, \alpha_2^*).$$

For the remainder of the section, we simplify notation by suppressing A subscripts and note that the procedure is performed once for each group $A \in \{H, L\}$.

Our approach will involve a transformation kernel, or *transport map*, that maps the distribution of the most accurate estimate of p^* to the distribution of our post-processed \hat{p} . We assume for simplicity that the p^* estimate has already been calibrated, and that it is discrete (which we justify in the appendix). In particular, p^* takes N ordered values $p = (p_1, p_2, \dots, p_N)$, each with probability mass given by $s = (s_1, s_2, \dots, s_N)$ where $\sum_i s_i = 1$. Furthermore, we will denote the post-processed \hat{p} as taking those same discrete values p but with different probability masses that we seek to optimize, $f = (f_1, f_2, \dots, f_N)$.

We call T the matrix that maps probability masses from the discrete distribution of p^* to that of \hat{p} . In particular, with probability T_{ij} , the kernel will map an individual with score p_i to the output score p_j . Therefore, the probability distribution of \hat{p} will be determined by

$$T^\top s = f. \tag{14}$$

In order to produce probability distributions, T must be right-stochastic: elements must take values between 0 and 1, and each row should sum to 1.

$$0 \leq T_{ij} \leq 1 \text{ and } \sum_{k=1}^N T_{ik} = 1 \quad \forall i, j \in \{1, \dots, N\}. \tag{15}$$

According to our fairness criteria, we further constrain T . To ensure that \hat{p} will be calibrated, we need the outcome of individuals assigned score f_j to satisfy $Y = 1$ with probability p_j . Given our assumption that p^* has itself been calibrated, this reduces to

$$\sum_{i=1}^N T_{ij} p_i s_i = p_j f_j \quad \forall j \in \{1, \dots, N\}. \tag{16}$$

4:10 Calibrated Scores and Equal Error Classifiers

The targeted false- and true-positive rates (α_1^*, α_2^*) derived in Section 3.1 similarly require:

$$\begin{aligned} \sum_{j=1}^N \sum_{i=1}^N T_{ij} p_i s_i (\mathbb{1}\{p_j \geq \bar{p}\} - \alpha_2^*) &= 0, \\ \sum_{j=1}^N \sum_{i=1}^N T_{ij} (1 - p_i) s_i (\mathbb{1}\{p_j \geq \bar{p}\} - \alpha_1^*) &= 0. \end{aligned} \quad (17)$$

Finally, we formulate an objective. Note that the mean-squared error of \hat{p} satisfies the bias-variance decomposition

$$\mathbb{E}[(\hat{p} - Y)^2] = \mathbb{E}[(\hat{p} - \mathbb{E}[Y|X, A])^2] + \mathbb{E}[(Y - \mathbb{E}[Y|X, A])^2],$$

and thus the \hat{p} that minimizes the left hand side is obtained by minimizing the first term on the right hand side. In particular, if the input score p^* is $\mathbb{E}[Y|X, A]$, then the post-processed score that minimizes mean-squared error will also minimize

$$\mathbb{E}[(\hat{p} - p^*)^2] = \sum_{i=1}^N \sum_{j=1}^N T_{ij} (p_i - p_j)^2 s_i. \quad (18)$$

Furthermore, even if p^* is not exactly equal to $\mathbb{E}[Y|X, A]$, the triangle inequality in $L^2(\mathbb{P})$ implies

$$\mathbb{E}[(\hat{p} - Y)^2]^{\frac{1}{2}} \leq \mathbb{E}[(p^* - Y)^2]^{\frac{1}{2}} + \mathbb{E}[(\hat{p} - p^*)^2]^{\frac{1}{2}}.$$

Thus, by minimizing the objective (18) we can effectively control the additional error due to post-processing. Combining this with the above constraints yields a straightforward linear program.

■ **Algorithm Stage 2** For each group, find calibrated scores achieving target rates.

Input: Raw scores $\{p_i^*\}$, number of bins N , target error rates from stage 1 of algorithm (α_1^*, α_2^*) , cutoff \bar{p} .

Step 1: Produce discrete score approximation of p^* : label N ordered values (p_1, p_2, \dots, p_N) with masses (s_1, s_2, \dots, s_N) .

Step 2: Find score transformation kernel T that minimizes (18) subject to the constraints (14), (15), (16) and (17).

Step 3: Map each individual with given raw score to a new post-processed score, based on probabilities given by kernel T .

Output: Scores \hat{p} from Step 3. By design these satisfy calibration and yield error rates (α_1^*, α_2^*) at cutoff \bar{p} .

3.3 Available Extensions

Our procedure can be modified to handle additional use cases. We can flexibly trade off the fairness and accuracy objectives, minimize error disparities rather than eliminate them when the feasible region is empty, accommodate a setting where the decision-maker's cutoff \bar{p} is estimated with error, and apply the procedure to more than two groups.

3.3.1 Relaxing the fairness criteria

An alternative formulation of our algorithm can accommodate multiple cases encountered in practice. By modifying Stage 1 to include a weighted error-rate penalty, users can flexibly trade off the fairness and accuracy objectives, minimize error disparities rather than eliminate them when the feasible region R is empty, and enforce just one error constraint as in the “equality of opportunity” criterion of Hardt et al. [12]. In general, the more flexible procedure will output group-specific optimal error rates: z_L^* and z_H^* . These group-specific targets are then inputted into Stage 2 which is otherwise unchanged.

To modify Stage 1, first we define a broader domain for the algorithm to search over in place of R . It contains all the error rates implementable by a calibrated score at the decision-maker’s cutoff, according to the inequalities (11), without regard to equal error rates. The domain is $R(H) \times R(L)$ where

$$R(A) = \left\{ (\alpha_1, \alpha_2) \in S(A) \mid \frac{1 - \alpha_2}{1 - \alpha_1} < \frac{\bar{p}/\beta_A}{(1 - \bar{p})} \leq \frac{\alpha_2}{\alpha_1} \right\}. \quad (19)$$

(Note that this is guaranteed to be nonempty, as it contains the error rates of the classifier $\mathbb{1}\{p^* \geq \bar{p}\}$.) We also replace the loss function (13) with a generalized version that includes both the decision-maker’s expected loss from the error rates as well as the groups’ rate disparities. The new loss function is

$$\gamma \ell(z_L) + (1 - \gamma) \ell(z_H) + (z_L - z_H)^\top \Lambda (z_L - z_H) \quad (20)$$

where $\ell(z_A)$ is the decision-maker’s expected loss $k\alpha_{1A}(1 - \mathbb{E}[Y|A]) + (1 - \alpha_{2A})\mathbb{E}[Y|A]$ and γ is the fraction of individuals in group L . Meanwhile, Λ is a positive semidefinite matrix that provides the flexibility of varying the enforcement of minimal error rate differences. For example, taking $\Lambda = \lambda I$ for arbitrarily large λ recovers the equal error rate solution when the feasible region R is nonempty, and otherwise outputs the solution that minimizes error rate disparities. Meanwhile a small choice of λ places relatively more weight on accuracy.

Alternatively, Λ could be chosen so that differences in the true and false positive rates are weighted differently. For example, we can achieve equal true positive rates and allow false positive rates to vary [12] by letting $\Lambda(2, 2)$ be large and assigning 0 to all other entries in Λ .

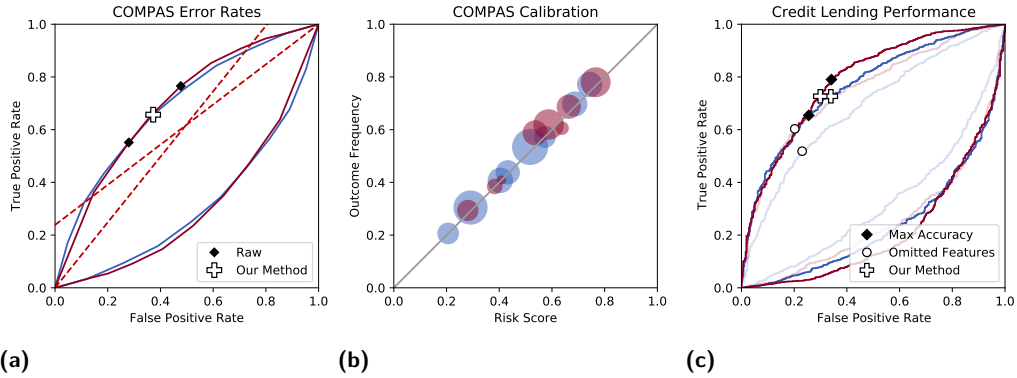
As a result of the flexible procedure, group-specific error rates z_L^* and z_H^* are identified to minimize the generalized loss function (20). The second stage of the algorithm can then be applied to identify a calibrated score that yields those target rates.

3.3.2 Accommodating an interval of possible k or \bar{p}

In settings where the exact \bar{p} is unknown or not fixed, users can adapt our algorithm to function for any cutoff in an interval $(\bar{p} - \epsilon, \bar{p} + \epsilon)$. It can be tailored to produce scores \hat{p} that are either below $\bar{p} - \epsilon$ or above $\bar{p} + \epsilon$, so that any cutoff applied within the interval would execute the same classifications.

In particular, we propose a couple modifications to generalize our algorithm to this setting. We wish for anyone receiving scores above $\bar{p} + \epsilon$ to be classified as $\hat{y} = 1$ and anyone receiving scores below $\bar{p} - \epsilon$ as $\hat{y} = 0$. Following the reasoning in Proposition 3, for such a score to be calibrated, the associated PPV should exceed $\bar{p} + \epsilon$ and the NPV should exceed $1 - (\bar{p} - \epsilon)$. Therefore, the feasible region previously defined in Theorem 4 by (11) is now defined by the points $(\alpha_1, \alpha_1) \in S(L) \cap S(H)$ that satisfy

$$\frac{\alpha_2}{\alpha_1} \geq \frac{\bar{p} + \epsilon}{\beta_L(1 - (\bar{p} + \epsilon))}, \quad \frac{(1 - \alpha_1)}{(1 - \alpha_2)} > \frac{\beta_H(1 - (\bar{p} - \epsilon))}{\bar{p} - \epsilon}. \quad (21)$$



■ **Figure 3** Evaluating algorithm performance. In each figure, maroon represents the high-mean group while blue represents the low-mean group. Panels (a) and (b) correspond to the criminal justice application, showing respectively that we can eliminate error rate disparities and maintain score calibration in COMPAS. Note that we define a true positive classification as correctly identifying someone who would not reoffend. Panel (c) covers the credit lending application, illustrating the empirical ROC curves from the rich feature set (opaque) and the limited feature set (translucent). Compared to a data omission strategy, our method raises the probability that creditworthy individuals from all education groups access loans.

This feasible region is used in Stage 1. In Stage 2, we add another constraint to specify that no post-processed scores be assigned values inside the interval of possible cutoffs:

$$T_{ik} = 0 \quad \forall k \text{ such that } p_k \in (\bar{p} - \epsilon, \bar{p} + \epsilon). \quad (22)$$

The rest of the procedure remains unchanged. The cost of the added flexibility is a tighter feasible region and higher MSE of the final score.

3.3.3 Satisfying the criteria for more than two protected groups

The algorithm can be modified to satisfy the fairness criteria for multiple groups, across multiple identifiers. First define each group as a unique combination of protected features. Then, the feasible set of error rates is given by the intersection of each $S(A)$ with the points satisfying the inequalities (11) where H is the highest-mean group and L is the lowest-mean group. Stage 1 of the algorithm proceeds to find the optimal set of error rates in that feasible region. Stage 2 proceeds as usual, implementing a separate program for each group.

4 Empirical Results

Let us take our procedure to data. In the first application, we post-process real COMPAS scores to demonstrate that risk assessments can be designed to output both calibrated risk scores as well as binary risk summaries satisfying equal error rates. Afterwards, we design a risk score to aid a lender's classification task to authorize loans, showing that it outperforms a common alternative strategy based on the omission of sensitive features. For interested readers, extensive detail about each application is presented in our online appendix [23].

4.1 Predicting criminal recidivism

Our procedure can design risk assessments that output both calibrated scores as well as binary “high” or “low” risk summaries satisfying equal error rates. We illustrate this in our first application, where we modify real criminal justice risk scores from COMPAS. As noted earlier, a *ProPublica* investigation showed that current COMPAS scores yield error imbalances across race, although they satisfy predictive parity overall [2, 1].

To check whether we can correct COMPAS error imbalances without sacrificing score calibration, we applied our post-processing technique to Broward County risk scores made public by *ProPublica* [16]. We define the outcome of interest as recidivism within two years, and we convert existing COMPAS scores that range from $[1, 10]$ to probabilities in $[0, 1]$. We define the classification cutoff as the minimum score of defendants classified as “high risk” in COMPAS, according to *ProPublica*’s influential analysis [17]. This corresponds to a cutoff of $\bar{p} = 0.54$ and loss parameter $k = 1.17$.

We compute the feasible region of achievable error rates according to Stage 1 of our algorithm and identify the loss-minimizing pair, as depicted in Figure 3a. Then, we use Stage 2 to post-process the COMPAS scores to achieve new calibrated scores yielding that optimal pair of error rates. The calibration of our scores is depicted in Figure 3b, where we group together by race defendants with the same post-processed scores and show that their corresponding recidivism outcomes lie on the main diagonal. Overall, our procedure eliminates the reported error disparities across racial groups (Figure 3a) while also preserving calibration (Figure 3b).

4.2 Predicting loan repayment

We next present an example of designing a risk score to inform a credit lender’s approvals of loan applicants. Our goal is to deliver to the lender calibrated scores for applications from two groups—one highly educated (H) and another less educated (L)—while ensuring that they yield classifications with equal group TPRs at the lender cutoff. That way, we know that qualified applicants will have the same probability of receiving a loan regardless of their education level. We suppose the lender in question views defaulting as highly costly and only authorizes loans to individuals with calibrated scores greater than $\approx .9$, corresponding to loss parameter $k = 10$.

We simulate this scenario by applying our algorithm to the Survey of Income and Program Participation (SIPP), a nationally-representative survey of the civilian population spanning multiple years [25]. We select as our outcome the ability to pay rent, mortgage, and utilities in 2016, and predict that outcome using survey responses from two years prior. We label individuals with at most a secondary school education as L and those with higher education as H .

The full dataset contains over 1,800 features spanning detailed financial variables (including work history, assets, and debts), as well as sensitive features (including demographic information). We apply our algorithm to the full feature set and derive calibrated scores that yield equal TPRs at the lender’s cutoff, using our algorithm extension that allows FPRs to vary. Then, we compare its performance to two accuracy-maximizing procedures: one based on the full feature set, and another commonly-practiced approach based on the omission of sensitive features. The results are summarized numerically in Table 1 and graphically in Figure 3c. Compared to prediction on all features and no post-processing, our algorithm raises the TPR of L and lowers that of H , while raising lender loss. Meanwhile, compared to the commonly used data omission strategy, our algorithm raises the probabilities that creditworthy applicants from *both* education groups are granted loans, and lowers loss for the lender.

■ **Table 1** Application to credit lending. Row [1] is based on raw scores. Row [2] summarizes the classifier that minimizes lender loss subject to equal true positive rates, given by the equal opportunity algorithm in Hardt et al. (2016). Row [3] summarizes our algorithm, which produces a calibrated score corresponding to equal true positive rate classifications; since it retrieves the same error rates as row [2], we see there is no added loss from enforcing score calibration. Row [4] summarizes the scores from the alternative procedure that omits sensitive features, displaying greater loss for the lender, lower true positive rates for both groups, and substantial error disparities across groups.

	Algorithmic Target	Lender Loss	TPR (H/L)	FPR (H/L)	Score MSE
<i>Trained on all features</i>					
[1]	Accuracy Maximizing	.517	(.795/.661)	(.341/.255)	.072
[2]	Eq. TPR Only	.532	(.727/.727)	(.299/.339)	N/A
[3]	*Eq. TPR + Calibration*	.532	(.727/.727)	(.299/.339)	.073
<i>Trained on limited features</i>					
[4]	Accuracy Maximizing	.591	(.603/.518)	(.202/.230)	.077

5 Conclusion

Decision-makers stand to benefit from algorithmic predictions. This paper studies fair prediction in the widespread setting in which a risk score is constructed to aid their classification tasks. We prove that it is possible to construct calibrated scores that lead to equal error rate classifications at group-blind cutoffs. We characterize exactly when it is possible and propose an algorithm that produces the most accurate score satisfying the fairness criteria and minimizing the decision-maker’s errors. Compared to a commonly practiced strategy of omitting sensitive data, we show that our algorithm can produce scores that enhance both efficiency and equity.

References

- 1 Julia Angwin and Jeff Larson. Bias in Criminal Risk Scores Is Mathematically Inevitable, Researchers Say. *ProPublica*, 2016.
- 2 Julia Angwin and Jeff Larson. Machine Bias. *ProPublica*, 2016.
- 3 Solon Barocas and Andrew D. Selbst. Big data’s disparate impact. *California Law Review*, 104(3):671–732, 2016.
- 4 Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, page 0049124118782533, 2018.
- 5 Alexandra Chouldechova. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data*, 5(2):153–163, 2017. doi:10.1089/big.2016.0047.
- 6 Sam Corbett-Davies and Sharad Goel. The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning. *arXiv:1808.00023 [cs]*, 2018. arXiv:1808.00023.
- 7 Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’17, pages 797–806, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3097983.3098095.
- 8 Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS ’12, pages 214–226, New York, NY, USA, 2012. Association for Computing Machinery. doi:10.1145/2090236.2090255.

- 9 Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Max Leiserson. Decoupled classifiers for group-fair and efficient machine learning. In Sorelle A. Friedler and Christo Wilson, editors, *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pages 119–133, New York, NY, USA, February 2018. PMLR.
- 10 Avi Feller, Emma Pierson, Sam Corbett-Davies, and Sharad Goel. A computer program used for bail and sentencing decisions was labeled biased against blacks. it’s actually not that clear. *The Washington Post*, 17, 2016.
- 11 Sumegha Garg, Michael P. Kim, and Omer Reingold. Tracking and improving information in the service of fairness. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC ’19, page 809–824, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3328526.3329624.
- 12 Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. Equality of Opportunity in Supervised Learning. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3315–3323. Curran Associates, Inc., 2016.
- 13 Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. Algorithmic Fairness. *AEA Papers and Proceedings*, 108:22–27, 2018. doi:10.1257/pandp.20181018.
- 14 Jon Kleinberg and Sendhil Mullainathan. Simplicity creates inequity: Implications for fairness, stereotypes, and interpretability. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC ’19, pages 807–808, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3328526.3329621.
- 15 Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPICs*, pages 43:1–43:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.ITCS.2017.43.
- 16 Jeff Larson. Data and analysis for “Machine bias”. *GitHub*, June 2017.
- 17 Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. How We Analyzed the COMPAS Recidivism Algorithm. <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>.
- 18 Marc Mezard and Andrea Montanari. *Information, Physics, and Computation*. Oxford University Press, Inc., USA, 2009.
- 19 Shira Mitchell, Eric Potash, Solon Barocas, Alexander D’Amour, and Kristian Lum. Algorithmic fairness: Choices, assumptions, and definitions. *Annual Review of Statistics and Its Application*, 8(1):141–163, 2021. doi:10.1146/annurev-statistics-042720-125902.
- 20 Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7599–7609. PMLR, 13–18 July 2020. URL: <http://proceedings.mlr.press/v119/perdomo20a.html>.
- 21 Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- 22 Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On Fairness and Calibration. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5680–5689. Curran Associates, Inc., 2017.
- 23 Claire Lazar Reich and Suhas Vijaykumar. A possibility in algorithmic fairness: Calibrated scores for fair classifications, 2020. arXiv:2002.07676.
- 24 Yonadav Shavit, Benjamin Edelman, and Brian Axelrod. Causal strategic linear regression. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 8676–8686. PMLR, 13–18 July 2020. URL: <http://proceedings.mlr.press/v119/shavit20a.html>.

- 25 U.S. Census Bureau. Survey of income and program participation, 2014. URL: <https://www.census.gov/programs-surveys/sipp/data/datasets.html>.
- 26 Cédric Villani. *Optimal Transport: Old and New*, volume 338. Springer Science & Business Media, 2008.
- 27 Marten Wegkamp. Model selection in nonparametric regression. *The Annals of Statistics*, 31(1):252–273, 2003. doi:10.1214/aos/1046294464.
- 28 Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, WWW '17, pages 1171–1180, Republic and Canton of Geneva, CHE, 2017. International World Wide Web Conferences Steering Committee. doi:10.1145/3038912.3052660.

A

 Appendix

A.1 Addendum to Theorem 1

Addendum. We relax condition (iii) of Theorem 1 and replace it with the weaker condition that $\text{Var}(Y|Z) > \epsilon$ almost surely. This will correspond to the assumption that Y cannot be perfectly predicted from any realization of Z .

We will make use of the criterion that Borel random variables R and R' are independent conditional on Σ iff for all bounded, continuous f and g we have

$$\mathbb{E}[f(R)g(R')|\Sigma] = \mathbb{E}[f(R)|\Sigma]\mathbb{E}[g(R')|\Sigma].$$

Now suppose that (Z, A, Y) are known to satisfy Theorem 1 conditions (i) and (ii), and that $\text{Var}(Y|Z) > 0$. Then let η be a $\text{Ber}(\epsilon)$ random variable independent of (Z, A, Y) . We consider a variable A_η that takes value A with probability $1 - \epsilon$ and otherwise flips the variable A with probability ϵ , that is,

$$A_\eta = A + \eta \pmod{2}.$$

This gives us a triple (Z, A_η, Y) that satisfies $\mathbb{E}[A|Z] \in (0, 1)$ and $\mathbb{E}[Y|A, Z] \in (0, 1)$ almost surely by construction, corresponding to condition (iii) from the Theorem. We can also show that the triple satisfies the other two conditions. For instance, to show that condition (i) holds, let S be an arbitrary set such that $S \in \sigma(Z)$.

We will use the fact that any $\sigma(Z)$ -measurable random variable V and any random variable U satisfy $\mathbb{E}[\mathbb{E}[U|Z]V] = \mathbb{E}[UV]$. In particular,

$$\begin{aligned} \mathbb{E}[\mathbb{E}[f(A_\eta)g(Y)|Z]\mathbb{1}_{Z \in S}] &= \mathbb{E}[f(A_\eta)g(Y)\mathbb{1}_{Z \in S}], (S \in \sigma(Z)) \\ &= \mathbb{E}_{(A, Z, Y)}[\mathbb{E}_\eta[f(A_\eta)]g(Y)\mathbb{1}_{Z \in S}], (\eta \perp (A, Y, Z)) \\ &= \mathbb{E}_{(A, Z, Y)}[\mathbb{E}[\mathbb{E}_\eta[f(A_\eta)]g(Y)|Z]\mathbb{1}_{Z \in S}], (S \in \sigma(Z)) \\ &= \mathbb{E}_{(A, Z, Y)}[\mathbb{E}[\mathbb{E}_\eta[f(A_\eta)]|Z]\mathbb{E}[g(Y)|Z]\mathbb{1}_{Z \in S}], (Y \perp A | Z) \\ &= \mathbb{E}_{(A, Z, Y)}[\mathbb{E}_\eta[f(A_\eta)]\mathbb{E}[g(Y)|Z]\mathbb{1}_{Z \in S}] \\ &= \mathbb{E}[f(A_\eta)\mathbb{E}[g(Y)|Z]\mathbb{1}_{Z \in S}], (\eta \perp (A, Y, Z)) \\ &= \mathbb{E}[\mathbb{E}[f(A_\eta)|Z]\mathbb{E}[g(Y)|Z]\mathbb{1}_{Z \in S}], \end{aligned}$$

where the last step follows since $\mathbb{E}[g(Y)|Z]\mathbb{1}_{Z \in S}$ is Z -measurable.

Because S was arbitrary and both $\mathbb{E}[f(A_\eta)|Z]\mathbb{E}[g(Y)|Z]$ and $\mathbb{E}[f(A_\eta)g(Y)|Z]$ are $\sigma(Z)$ -measurable, we can conclude that $\mathbb{E}[f(A_\eta)|Z]\mathbb{E}[g(Y)|Z] = \mathbb{E}[f(A_\eta)g(Y)|Z]$ almost surely so (i) is satisfied. A very similar argument shows that (ii) holds. Therefore, by Theorem 1, A_η is independent of (Z, Y) . Then given arbitrary bounded and continuous functions f and g ,

$$\mathbb{E}[f(A_\eta)g(Z, Y)] = \mathbb{E}[f(A_\eta)]\mathbb{E}[g(Z, Y)].$$

Using the fact that $A_\eta \rightarrow A$ as $\eta \downarrow 0$ in $L^2(\mathbb{P})$, and that $h \mapsto \mathbb{E}[h]$ and $(h, h') \mapsto \mathbb{E}[hh']$ are continuous in $L^2(\mathbb{P})$, we conclude by continuity that

$$\mathbb{E}[f(A)g(Z, Y)] = \mathbb{E}[f(A)]\mathbb{E}[g(Z, Y)].$$

Since f and g were arbitrary, we have in fact shown that A is independent of (Z, Y) , as wanted.

Thus, we have succeeded in proving the following refinement: under Theorem 1 assumptions (i) and (ii), if Y cannot be perfectly predicted from any realization of Z , then the random variables A and (Y, Z) must be independent.

Since assumptions (i) and (ii) continue to hold if we condition on $Z \in S$ for any S , we can say further that if Theorem 1 conditions (i) and (ii) hold and P is the set of values of Z from which perfect prediction is not possible, i.e. $\text{Var}(Y|Z) > 0$ then A and Y are independent conditionally on $Z \in P$. ◀

A.2 Proof of Proposition 2

Proof. First we can prove a lemma stating that $S(A)$ is convex. To see this, let ξ be an independent $\text{Ber}(\lambda)$ random variable. Then, by iterating expectations, one sees that

$$\alpha(\hat{y} + \xi(\hat{z} - \hat{y}), A) = \lambda\alpha(\hat{z}, A) + (1 - \lambda)\alpha(\hat{y}, A).$$

Using this convexity, we can prove the proposition. Note that the points $\alpha(\mathbb{1}\{p^* \geq c\}, A)$ that make up the group- A ROC curve of p^* describe the error rates achieved by all cutoff classifiers based on p^* , and so they are in $S(A)$. Meanwhile, since

$$\alpha(1 - \hat{y}, A) = (1, 1) - \alpha(\hat{y}, A),$$

the points $(1, 1) - \alpha(\mathbb{1}\{p^* \geq c\}, A)$ must also be in $S(A)$. This corresponds to the group- A ROC curve of the scores $1 - p^*$. Any point in the convex hull of these two ROC curves can be achieved by randomization as in the aforementioned lemma. For further details and intuition, see Section 4 in Hardt et al. [12]. Note that Hardt et al. choose not to illustrate the feasible region below the main diagonal as it corresponds to classifiers that are worse than random.

To show that *all* attainable error rates belong to this set, we use the convexity of $S(A)$ to note that the support points of $S(A)$ correspond to all classifiers that yield extrema of $\gamma_1\alpha_1(\hat{y}, A) + \gamma_2\alpha_2(\hat{y}, A)$ where (γ_1, γ_2) are arbitrary weights. To describe these support points tractably, we can use the result derived later in the appendix (Proposition 7) that shows that optimal classifications can be chosen to depend on only p^* and A , where $p^* = \mathbb{E}[Y|X, A]$. Thus the extrema of $\gamma \cdot \alpha(\hat{y}, A)$ are achieved by cutoff rules $f(p^*, A) = \mathbb{1}\{p^* \geq c\}$ and $f(p^*, A) = \mathbb{1}\{p^* < c\}$, giving support points

$$\bigcup_{c \in [0, 1]} \left\{ \alpha(\mathbb{1}\{p^* \geq c\}, A), (1, 1) - \alpha(\mathbb{1}\{p^* \geq c\}, A) \right\},$$

which as we have shown are contained in $S(A)$. Finally, we use the fact that a convex set containing all of its support points is equal to the convex hull of its support points. ◀

A.3 Proof of Extension of Proposition 3

Proof. Suppose that i holds and call \hat{p}_f the fair score for which $\hat{y} = \mathbb{1}\{\hat{p}_f \geq \bar{p}\}$ satisfies equal error rates. Then since \hat{p}_f is calibrated,

$$\begin{aligned}\mathbb{P}(Y = 1|\hat{y} = 1, A) &= \mathbb{E}[Y|\hat{p}_f \geq \bar{p}, A] = \mathbb{E}[\hat{p}_f|\hat{p}_f \geq \bar{p}, A] \geq \bar{p}, \\ \mathbb{P}(Y = 1|\hat{y} = 0, A) &< \bar{p}.\end{aligned}$$

So in addition to satisfying equal error rates, \hat{y} satisfies (9) and (10), which are equivalent to the two conditions in (8). Thus ii is a necessary condition for fairness.

Now we show the converse; ii is also sufficient for fairness. Suppose that ii holds and let \hat{y}_f be a classifier satisfying equal error rates and (8). Choose $\hat{p}(\hat{y}_f, A) = \mathbb{P}(Y = 1|\hat{y}_f, A)$. These scores are calibrated by construction. Also, since they satisfy $\hat{p}(\hat{y}_f = 0, A) < \bar{p}$ and $\hat{p}(\hat{y}_f = 1, A) \geq \bar{p}$, they exactly implement the classifier \hat{y}_f at the cutoff \bar{p} . ◀

A.4 Proof of Theorem 4

Proof. Building on the above extension of Proposition 3, it is enough for us to show that the existence of the point $(\alpha_1, \alpha_2) \in S(L) \cap S(H)$ satisfying (11) is equivalent to the following: There exists a classifier \hat{y} satisfying equal error rates and (8).

First note that $S(L) \cap S(H)$ is nonempty, since for example $(0, 0)$ and $(1, 1)$ are points in both $S(L)$ and $S(H)$. So we can consider some arbitrary (α_1, α_2) that is in $S(L) \cap S(H)$ and is therefore implementable by an equal error rate classifier that we call \hat{y}_e . We need to show that \hat{y}_e satisfying the conditions in (8) $\forall A$ is equivalent to its corresponding true and false positive rates $(\alpha_1(\hat{y}_e, A), \alpha_2(\hat{y}_e, A))$ satisfying (11) $\forall A$. Recall that the PPV condition in (8) required

$$\mathbb{P}(Y = 1|\hat{y}_e = 1, A) \geq \bar{p}.$$

Applying Bayes' rule to the inequality, we have

$$\begin{aligned}\mathbb{P}(Y = 1|\hat{y}_e = 1, A) &= \frac{\mathbb{P}(\hat{y}_e = 1|Y = 1, A)\mathbb{P}(Y = 1|A)}{\mathbb{P}(\hat{y}_e = 1|A)} \\ &= \frac{\alpha_2(\hat{y}_e, A)\mu_A}{\alpha_2(\hat{y}_e, A)\mu_A + \alpha_1(\hat{y}_e, A)(1 - \mu_A)} \geq \bar{p}.\end{aligned}$$

After algebraic manipulation, the restriction can be written

$$\frac{\alpha_2(\hat{y}_e, A)}{\alpha_1(\hat{y}_e, A)} \geq \frac{\bar{p}(1 - \mu_A)}{(1 - \bar{p})\mu_A} = \frac{\bar{p}}{(1 - \bar{p})\beta_A}.$$

where $\beta_A \equiv \mu_A/(1 - \mu_A)$. Therefore $(\alpha_1(\hat{y}_e, A), \alpha_2(\hat{y}_e, A))$ must satisfy the following

$$\frac{\alpha_2(\hat{y}_e, A)}{\alpha_1(\hat{y}_e, A)} \geq \frac{\bar{p}}{(1 - \bar{p})\beta_A}$$

Since $\beta_L < \beta_H$, the condition is more restrictive when $A = L$, giving the first condition in (11). We next similarly transform the NPV condition in (8), recalling it requires $\mathbb{P}(Y = 0|\hat{y} = 0, A) > 1 - \bar{p}$. By Bayes' rule,

$$\begin{aligned}\mathbb{P}(Y = 0|\hat{y} = 0, A) &= \frac{\mathbb{P}(\hat{y} = 0|Y = 0, A)\mathbb{P}(Y = 0|A)}{\mathbb{P}(\hat{y} = 0|A)} \\ &= \frac{(1 - \alpha_1(\hat{y}, A))(1 - \mu_A)}{(1 - \alpha_1(\hat{y}, A))(1 - \mu_A) + (1 - \alpha_2(\hat{y}, A))\mu_A} > 1 - \bar{p}.\end{aligned}$$

After algebraic manipulation, this becomes $\forall A$

$$\frac{(1 - \alpha_1(\hat{y}, A))}{(1 - \alpha_2(\hat{y}, A))} > \frac{(1 - \bar{p})\beta_A}{\bar{p}}.$$

Since $\beta_H > \beta_L$, the most restrictive case is when $A = H$, giving the second condition in (11).

Note that special attention should be given to the corner solutions. At point $(0, 0)$, the first condition in (11) becomes irrelevant and so the second condition in (11) is necessary and sufficient. Meanwhile at $(1, 1)$, the second condition in (11) becomes irrelevant so the first condition in (11) is necessary and sufficient. ◀

A.5 Proof of Corollary 5

Proof. Let F and G denote the lines for which the inequalities (11) hold with equality. That is to say, $F, G \subset \mathbb{R}^2$ are given by

$$F = \left\{ (\alpha_1, \alpha_2) \in \mathbb{R}^2 \mid \frac{\alpha_2}{\alpha_1} = \frac{\bar{p}}{\beta_L(1 - \bar{p})} \right\}, \quad G = \left\{ (\alpha_1, \alpha_2) \in \mathbb{R}^2 \mid \frac{(1 - \alpha_1)}{(1 - \alpha_2)} = \frac{\beta_H(1 - \bar{p})}{\bar{p}} \right\}$$

The lines intersect at $(\check{\alpha}_1, \check{\alpha}_2)$ given by (12). Our proof will rest on a few basic facts: $S(L) \cap S(H)$ is convex, F contains $(0, 0)$, G contains $(1, 1)$, and both lines have positive slope. First we prove that if $\check{\alpha}_1 \leq 0$, $\check{\alpha}_1 \geq 1$, or both ROC curves lie above the intersection $(\check{\alpha}_1, \check{\alpha}_2)$, then there exists a point (α_1, α_2) satisfying the feasibility conditions in Theorem 4.

Case I: $0 < \check{\alpha}_1 < 1$ and $(\check{\alpha}_1, \check{\alpha}_2)$ lies below both ROC curves. Note that increasing α_2 slackens both inequalities (11). Thus, if $0 < \check{\alpha}_1 < 1$ and $(\check{\alpha}_1, \check{\alpha}_2)$ lies below both ROC curves, there then exists a point $(\check{\alpha}_1, \alpha_2)$ with $\alpha_2 > \check{\alpha}_2$ that lies on the minimum of the two ROC curves, hence in $S(H) \cap S(L)$, and moreover the inequalities (11) hold at $(\check{\alpha}_1, \alpha_2)$. This is a feasible point.

Case II: $\check{\alpha}_1 \leq 0$. On the other hand, if $\check{\alpha}_1 \leq 0$, then in $(0, 1) \times \mathbb{R}$ the line F lies strictly above G . Then the point $(0, 0) \in S(L) \cap S(H) \cap F$ lies above G , meaning that the second condition in (11) holds and the point is feasible.

Case III: $\check{\alpha}_1 \geq 1$. If $\check{\alpha}_1 \geq 1$, then in $(0, 1) \times \mathbb{R}$ the line G lies strictly above F . Then the point $(1, 1) \in S(L) \cap S(H) \cap G$ lies above F , so the first condition in (11) holds and the point is feasible.

Finally, we prove the converse that if $0 < \check{\alpha}_1 < 1$ and $\check{\alpha}_2$ lies above at least one of the ROC curves, then the feasible region is empty. Let the intersection of $S(L) \cap S(H)$ with the half-space above F be denoted by I_F , and the intersection of $S(L) \cap S(H)$ with the half-space above G be denoted by I_G . We need to show that $I_F \cap I_G$ is empty. The argument follows from the convexity of $S(L) \cap S(H)$ and the fact that both F and G have positive slopes. In particular, due to the convexity of $S(L) \cap S(H)$, the positive slope of F , and the fact that $(0, 0)$ is in F , we know the line F must intersect the boundary of $S(L) \cap S(H)$ strictly to the left of $\check{\alpha}_1$. Meanwhile, G must intersect the boundary of $S(L) \cap S(H)$ strictly to the right of $\check{\alpha}_1$. Thus the rightmost point of I_F lies strictly to the left of the leftmost point of I_G , and the intersection of $S(L) \cap S(H)$ with both half-spaces above F and G must be empty. ◀

A.6 Justification for post-processing p^* in algorithm

First we justify post-processing the Bayes optimal p^* to arrive at the optimal fair \hat{p} . To do so we adapt Proposition 5.2 from Hardt et al. [12] to our setting and prove the following

► **Proposition 7.** *For any source distribution over (Y, X, A) with Bayes optimal regressor given by $p^*(X, A) = \mathbb{E}[Y|X, A]$ and loss function ℓ , there exists a predictor $\hat{p}(p^*, A)$ such that*

- (i) \hat{p} is an optimal predictor satisfying our fairness properties of calibration and equal error rates. That is, $\mathbb{E}[\ell(\mathbb{1}_{\hat{p}>p}, Y)] \leq \mathbb{E}[\ell(\mathbb{1}_{\hat{g}>p}, Y)]$ for any \hat{g} that satisfies the properties.
- (ii) \hat{p} is derived from (p^*, A) . In particular, it is a (possibly random) function of the random variables (p^*, A) alone, and is independent of X conditional on (p^*, A) .

Proof. To start, first note that our fairness properties of calibration and equal error rates on a score p and classifications $\mathbb{1}\{p \geq \bar{p}\}$ are “oblivious.” That is, they depend only on the joint distribution of (Y, A, p) given the known cutoff \bar{p} . We will show that for any arbitrary \hat{g} that satisfies the fairness properties, we can construct a \hat{p} that also satisfies fairness, yields the same expected loss, and is derived from (p^*, A) .

Consider an arbitrary $\hat{g} = f(X, A)$ satisfying the fairness properties. We can define $\hat{p}(p^*, A)$ as follows: draw a vector X' independently from the conditional distribution of X given the realized values of p^* and A , and set $\hat{p} = f(X', A)$. Note this \hat{p} satisfies (ii) by construction.

To show that this \hat{p} satisfies the fairness properties and yields the same expected loss as \hat{g} , note that since Y is binary with conditional expectation equal to the Bayes optimal p^* , we know Y is independent of X conditional on p^* . Therefore (Y, p^*, X, A) and (Y, p^*, X', A) have the same joint distribution, and so must $(f(X, A), A, Y)$ and $(f(X', A), A, Y)$. Since the fairness properties are oblivious and depend only on these latter joint distributions, then we know that as long as \hat{g} satisfies them then so will \hat{p} . Finally, we can deduce that (Y, \hat{g}) and (Y, \hat{p}) also have the same joint distribution, meaning that (i) is satisfied with equality. ◀

A.7 Our algorithm as a mean-preserving contraction of scores

We observe that a calibrated score derived from another is a mean-preserving contraction. Since the Bayes optimal p^* that serves as input to our algorithm frequently satisfies calibration (see Liu et al. 2019), then our post-processing method can be viewed as finding its smallest mean preserving contraction that achieves equal error rates at the decision-maker’s cutoff.

The relationship between calibrated scores related by post-processing is characterized by our proposition below.

► **Proposition 8.** *Let p_A be any calibrated score of group A , i.e. satisfying $\mathbb{E}[Y|p_A] = p_A$ for members of A , and let $\hat{p}_A = f(p_A, \zeta)$ be a score post-processed from p_A that is also calibrated, where ζ is independent of Y conditional on p_A . Then, \hat{p}_A is a mean-preserving contraction of p_A , with $p_A = \hat{p}_A + Z$ and $\mathbb{E}[Z|\hat{p}_A] = 0$. Conversely, any \tilde{p}_A that satisfies $p_A = \tilde{p}_A + Z$ with $\mathbb{E}[Z|\tilde{p}_A] = 0$ is calibrated.*

Proof. We first show that \hat{p}_A is a mean-preserving contraction of p_A . To start, note that the post-processed \hat{p}_A is assumed to be calibrated, so $\mathbb{E}[Y|\hat{p}_A] = \hat{p}_A$. Moreover, since $\hat{p}_A = f(p_A, \zeta)$, we have $\sigma(\hat{p}_A) \subseteq \sigma(p_A, \zeta)$. Therefore by the tower property of conditional expectation,

$$\begin{aligned} \hat{p}_A &= \mathbb{E}[Y|\hat{p}_A] = \mathbb{E}[\mathbb{E}[Y|p_A, \zeta]|\hat{p}_A] \\ &= \mathbb{E}[\mathbb{E}[Y|p_A]|\hat{p}_A], \text{ (by conditional independence of } \zeta) \\ &= \mathbb{E}[p_A|\hat{p}_A], \text{ (by calibration of } p_A). \end{aligned}$$

Then $p_A = p_A + (\hat{p}_A - \mathbb{E}[p_A|\hat{p}_A]) = \hat{p}_A + (p_A - \mathbb{E}[p_A|\hat{p}_A])$ where the second term is by construction mean independent of \hat{p}_A , so \hat{p}_A is a mean-preserving contraction of p_A .

Now we show that if the score \tilde{p}_A is a mean-preserving contraction of p_A such that $p_A = \tilde{p}_A + Z$ for some Z satisfying $\mathbb{E}(Z|\tilde{p}_A) = 0$, then \tilde{p}_A is calibrated. Observe that

$$\mathbb{E}[p_A|\tilde{p}_A] = \mathbb{E}[\tilde{p}_A + Z|\tilde{p}_A] = \mathbb{E}[\tilde{p}_A|\tilde{p}_A] + \mathbb{E}[Z|\tilde{p}_A] = \tilde{p}_A$$

which is sufficient to show that \tilde{P}_A is calibrated. To see why, recall that p_A is calibrated and note that by the tower property of conditional expectation with $\sigma(\tilde{p}_A) \subseteq \sigma(p_A)$,

$$\mathbb{E}[p_A|\tilde{p}_A] = \mathbb{E}[\mathbb{E}(Y|p_A)|\tilde{p}_A] = \mathbb{E}[Y|\tilde{p}_A]. \quad \blacktriangleleft$$

A.8 Justification for discretizing p^*

Our algorithm uses the discretization of p^* to construct a linear program that maps probability masses from p^* to \hat{p} . Note that even if the original p^* is not discrete, it can easily be discretized into N bins by taking $p' = \lfloor Np^* \rfloor / N$. The discretized score will satisfy $|p' - p^*| \leq N^{-1}$ almost surely, so for large values of N , the discretization p' approximates p^* well.