

2019

From soft classifiers to hard decisions

Ran Canetti, Aloni Cohen, Nishanth Dikkala, Govind Ramnarayan, Sarah Scheffler, Adam Smith. 2019. "From Soft Classifiers to Hard Decisions." Proceedings of the Conference on Fairness, Accountability, and Transparency - FAT* '19, <https://doi.org/10.1145/3287560.3287561>
<https://hdl.handle.net/2144/40972>

"Downloaded from OpenBU. Boston University's institutional repository."

From Soft Classifiers to Hard Decisions: How fair can we be?*

Ran Canetti[†] Aloni Cohen[‡] Nishanth Dikkala[§] Govind Ramnarayan[¶]
 Sarah Scheffler^{||} Adam Smith^{**}

January 23, 2019

Abstract

A popular methodology for building binary decision-making classifiers in the presence of imperfect information is to first construct a calibrated non-binary “scoring” classifier, and then to post-process this score to obtain a binary decision. We study various group fairness properties of this methodology, when the non-binary scores are calibrated over all protected groups, and with a variety of post-processing algorithms. Specifically, we show:

- There does not exist a general way to post-process a calibrated classifier to equalize protected groups’ positive or negative predictive value (PPV or NPV). For certain “nice” calibrated classifiers, either PPV or NPV can be equalized when the post-processor uses different thresholds across protected groups. Still, when the post-processing consists of a single global threshold across all groups, natural fairness properties, such as equalizing PPV in a nontrivial way, do not hold even for “nice” classifiers.
- When the post-processing stage is allowed to *defer* on some decisions (that is, to avoid making a decision by handing off some examples to a separate process), then for the non-deferred decisions, the resulting classifier can be made to equalize PPV, NPV, false positive rate (FPR) and false negative rate (FNR) across the protected groups. This suggests a way to partially evade the impossibility results of Chouldechova and Kleinberg et al., which preclude equalizing all of these measures simultaneously. We also present different deferring strategies and show how they affect the fairness properties of the overall system.

We then evaluate our post-processing techniques using the COMPAS data set from 2016.

*The order of authors is alphabetical and does not convey any information on the authors’ relative contributions.

[†]Boston University and Tel Aviv University. Member of CPIIS. Supported by NSF awards 1413920 & 1801564, ISF award 1523/14.

[‡]MIT. Supported by NSF award CNS-1413920.

[§]MIT. Supported by NSF award CCF-1617730, CCF-1650733, and ONR N00014-12-1-0999.

[¶]MIT. Supported by NSF award CCF-1665252 and NSF award DMS-1737944.

^{||}Boston University. Supported by the Clare Boothe Luce Graduate Research Fellowship and NSF award 1414119.

^{**}Computer Science Department, Boston University. Supported in part by NSF awards IIS-1447700 and AF-1763786 and a Sloan Foundation Research Award. ads22@bu.edu.

Contents

1	Introduction	3
1.1	Related work	6
2	Preliminaries	6
2.1	Calibration	7
2.2	Accuracy Profiles (APs)	8
2.3	Group Fairness Measures	9
3	The Limits of Post-Processing	10
3.1	Fairness Conditions for Post-Processors	10
3.2	General impossibility of equalizing PPV, NPV	12
3.3	A niceness Condition for APs	12
3.4	Equalizing PPV or NPV by Thresholding	13
3.4.1	Group Blind Thresholds	13
3.4.2	Group-Aware Thresholds	15
3.4.3	The Limitations of Thresholding	17
3.5	Equalizing APs	20
3.5.1	Mass Averaging	21
4	Post-Processing Calibrated Classifiers with Deferrals	22
4.1	Thresholding with deferrals	22
4.2	Equalizing APs with deferrals	26
5	Experiments on COMPAS Data	29
5.1	Thresholding with Deferrals	29
5.2	Converting one AP into Another	30
5.3	Equalizing APs	30
6	Models of Deferring	31
6.1	Technical implications of deferral model	32
A	Results for NPV and Continuous, Full Support APs	35
A.1	Results for Negative Predictive Value (NPV)	35
A.2	Results for Continuous, Full Support APs	36

1 Introduction

The concept of *fairness* is deeply ingrained in our psyche as a fundamental, essential ingredient of Human existence. Indeed the perception of fairness, broadly construed as accepting each others’ equal right for well being, is arguably one of the most basic tenets of cooperative societies of individuals in general, and Human existence in particular.

However, as fundamental as this concept may be, it is also extremely elusive: different societies have developed very different notions of fairness and equality among individuals, subject to different religious, ethical, and social beliefs; in particular, the intricate interplay between fairness and *justice*, which is yet another somewhat elusive concept, is often not well-defined and a matter of subjective interpretation.

The concept is further complicated by the fact that human decisions are often made with incomplete information and limited resources. These two factors must be taken into account when evaluating whether decision-making processes are “fair.” Indeed, these two aspects of the problem have become increasingly prominent as societies grow and decision processes become more complex and algorithmic.

One way that researchers are responding to these growing concerns is by attempting to formulate precise notions for *fairness* of decisions processes, e.g. [DHP⁺11, KMR16]. While these notions do not intend to capture the complexities of the ethical, socio-economic, or religious aspects of fairness, they do consider the fairness aspects of statistical decision-making processes with incomplete information. Essentially, these notions accept the fact that a decision process with incomplete information will inevitably make errors relative to the desired full-information notion (which is treated as a given), and provide guidelines on how to “distribute the errors fairly” across individuals, or alternatively across groups of individuals. These definitions have proven to be meaningful and eye opening; in particular, it has been demonstrated that some very natural notions of “fair distribution of errors” are mutually inconsistent: No decision mechanism with incomplete information can satisfy all, except for in trivial cases [Cho17, KMR16].

Faced with this basic impossibility, we would like to better understand the process of decision making with incomplete information, and use this understanding to propose ways to circumvent this impossibility.

Specifically, we concentrate on the task of post-processing a calibrated soft classifier under group fairness constraints. We suppose that individuals belong to one of two or more disjoint *protected groups*. Our overall task is to decide whether a given individual has some hidden binary property B in a way that ensures “fair balancing of errors” across the groups.

For that purpose, we consider the following two-stage mechanism. The first stage consists of constructing a classifier \hat{S} that outputs for each individual x a score $s \in [0, 1]$ that is related to the chance that x has property B . The only requirement we make of \hat{S} is group-wise calibration: for both g_1 and g_2 , and for each $s \in [0, 1]$, the fraction of individuals in the group that get score s and have the property, out of all individuals in the group that get score s , is s . The second stage takes as input the output $s = \hat{S}(x)$ of the first stage and the group to which x belongs, and outputs a binary decision: its best guess at whether x has property B .

An attractive aspect of this two-stage mechanism is that each stage can be viewed as aimed at a different goal: The first stage is aimed at gathering information and providing the best accuracy possible, with only minimal regard to fairness (i.e only group-wise calibration). The second stage is aimed to extract a decision from the information collected in the first stage, while making sure that the errors are distributed “fairly.”

To further focus our study, we take the first stage as a given and concentrate on the second. That is, we consider the problem of *post-processing* the scores given by the calibrated soft classifier \hat{S} into binary predictions. A representative example is a judge making a bail decision based on a score provided by a software package. Following [Cho17, HPS16] we consider the following four performance measures for the resulting binary classifier: the *positive predictive value (PPV)*, namely the fraction of individuals that have the property among all individuals that the classifier predicted to have the property; The *false positive rate (FPR)*, namely the fraction of individuals that were predicted to have the property among all individuals that don't have the property; The *negative predictive value (NPV)* and *false negative rate (FNR)*, which are defined analogously. Ideally, we would like to equalize each one of the four measures across the groups, i.e. the measure will have the same value when restricted to samples from each group. Unfortunately, however, we know that this is impossible in general [Cho17, KMR16]. This leads us to a broad question that motivates our work:

Under what conditions can we post-process a calibrated soft classifier's outputs so that the resulting hard classifier equates a subset of $\{\text{PPV}, \text{NPV}, \text{FNR}, \text{FPR}\}$ across a set of protected groups? How can we balance these conflicting goals?

Results: Post-Processing With Thresholds. In a first set of results we consider the properties obtained by post-processing via a “threshold” mechanism. Naively, a threshold post-processing mechanism would return 1 for individual x whenever the calibrated score $s(x)$ is above some fixed threshold, and return 0 otherwise. We somewhat extend this mechanism by allowing the post-processor “fine-tune” its decision by choosing the output probabilistically whenever the result of the soft classifier is exactly the threshold.

We first observe that the popular and natural post-processing method of using a single threshold across all groups has an inherent deficiency: No such mechanism can in general guarantee equality of either PPV or NPV across the protected groups.

We then show that, when using different thresholds for the different groups, one can equalize *either* PPV or NPV (but not both) across the two groups, assuming the profile of \hat{S} has some non-degeneracy property.

The combination of the impossibility of single threshold and the possibility of per-group threshold also stands in contrast to the belief that a soft classifier that is calibrated across both groups allows “ignoring” group-membership information in any post-processing decision [MP17]. Indeed, the conversion to a binary decision “loses information” in different ways for the two groups, and so group membership becomes relevant again after post-processing.

Results: Adding deferrals. For the second set of results we consider post-processing strategies that do not always output a decision. Rather, with some probability the output is \perp , or “I don't know”, which means that the decision is deferred to another (hopefully higher quality, even if more expensive) process. Let us first present our technical results and then discuss potential interpretations and context.

The first strategy is a natural extension of the per-group threshold: we use *two* thresholds per group, returning 1 above the right threshold, 0 below the left threshold, and \perp between the thresholds. We show that there always exists a way to choose the thresholds such that, conditioned on the decision not being \perp , both the PPV and NPV are equal across groups.

Next we show a family of post-processing strategies where, conditioned on the decision not being \perp , *all four quantities* (PPV, NPV, FPR, FNR) are equal across groups.

All strategies in this family have the following structure: Given an individual x , the strategy first makes a randomized decision whether to defer on x , where the probability depends on $\hat{S}(x)$ and the group membership of x . If not deferred, then the decision is made via another post-processing technique.

One method for determining the probabilities of deferrals is to make sure that, the distribution of scores returned by the calibrated soft classifier *conditioned on not deferring*, is equal for the two groups (That is, let $p_{s,g}$ denote the probability, restricted to group g , that an element gets score s conditioned on not deferring. Then for any s , we choose deferral probabilities so that $p_{s,g_1} = p_{s,g_2}$.) The resulting classifier can then be post-processed in *any* group-blind way (say, via a single threshold mechanism as described above).

Of course, the fact that all four quantities are equalized conditioned on not deferring does not, in and of itself, provide any guarantees regarding the fairness properties of the overall decision process — which includes also the downstream decision mechanism. For one, it would be naive to simply assume that fairness “composes” [DI18]. Furthermore, the impossibility of [Cho17, KMR16] says that the overall decision-making process cannot possibly equalize all four measures.

However, in some cases one can provide alternative (non-statistical) justification for the fairness of the overall process: For instance, if the downstream decision process never errs, the overall process might be considered “procedurally fair.” We present more detailed reflections on our deferral-based approach in Section 6.

We note that deferring was considered in machine learning in a number of contexts, including the context of fairness-preservation [MPZ17]. In these works, the classifier typically punts only when its confidence regarding some decision is low. By contrast, we use deferrals in order to “equalize” the probability mass functions of the soft classifier over the two groups, which may involve deferring on individuals for whom there is higher confidence. Indeed, deferring on some higher-confidence individuals seems inherent to our goal of equalizing PPV, NPV, FPR, and FNR while keeping the deferral rate low. Furthermore, our framework allows for a wide range of deferral strategies which might be used to promote additional goals. Pursuing alternate strategies for deferral is an interesting direction for future work.

Experimental results. We test our methodology on the Broward county dataset with COMPAS scores made public by ProPublica [ALMK16] in order to better understand its strengths and limitations. Indeed, it has been shown that the COMPAS scoring mechanism is an approximately calibrated soft classifier. We first ran our two-threshold post-processing mechanism and obtained a binary decision algorithm which equalizes both PPV and NPV across Caucasians and African-Americans.

We then ran our post-processing mechanism with deferrals to equalize all four of PPV, NPV, FPR, FNR across the two groups, with three different methods for deciding how to defer: In the first method, decisions are deferred only for Caucasians; in the second, decisions are deferred only for African Americans; in the third method, decisions are deferred for an equal fraction of Caucasians and of African Americans. This fraction is precisely equal to the statistical (total variation) distance between the distributions of scores produced by the soft classifier on the two groups. More details about the results along with figures are given in Section 5.

Extensions and open problems. As just mentioned, a natural question is to find alternative ways for deciding when to defer, along with ways to argue fairness properties for the overall combined process.

We also leave open the setting where individuals belong to multiple, potentially intersecting groups as in [HKRR17, KNRW18].

Yet another question is to consider additional (or alternative) properties of soft classifiers that will make for more efficient or effective post-processing.

1.1 Related work

We briefly describe the works most closely related to ours, though both the list of works and their summaries are inevitably too short. Our work fits in a research program on group fairness notions following the work of Chouldechova [Cho17] and Kleinberg et al. [KMR16]. Those works demonstrate the inherent infeasibility of simultaneously equalizing a collection of measures of group accuracy. Our work considers the notions of calibration as formalized in [PRW⁺17] and those of PPV, NPV, FPR, and FNR from [Cho17] and [KMR16].

The power of post-processing calibrated scores into decisions using threshold classifiers in the context of fairness has been previously studied by Corbett-Davies, Pierson, Feller, Goel, and Huq [CPF⁺17]. As in our work, they show that it is feasible to equalize certain statistical fairness notions across groups using (possibly different) thresholds. They additionally show that these thresholds are in some sense optimal. Whereas [CPF⁺17] focuses on statistical parity, conditional statistical parity, and false positive rate, our most comparable results consider PPV. In our work, we further show that in some cases thresholds fail to equalize both PPV and NPV (called *predictive parity* by [Cho17]), unless we also allow our post-processor to defer on some inputs. Our work also studies methods of post-processing that are much more powerful than thresholding, especially when allowing deferrals. On the technical side, [CPF⁺17] assumes that their soft classifiers are supported on the continuous interval $[0, 1]$, simplifying the analyses. We instead study classifiers with finite support as it is closer to true practice in many settings (e.g., COMPAS risk scores).

Using deferrals to promote fairness has been considered also in the work of Madras, Pitassi, and Zemel [MPZ17]. Specifically they consider how deferring on some inputs may promote a combination of accuracy and fairness, especially when taking explicit account of the downstream decision maker. They make use of two-threshold deferring post-processors like those discussed in Section 5. While it helped inform our work, [MPZ17] takes a more experimental approach and focuses on minimizing the “disparate impact,” a measure of total difference in classification error between groups, while maximizing accuracy. One important difference between our works is that Madras et al. distinguish between “rejecting” and “deferring.” Rejecting is oblivious as to properties of the downstream decision maker, while deferring tries to counteract the biases of the decision maker. Our work considers only the former notion, but uses the term “defer” instead of “reject.”

2 Preliminaries

We study the problem of binary classification. An *instance* is an element, usually denoted x , of a universe \mathcal{X} . We restrict our attention to instances sampled uniformly at random from the universe, denoted $X \sim \mathcal{X}$. Our theory extends directly to any other distribution on \mathcal{X} ; that distribution does

not need to be known to the classifiers. Each instance x is associated with a *true type* $Y(x) \in \{0, 1\}$. Each instance x is also associated with a *group* $G(x) \in \mathcal{G}$, where \mathcal{G} is the set of groups. We restrict our attention to sets \mathcal{G} that form a partition of the universe \mathcal{X} . We denote by \mathcal{X}_g the set of instances x in group g , and by X_g the random variable distributed uniformly over \mathcal{X}_g . Note that for any events E_1 and E_2 , $\Pr_{X \sim \mathcal{X}_g}[E_1 \mid E_2] = \Pr_{X \sim \mathcal{X}}[E_1 \mid E_2, G(X) = g]$.

Definition 2.1 (Base rate (BR)). The *base rate* of a group $g \in \mathcal{G}$, is

$$\text{BR}_g = \Pr[Y(X_g) = 1] = \mathbb{E}[Y(X_g)]. \quad (1)$$

When \mathcal{X} is finite, BR_g is simply the fraction of individuals x in the group g for whom $Y(x) = 1$.

A *classifier* is a randomized function with domain $\mathcal{X} \times \mathcal{G}$.¹ A *hard classifier*, denoted \hat{Y} , outputs a *prediction* in $\{0, 1\}$, interpreted as a guess of the true type $Y(x)$. A *soft classifier*, denoted \hat{S} , outputs a *score* $s \in [0, 1]$, interpreted as a measure of confidence that $Y(x) = 1$. We restrict our attention to soft classifiers with finite image. We call a classifier *group blind* if its output is independent of the input group g . For all groups $g \in \mathcal{G}$, we call a hard classifier \hat{Y} *non-trivial on* g if $\Pr[\hat{Y}(X_g) = 1] > 0$ and $\Pr[\hat{Y}(X_g) = 0] > 0$. Hard classifiers are *trivial on* g if they are not non-trivial on g .

A *post-processor* is a randomized function with domain $[0, 1] \times \mathcal{G}$. As with classifiers, a post-processor can be *hard* or *soft*. A hard post-processor, denoted \hat{D} , outputs a prediction in $\{0, 1\}$. A soft post-processor, denoted \hat{D}^{soft} , outputs a score $s \in [0, 1]$. Observe that for a soft classifier \hat{S} , $\hat{D} \circ \hat{S}$ is a hard classifier, and $\hat{D}^{\text{soft}} \circ \hat{S}$ is a soft classifier. As with classifiers, we call a post-processor group blind if its output is independent of the group g , and we restrict our attention to post-processors with finite image. The restriction to finite image is for mathematical convenience (and also because digital memory leads to discrete universes); our results generalize to infinite images as well.

In Section 4, we expand the definitions of both classifier and post-processors to allow an additional input or output: the special symbol \perp .

2.1 Calibration

Several works concerning algorithmic fairness focus on various notions of *calibration*. The following calibration notions are defined only over soft classifiers:

Definition 2.2 (Calibration (Soft)). We say a soft classifier \hat{S} is *calibrated* if $\forall s \in [0, 1]$ for which $\Pr_{X \sim \mathcal{X}}[\hat{S}(X) = s] > 0$,

$$\Pr_{X \sim \mathcal{X}}[Y(X) = 1 \mid \hat{S}(X) = s] = \mathbb{E}_{X \sim \mathcal{X}}[Y(X) \mid \hat{S}(X) = s] = s.$$

The probability above is taken over the sampling of X , as well as random choices made by \hat{S} at classification time.

¹As the focus of this paper is on the post processing of classifiers, we set aside questions such as the origin of the given classifier, including the randomness used in training, the origin or quality of the training data, and societal factors affecting the classifier. In particular, the classifiers we consider in this work are memoryless: they do not remember inputs or random choices from previous invocations. That is, we assume that if X, X' are two independent random variables drawn from \mathcal{X} then $\hat{S}(X)$ and $\hat{S}(X')$ (respectively $\hat{Y}(X)$ and $\hat{Y}(X')$) are also independent random variables. (Our formalism can be naturally extended also to classifiers with initial randomized preprocessing, by considering the family of derivative classifiers, where for each derivative classifier the random choices made at preprocessing are fixed to some value. The formalism can then be applied separately to each derivative classifier.)

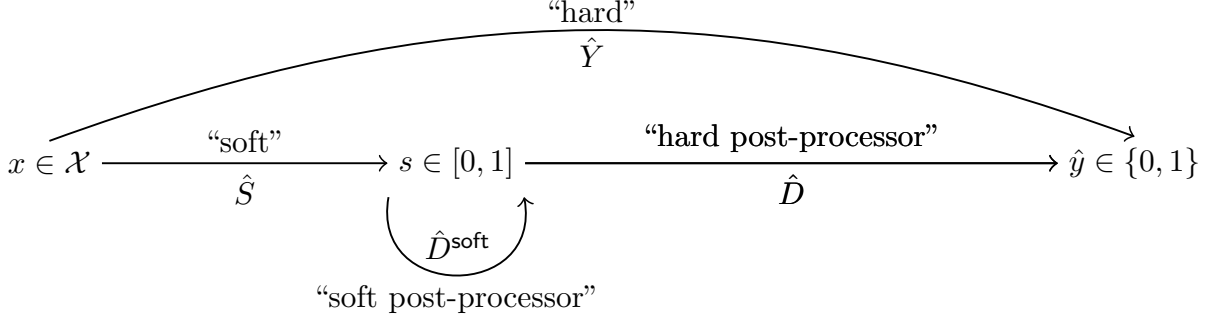


Figure 1: We call a classifier that returns results in $[0, 1]$ a soft classifier to differentiate it from those which return results in $\{0, 1\}$, which we call hard classifiers. We refer to classifiers that take as input the output of a soft classifier as post-processors.

Definition 2.3 (Groupwise Calibration (Soft)). We say that a soft classifier \hat{S} is *groupwise calibrated* if it is calibrated within all groups. That is, $\forall g \in \mathcal{G}$ and $\forall s \in [0, 1]$ for which $\Pr[\hat{S}(X_g) = s] > 0$, we have that

$$\Pr[Y(X_g) = 1 \mid \hat{S}(X_g) = s] = s.$$

Groupwise calibration is essentially the same notion as *multicalibration* [HKRR17] with the difference that in their case the true types are values in $[0, 1]$. We use a different term to emphasize that we restrict our attention to collections of groups \mathcal{G} that form a partition of the universe \mathcal{X} .

The two definitions above are stated for soft classifiers whose output distribution is discrete, since we must be able to condition on the event $\hat{S}(X) = s$ or $\hat{S}(X_g) = s$. That said, it extends naturally to classifiers with continuously-distributed outputs provided that the conditional probabilities are well defined.

2.2 Accuracy Profiles (APs)

Throughout this work, we make repeated reference to the probability mass function of the random variable $\hat{S}(X_g)$ for a calibrated soft classifier \hat{S} acting on a randomly distributed input X_g . We call this distribution on calibrated scores an *accuracy profile* (AP).

Definition 2.4 (Accuracy Profile (AP)). The *accuracy profile* (AP) of a calibrated soft classifier \hat{S} for a group g , denoted by $\hat{\mathcal{P}}_g$, is the PMF of $\hat{S}(X_g)$. That is, for $s \in [0, 1]$, $\hat{\mathcal{P}}_g(s) = \Pr[\hat{S}(X_g) = s]$.

Abusing notation, we denote by $\hat{\mathcal{P}}$ the collection $\{\hat{\mathcal{P}}_g\}_{g \in \mathcal{G}}$, and call it the *AP of \hat{S}* . We denote by $\text{Supp}(\hat{\mathcal{P}}_g)$ the support of the AP $\hat{\mathcal{P}}_g$, namely the set $\text{Supp}(\hat{\mathcal{P}}) = \{s : \exists x \in \mathcal{X}_g, \exists r \text{ s.t. } \hat{S}(x, r) = s\} \subseteq [0, 1]$.

An accuracy profile is a distribution of scores for a calibrated classifier \hat{S} . Because \hat{S} is calibrated, the AP conveys information about the performance of \hat{S} , and is constrained by properties of the underlying distribution on X . For example, the AP's expectation is exactly the base rate for the population:

Proposition 2.1. For any groupwise calibrated soft classifier \hat{S} , for all groups $g \in \mathcal{G}$: $BR_g = \mathbb{E}[\hat{S}(X_g)]$.

Proof of Proposition 2.1.

$$\begin{aligned}
\text{BR}_g &= \Pr[Y(X_g) = 1] \\
&= \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \Pr[Y(X_g) = 1 \mid \hat{S}(X_g) = s] \Pr[\hat{S}(X_g) = s] \\
&= \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} s \Pr[\hat{S}(X_g) = s] \\
&= \mathbb{E}[\hat{S}(X_g)]
\end{aligned}$$

where the third line follows from the definition of a calibrated classifier (Definition 2.3). \square

Accuracy profiles also provide useful geometric intuition for reasoning about the effects of post-processing calibrated scores. We elaborate on this in Section 3.1 (see Figure 2).

2.3 Group Fairness Measures

Several well-studied measures of statistical “fairness” (e.g., [HPS16, Cho17, KMR16, PRW⁺17, HKRR17, KNRW18]) look at how the following key performance measures of a classifier differ across groups. The *false positive rate* (FPR) of a hard classifier \hat{Y} for a group g is the rate at which \hat{Y} gives a positive classification among instances $x \in \mathcal{X}_g$ with true type 0. The *false negative rate* (FNR) is defined analogously for predicted negative instances with true type 1. *Positive predictive value* (PPV) and *negative predictive value* (NPV) track the rate of mistakes within instances that share a predicted type. Informally, positive predictive value captures how much meaning can be given to a predicted 1, and negative predictive value is similar for predicted 0. We now define these statistics formally.

Definition 2.5. Given a hard classifier \hat{Y} and a group g , we define
the *false positive rate* of \hat{Y} for g : $\text{FPR}_{\hat{Y},g} = \Pr[\hat{Y}(X_g) = 1 \mid Y(X_g) = 0]$;
the *false negative rate* of \hat{Y} for g : $\text{FNR}_{\hat{Y},g} = \Pr[\hat{Y}(X_g) = 0 \mid Y(X_g) = 1]$;
the *positive predictive value* of \hat{Y} for g : $\text{PPV}_{\hat{Y},g} = \Pr[Y(X_g) = 1 \mid \hat{Y}(X_g) = 1]$;
the *negative predictive value* of \hat{Y} for g : $\text{NPV}_{\hat{Y},g} = \Pr[Y(X_g) = 0 \mid \hat{Y}(X_g) = 0]$.

The probability statements in the definitions above reflect two sources of randomness: the sampling of X_g from the group g and any random choices made by the classifier \hat{Y} .

Among previous works, some [HPS16, KMR16] focus on equalizing only one or both of the false positive rates and false negative rates across groups, called *balance* for the negative and positive classes, respectively. Equalizing positive and negative predictive value across groups is often combined into one condition called *predictive parity* [Cho17]. We split the value out to be a separate condition for the positive and negative predictive classes. Predictive parity appears to be a hard-classifier analogue of calibration: both can be interpreted as saying that the output of the classifier (hard or soft) contains all the information contained in group membership. Our results highlight that the relationship between these notions is more subtle than it first appears; see Section 3 for further discussion.

3 The Limits of Post-Processing

Suppose throughout this section that \hat{S} is a groupwise calibrated soft classifier. Our goal in this section is to make binary predictions based on $\hat{S}(x)$ — and possibly the group $G(x)$ — subject to equalizing PPV and/or NPV among groups. That is, we wish to make a prediction using a hard post-processor \hat{D} such that $\hat{Y} = \hat{D} \circ \hat{S}$ equalizes PPV and/or NPV among groups. We chose to concentrate first on (the limitations of) equalizing PPV and NPV rather than FPR and FNR due to the conceptual similarity of PPV and NPV to calibration. Also, the case of equalizing false positive rates with thresholds is addressed in [CPF⁺17].

3.1 Fairness Conditions for Post-Processors

We begin by making a simple observation about post-processing that provides some geometric intuition for the rest of this section. Just as in Proposition 2.1, we can express $\text{PPV}_{\hat{Y},g}$ and $\text{NPV}_{\hat{Y},g}$ succinctly in terms of conditional expectations over the AP $\hat{\mathcal{P}}_g$.

Proposition 3.1. *Let $\hat{Y} = \hat{D} \circ \hat{S}$ be a hard classifier that is non-trivial for all $g \in \mathcal{G}$ where \hat{S} is groupwise calibrated with respect to \mathcal{G} . For any $g \in \mathcal{G}$ we have:*

$$\begin{aligned}\text{PPV}_{\hat{Y},g} &= \mathbb{E}[\hat{S}(X_g) \mid \hat{Y}(X_g) = 1] \\ \text{NPV}_{\hat{Y},g} &= 1 - \mathbb{E}[\hat{S}(X_g) \mid \hat{Y}(X_g) = 0]\end{aligned}$$

Proof of Proposition 3.1. We first observe that the output of a post-processor is conditionally independent of the true type, conditioned on the output of the soft classifier it is post-processing and the group membership:

Fact 3.1. *Consider any randomized function $\hat{D} : [0, 1] \times \mathcal{G} \rightarrow \{0, 1\}$. Since \hat{D} is a randomized function with inputs $s \in [0, 1]$ and $g \in \mathcal{G}$, we have that*

$$(\hat{D}(\hat{S}(X), G(X)) \perp Y(X) \mid (\hat{S}(X), G(X)) \quad (2)$$

or in other words that $\hat{D}(\hat{S}(X), G(X))$ is conditionally independent of the true type $Y(X)$, since fixing the inputs to \hat{D} makes its output purely a function of its random string.

Now recall that $\text{PPV}_{\hat{Y},g}$ and $\text{NPV}_{\hat{Y},g}$ are well-defined for all groups because \hat{Y} is non-trivial on all groups. We then have

$$\begin{aligned}\text{PPV}_{\hat{Y},g} &= \Pr[Y(X_g) = 1 \mid \hat{Y}(X_g) = 1] \\ &= \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \Pr[Y(X_g) = 1, \hat{S}(X_g) = s \mid \hat{D}(\hat{S}(X_g), g) = 1] \\ &= \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \Pr[Y(X_g) = 1 \mid \hat{S}(X_g) = s, \hat{D}(\hat{S}(X_g), g)] \\ &\quad \cdot \Pr[\hat{S}(X_g) = s \mid \hat{D}(\hat{S}(X_g), g) = 1] \\ &= \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} s \Pr[\hat{S}(X_g) = s \mid \hat{D}(\hat{S}(X_g), g) = 1] \\ &= \mathbb{E}[\hat{S}(X_g) \mid \hat{Y}(X_g) = 1]\end{aligned}$$




Figure 2: Accuracy Profiles (APs, definition 2.4) yield useful geometric intuitions, which come from the calibration property (definition 2.2). With a threshold, the expected PPV, NPV, FPR, and FNR can be seen visually.

where the fourth line follows from the fact that the group g is fixed within \mathcal{X}_g , which lets us apply Fact 3.1, and the fact that \hat{S} is calibrated on g . Similar simplifications give us that

$$\begin{aligned} \text{NPV}_{\hat{Y},g} &= \Pr[Y(X_g) = 0 \mid \hat{D}(\hat{S}(X_g), g) = 0] \\ &= 1 - \Pr[Y(X_g) = 1 \mid \hat{D}(\hat{S}(X_g), g) = 0] \\ &= 1 - \mathbb{E}[\hat{S}(X_g) \mid \hat{D}(\hat{S}(X_g), g) = 0] \end{aligned}$$

□

Using Proposition 3.1, we can geometrically see how certain post-processing decision rules will interact with the AP for a group g . For example, using a threshold, the expected true positives, true negatives, false positives, and false negatives can be estimated, as shown in Figure 2.

Proposition 3.2 below gives a characterizations of the false positive and false negative rates in a manner analogous to how Proposition 3.1 describes PPV and NPV:

Proposition 3.2. *Let \hat{Y} and \hat{S} be hard and soft classifiers as in Proposition 3.1. Then for any $g \in \mathcal{G}$,*

$$\begin{aligned} \text{FPR}_{\hat{Y},g} &= \Pr[\hat{Y}(X_g) = 1] \cdot \frac{1 - \mathbb{E}[\hat{S}(X_g) \mid \hat{Y}(X_g) = 1]}{1 - \mathbb{E}[\hat{S}(X_g)]} \\ \text{FNR}_{\hat{Y},g} &= \Pr[\hat{Y}(X_g) = 0] \cdot \frac{\mathbb{E}[\hat{S}(X_g) \mid \hat{Y}(X_g) = 0]}{\mathbb{E}[\hat{S}(X_g)]} \end{aligned}$$

Assume that $\Pr[Y(X_g) = 1] > 0$ and $\Pr[Y(X_g) = 0] > 0$ (that is, assume $0 < BR_g < 1$) so that FPR and FNR are well-defined.

Proof of Proposition 3.2. We give the proof for FPR, and the proof for FNR is similar. By applying

Bayes' rule, we can write

$$\begin{aligned} \text{FPR}_{\hat{Y},g} &= \Pr[\hat{Y}(X_g) = 1 \mid Y(X_g) = 0] \\ &= \Pr[Y(X_g) = 0 \mid \hat{Y}(X_g) = 1] \cdot \frac{\Pr[\hat{Y}(X_g) = 1]}{\Pr[Y(X_g) = 0]} \end{aligned} \quad (3)$$

Noting that $\Pr[Y(X_g) = 0 \mid \hat{Y}(X_g) = 1] = 1 - \text{PPV}_{\hat{Y},g}$, we can apply Proposition 3.1 and rearrange to write the RHS of Equation 3 as follows.

$$\text{RHS of (3)} = \Pr[\hat{Y}(X_g) = 1] \cdot \frac{1 - \mathbb{E}[\hat{S}(X_g) \mid \hat{Y}(X_g) = 1]}{\Pr[Y(X_g) = 0]} \quad (4)$$

We note that $\Pr[Y(X_g) = 0] = 1 - \mathbb{E}[\hat{S}(X_g)]$ (Proposition 2.1). Substituting this in to the RHS of Equation 4, we conclude the result. \square

3.2 General impossibility of equalizing PPV, NPV

It is not always possible to directly post-process a soft groupwise calibrated classifier into a hard one with equalized PPV (or NPV) for all groups, as we demonstrate by counterexample in Proposition 3.3. Before proceeding, we note that our counterexample is somewhat contrived—in particular, the AP induced by the soft classifier \hat{S} in the proof of Proposition 3.3 takes only one value on each group. When the AP of \hat{S} is more nicely structured on each group, we will see that there are general methods to equalize PPV (or NPV).

Proposition 3.3. *Fix two disjoint groups g_1 and g_2 with respective base rates BR_1 and BR_2 such that $BR_1 \neq BR_2$. Then there exists a soft classifier \hat{S} that is groupwise calibrated, but for which there is no post-processor $\hat{D} : [0, 1] \times \mathcal{G} \rightarrow \{0, 1\}$ such that $\hat{D} \circ \hat{S}$ equalizes PPV, unless $\Pr[\hat{D}(BR_i, g_i) = 1] = 0$ for $i = 1$ or 2 .*

Proof of Proposition 3.3. Consider the classifier \hat{S} such that $\hat{S}(x) = BR_1$ if $x \in g_1$ and $\hat{S}(x) = BR_2$ if $x \in g_2$. This classifier is trivially groupwise calibrated. Since $\Pr[\hat{D}(BR_i, g_i) = 1] > 0$ for $i = 1$ and 2 , we conclude that $\text{PPV}_{\hat{Y},g_i}$ is well-defined for g_1 and g_2 . The proof now follows from the characterization of PPV in Proposition 3.1. This is because $\text{PPV}_{\hat{Y},g_i}$ is equal to the expectation of $\hat{S}(X)$ where X is drawn from a distribution with support contained in g_i , and hence it is equal to BR_i , and $BR_1 \neq BR_2$. \square

The analogous statement regarding impossibility of equalizing NPV is formulated as Proposition A.1 in Appendix A.1.

3.3 A niceness Condition for APs

We now give a non-degeneracy condition on APs motivated by the impossibility result for post-processing given by Proposition 3.3.

Definition 3.1 (Niceness of APs). Let \mathcal{G} be a set of groups. A distribution on calibrated scores $\hat{\mathcal{P}}$ is *nice* if $\text{Supp}(\hat{\mathcal{P}}_g)$ is the same for all $g \in \mathcal{G}$.

Note that this condition rules out the counterexample given by Proposition 3.3, since the APs in the counterexample had different (in fact, disjoint) supports for different groups. Hence, we can hope to successfully post-process soft classifiers with nice APs.

3.4 Equalizing PPV or NPV by Thresholding

We pay special attention to thresholds because they are simple to understand and therefore very widely used. We use one slight modification to deterministic thresholds that adds an element of randomness: if a score is *at* the threshold, we randomly determine which side of the threshold it falls on, according to a distribution defined below.

Definition 3.2 (Threshold Post-Processor). A threshold post-processor $\hat{D}_{(\tau, \mathcal{R})} : [0, 1] \times \mathcal{G} \rightarrow \{1, 0\}$ is a function from a score $s \in [0, 1]$ and a group $g \in \mathcal{G}$, parameterized by τ and \mathcal{R} . The threshold parameter $\tau : \mathcal{G} \rightarrow [0, 1]$ specifies the threshold for the group g , and $\mathcal{R} : \mathcal{G} \rightarrow [0, 1]$ is the probability of returning 1 when the input score s is on the threshold $\tau(g)$. It returns the following outputs:

$$\hat{D}_{(\tau, \mathcal{R})}(s, g) = \begin{cases} 1 & s > \tau(g) \\ 0 & s < \tau(g) \\ 1 \text{ w.p. } \mathcal{R}(g) \text{ else } 0 & s = \tau(g) \end{cases}$$

In the setting of an infinite number of scores and a continuous domain (i.e. scores are represented by a probability density function instead of a probability mass function), we can use purely deterministic threshold functions in which $\mathcal{R} \equiv 1$, and achieve very similar results for the rest of this section.

If both $\tau(g)$ and $\mathcal{R}(g)$ do not vary across groups $g \in \mathcal{G}$, then the post-processor is the same across groups. In this case, we will call the post-processor a *group blind* threshold post-processor, and will overload τ and \mathcal{R} to be constants.

We now study the effectiveness of thresholds for post-processing soft classifiers with nice APs. The main takeaways are:

1. If the APs are nice, then threshold post-processors can equalize PPV (Propositions 3.4 and 3.6).
2. However, group blind threshold post-processors are rather limited in their ability to equalize PPV (Proposition 3.5).
3. Furthermore, equalizing PPV with thresholds (group blind or otherwise) may have undesirable social consequences.
4. Thresholds cannot always equalize PPV and NPV simultaneously, even for nice APs (Proposition 3.7).

Results 1-3 also apply to NPV (see Proposition A.3).

3.4.1 Group Blind Thresholds

We begin by classifying which group-blind threshold post-processors can equalize PPVs across all groups (Propositions 3.4 and 3.5). By symmetry, our arguments give a similar characterization for equalizing NPVs.

Proposition 3.4. *For every nice groupwise calibrated soft classifier \hat{S} and for every group-blind threshold post-processor $\hat{D}_{(\tau, \mathcal{R})}$ such that $\tau(g) = \max(\text{Supp}(\hat{P}_g))$ for all g , then the composed classifier $\hat{Y} = \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}$ equalizes PPVs across all groups for which \hat{Y} is non-trivial.*

The existence of the threshold post-processors in Proposition 3.4 follows from the assumed finiteness of the range of the soft classifier. In the case where the range of the soft classifier is infinite, such post-processors may not exist.

Proof of Proposition 3.4. Any of the given post-processors only ever maps the largest score in the support of $\hat{\mathcal{P}}_g$ to 1, for all groups g . Hence, PPV_g is exactly the largest score in $\text{Supp}(\hat{\mathcal{P}}_g)$. By the assumption that $\hat{\mathcal{P}}$ is nice, $\text{Supp}(\hat{\mathcal{P}}_g)$ is the same for all groups g , and hence the PPV is equalized across groups. \square

We prove the analogous statement for NPV in Proposition A.2 in the Appendix. We proceed to show that the post-processors described in Proposition 3.4 are the *only* non-trivial, group blind post-processors that equalize PPV across groups in general, as we prove in Proposition 3.5.

Proposition 3.5. *There exists a groupwise-calibrated soft classifier with a nice AP for which no non-trivial group blind threshold post-processor, other than the ones in Proposition 3.4, can equalize PPV across groups.*

At a high level, the proof of Proposition 3.5 works as follows: We can make the AP on one group uniform, and the AP of another group strictly increasing. Then, threshold post-processors naturally favor the latter group, as the AP for that group gives more weight to higher scores than lower ones when compared to the former AP. Our characterization of PPV (Proposition 3.1) features prominently in the proof.

In preparation to proving Proposition 3.5, we first prove the following lemma:

Lemma 3.1. *Let $g_1, g_2 \in \mathcal{G}$ be two different groups, and fix a group-blind threshold post-processor $\hat{D}_{(\tau, \mathcal{R})}$. Let $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$ be the expected conditional AP on scores $\geq \tau$ that results from starting with the AP $\hat{\mathcal{P}}_{g_1}$ over scores in group g_1 and conditioning on the scores that $\hat{D}_{(\tau, \mathcal{R})}$ sends to 1, and similarly let $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ denote the same type of conditional AP when starting with the $\hat{\mathcal{P}}_{g_2}$ over scores in group g_2 . If $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ strictly stochastically dominates $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$, then*

$$\text{PPV}_{\hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}, g_1} < \text{PPV}_{\hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}, g_2}$$

Proof. We use the characterization of PPV given in Proposition 3.1, for the special case where the post-processor thresholds as described above. We can write the PPV for group g_1 as follows:

$$\begin{aligned} \text{PPV}_{\hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}, g_1} &= \mathbb{E}_{X_1 \sim \mathcal{X}_{g_1}} [\hat{S}(X_1) \mid \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}(X_1) = 1] \\ &= \mathbb{E}_{s \sim \hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}} [s] \end{aligned} \tag{5}$$

where the second line follows from the definition of $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$.

Similarly, we have that

$$\text{PPV}_{\hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}, g_2} = \mathbb{E}_{s \sim \hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}} [s] \tag{6}$$

Since $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ stochastically dominates $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$, the expectation on the RHS of Equation 6 is larger than the expectation on the RHS of Equation 5, yielding the result. \square

Proof of Proposition 3.5. Fix two groups g_1 and g_2 and a finite set of points $S \subset [0, 1]$ such that the PMFs of the soft classifier \hat{S} on g_1 and g_2 have support equal to S - that is, $\text{Supp}(\hat{\mathcal{P}}_{g_1}) = \text{Supp}(\hat{\mathcal{P}}_{g_2}) = S$. For concreteness, suppose that $|S| = 10$.

Let the PMFs of the soft classifier \hat{S} on these two groups respectively be given by $\hat{\mathcal{P}}_{g_1}(s) = 1/|\text{Supp}(\hat{\mathcal{P}}_{g_1})|$ and $\hat{\mathcal{P}}_{g_2}(s) \propto s$ for all $s \in S$, where $\hat{\mathcal{P}}_{g_2}$ is normalized with a constant such that it sums to 1. Fix a group blind threshold post-processor $\hat{D}_{(\tau, \mathcal{R})}$ that is not one of the ones mentioned in Proposition 3.4. Since $\hat{S}_{(\tau, \mathcal{R})}$ is group blind, its threshold function is a constant which we name τ .

Let $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$ be the expected conditional AP on scores $\geq \tau$ that results from starting with the AP $\hat{\mathcal{P}}_{g_1}$ over scores in group g_1 and conditioning on the scores that $\hat{D}_{(\tau, \mathcal{R})}$ sends to 1. We can get this conditional PMF by removing scores $s < \tau$, multiplying $\hat{\mathcal{P}}_{g_1}(\tau)$ by \mathcal{R} , and re-normalizing the remaining values to get a distribution. Let $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ be defined similarly.

We claim that $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ strictly stochastically dominates $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$, which allows us to invoke Lemma 3.1 to conclude that the PPV on the two groups are unequal. We now show that $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ strictly stochastically dominates $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$. This is clearly true by design if $\mathcal{R} = 0$ or 1 : in this case, the post-processor is simply a deterministic threshold function, and we know by design that $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$ is uniform while $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ is a strictly increasing function. If $\mathcal{R} = r$ for some $r \in (0, 1)$, then we can write $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$ as a convex combination of $\hat{\mathcal{P}}_{g_1, \tau, 0}$ and $\hat{\mathcal{P}}_{g_1, \tau, 1}$ (with weight r on the distribution where $\mathcal{R} = 1$, and weight $1 - r$ on the distribution where $\mathcal{R} = 0$). We can write $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ as the same convex combination of the conditional distributions over g_2 where $\mathcal{R} = 0$ and $\mathcal{R} = 1$. Since we already established stochastic domination for the cases where $\mathcal{R} = 0$ and $\mathcal{R} = 1$, this establishes stochastic domination for the case where $\mathcal{R} \in (0, 1)$. \square

We achieve the same result for NPV in Proposition A.3. In the setting where the range of the soft classifier is infinite and continuous, we show in Proposition A.5 that a similar negative result holds, but without the existence of the classifiers in Proposition 3.4.

Propositions 3.4 and 3.5 demonstrate the limitations of group blind thresholds on calibrated scores. Though this method of post-processing has social appeal, it does not actually preserve the fairness properties that one would expect. In the next section we repeat our analysis but relax our group blindness requirement.

3.4.2 Group-Aware Thresholds

If we allow the different groups to have different thresholds, then we grant ourselves more degrees of freedom to be able to satisfy binary fairness constraints. In particular, we can equalize PPV across groups in a more meaningful way than done in Proposition 3.4.

Recall that the group blind threshold post-processors in Proposition 3.4 are the only group blind threshold post-processors that work on certain nice APs (shown in Proposition 3.5). However, these post-processors have the property that the only score they map to 1 is the largest score in the support, which can be undesirable for many applications.

In particular, all classifiers in Proposition 3.5 make the PPV on each group g_i equal to the maximum score in the support of $\hat{\mathcal{P}}_{g_i}$. However, the (not-necessarily-group-blind) threshold post-processors in Proposition 3.6 below can make the PPV on each group equal to any fixed value between the maximum base rate of g_i and the maximum score in $\text{Supp}(\hat{\mathcal{P}}_{g_i})$.

Proposition 3.6. *Let \mathcal{G} be a set of groups. For any soft classifier \hat{S} with a nice AP $\hat{\mathcal{P}}$ such that \hat{S} is groupwise calibrated over \mathcal{G} and $|\text{Supp}(\hat{\mathcal{P}}_g)| \geq 2$ for all $g \in \mathcal{G}$, then there exists a non group*

blind, non-trivial threshold post-processor $\hat{D}_{(\tau, \mathcal{R})}$ that is not one of the ones from Proposition 3.4 such that the hard classifier $\hat{Y} = \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}$ equalizes PPV across \mathcal{G} .

This holds even if we require that the PPV of all the groups is equal to an arbitrary value in $(\max_i \text{BR}_{g_i}, s_{\max})$, where $\max_i \text{BR}_{g_i}$ is the maximum base rate among the groups $g_i \in \mathcal{G}$ and s_{\max} is the maximum score in the support of $\hat{\mathcal{P}}_{g_i}$.²

Moreover, since this post-processor is not group blind, it is not one of the post-processors described in Proposition 3.4.

In preparation to proving Proposition 3.6, we first prove the following claim:

Claim 3.1 (Monotonicity of PPV and NPV). *Fix a soft classifier \hat{S} and corresponding AP $\hat{\mathcal{P}}$, as well as a group g . Fix group blind threshold post-processors $\hat{D}_{\tau_1, \mathcal{R}_1}$ and $\hat{D}_{\tau_2, \mathcal{R}_2}$ such that either $\tau_1 < \tau_2$ or $\tau_1 = \tau_2$ and $\mathcal{R}_1 \geq \mathcal{R}_2$. Then:*

- (a) $\text{PPV}_{\hat{D}_{\tau_1, \mathcal{R}_1} \circ \hat{S}, g} \leq \text{PPV}_{\hat{D}_{\tau_2, \mathcal{R}_2} \circ \hat{S}, g}$
- (b) $\text{NPV}_{\hat{D}_{\tau_1, \mathcal{R}_1} \circ \hat{S}, g} \leq \text{NPV}_{\hat{D}_{\tau_2, \mathcal{R}_2} \circ \hat{S}, g}$

Proof. We show conclusion (a); conclusion (b) is shown analogously. Define $\hat{\mathcal{P}}_{g, \tau_1, \mathcal{R}_1}$ to be the conditional PMF on scores $\geq \tau_1$ that results from starting with the AP $\hat{\mathcal{P}}_g$ over scores in group g and conditioning on the scores that $\hat{D}_{\tau_1, \mathcal{R}_1}$ sends to 1, and let $\hat{\mathcal{P}}_{g, \tau_2, \mathcal{R}_2}$ be defined similarly (but for the threshold post-processor $\hat{D}_{\tau_2, \mathcal{R}_2}$).

We claim that $\hat{\mathcal{P}}_{g, \tau_2, \mathcal{R}_2}$ stochastically dominates $\hat{\mathcal{P}}_{g, \tau_1, \mathcal{R}_1}$, which yields the desired result by the characterization of PPV given in Proposition 3.1 (and more explicitly written in Equations 5 and 6). \square

Proof of Proposition 3.6. Fix a soft classifier \hat{S} with a nice AP $\hat{\mathcal{P}}$ that is group-wise calibrated over g_1, \dots, g_n , and fix a desired value $v \in (\max_i \text{BR}_{g_i}, s_{\max})$. We will show that we can design a threshold post-processor (τ, \mathcal{R}) such that $\text{PPV}_{g, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}} = v$ for all groups g .

Fix an arbitrary group g_j . We proceed via a continuity argument to show that we can tune the threshold on g_j to achieve PPV equal to v . The maximum possible value for $\text{PPV}_{g_j, \hat{D}_{(\tau, \mathcal{R})}}$ is s_{\max} (achieved when $\tau = s_{\max}$, by Claim 3.1), where s_{\max} is the largest score in the support, as defined in the proposition statement³.

Furthermore, note that, for any group, a lower bound on the PPV of a hard classifier on that group is the base rate of the group, where the lower bound is matched by the trivial post-processor that sends every score to 1. This follows from Claim 3.1.

We now claim that there is a setting of $\tau(g_j)$ and $\mathcal{R}(g_j)$ that achieves $\text{PPV}_{g, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}} = v$. We accomplish this by showing that there is a way to change $(\tau(g_j), \mathcal{R}(g_j))$ such that the PPV decreases continuously. We first show:

Claim 3.2 (Continuity of PPV). *Fix a soft classifier \hat{S} and corresponding AP $\hat{\mathcal{P}}$, as well as a group g . Suppose we have two post-processing algorithms, \hat{D}_1 and \hat{D}_2 . Let $\hat{\mathcal{P}}_{g, \hat{D}_1}$ be the expected conditional AP that results from starting with the AP $\hat{\mathcal{P}}_g$ over scores in group g and conditioning on the scores that \hat{D}_1 sends to 1, and define $\hat{\mathcal{P}}_{g, \hat{D}_2}$ similarly. If $d_{TV}(\hat{\mathcal{P}}_{g, \hat{D}_1}, \hat{\mathcal{P}}_{g, \hat{D}_2}) < \epsilon$, then $|\text{PPV}_{g, \hat{D}_1 \circ \hat{S}} - \text{PPV}_{g, \hat{D}_2 \circ \hat{S}}| < O(\epsilon)$. Or in words, if the distance between the conditional APs is small, then the difference in PPV is small.*

²For the case where the support of $\hat{\mathcal{P}}_{g_i}$ is infinite, s_{\max} should be the supremum of scores.

³We ignore the trivial post-processor that never maps anything to 1, and hence leaves the PPV undefined.

Proof. Recall the characterization of PPV given in Proposition 3.1 (and more explicitly written in Equation 5). This tells us that the PPV of group g for the classifier $\hat{D}_1 \circ \hat{S}$ is exactly the expectation of a random variable distributed according to $\hat{\mathcal{P}}_{g, \hat{D}_1}$. Similarly, the PPV of group g for the classifier $\hat{D}_2 \circ \hat{S}$ is the expectation of a r.v. distributed according to $\hat{\mathcal{P}}_{g, \hat{D}_2}$. Since both $\hat{\mathcal{P}}_{g, \hat{D}_1}$ and $\hat{\mathcal{P}}_{g, \hat{D}_2}$ have support bounded between 0 and 1, their expectations can differ by at most ϵ , from which the claim follows. For completeness, we prove this below.

Suppose wlog that $\hat{\mathcal{P}}_{g, \hat{D}_1}$ has the larger expectation. Let $S = \{s \in \text{Supp}(\hat{\mathcal{P}}_g) : \hat{\mathcal{P}}_{g, \hat{D}_1}(s) > \hat{\mathcal{P}}_{g, \hat{D}_2}(s)\}$. Then:

$$\begin{aligned} \text{PPV}_{g, \hat{D}_1 \circ \hat{S}} &= \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} s \hat{\mathcal{P}}_{g, \hat{D}_1}(s) \\ &= \text{PPV}_{g, \hat{D}_2 \circ \hat{S}} + \sum_{s \in S} s(\hat{\mathcal{P}}_{g, \hat{D}_1}(s) - \hat{\mathcal{P}}_{g, \hat{D}_2}(s)) \\ &< \text{PPV}_{g, \hat{D}_2 \circ \hat{S}} + \epsilon \end{aligned}$$

where in the second line we use the fact that $\text{PPV}_{g, \hat{D}_2 \circ \hat{S}}$ is the expectation of $\hat{\mathcal{P}}_{g, \hat{D}_2}$, and in the last line we use the fact that $s \in [0, 1]$ and that the TV-distance between the two distributions is less than ϵ . \square

Now, consider the following way to change $(\tau(g_j), \mathcal{R}(g_j))$. Fix $\epsilon > 0$, and an initial setting for $(\tau(g_j), \mathcal{R}(g_j))$ s.t. $\tau(g_j)$ is not the smallest item in the support or $\mathcal{R}(g_j) > \epsilon$. Reduce $\mathcal{R}(g_j)$ by ϵ , wrapping around on the interval $(0, 1]$ and decreasing $\tau(g_j)$ to the next largest item in the support when this would otherwise make $\mathcal{R}(g_j)$ negative.

This very minor transformation to the threshold changes the AP conditional on outputting 1 very slightly - so slightly that the TV distance between the old conditional AP and the new AP is at most some ϵ' which is a function of ϵ . This lets us apply Claim 3.2 to show that the PPV changes by at most a function of ϵ . So as we take ϵ going towards 0, this shows that the PPV changes by an amount going towards 0. This establishes that the PPV changes “continuously” with respect to this deforming procedure.

By Claim 3.1, we have that the above deforming procedure can only decrease the PPV. Therefore, we can continuously decrease the PPV, starting from s_{\max} , by continuously deforming the threshold post-processor with the method above. Note that $s_{\max} > v > \max_i \text{BR}_{g_i} \geq \text{BR}_{g_j}$. By the Intermediate Value Theorem, there must be a setting of $(\tau(g_j), \mathcal{R}(g_j))$ such that $\text{PPV}_{g_j, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}} = v$. \square

We assert the analogous statement for the case of NPV in Claim A.1. The corresponding statement for the case of soft classifiers with infinite range is asserted in Proposition A.6.

3.4.3 The Limitations of Thresholding

While Proposition 3.6 shows that a threshold post-processor can equalize the PPV across n groups, this threshold post-processor can be unsatisfying from a social justice standpoint. Consider an example with two groups g_1 and g_2 , where group g_2 is “privileged” by having a higher base rate of, say, credit worthiness. Suppose that we have an AP that is decreasing with respect to score on group g_1 , and increasing with respect to score on group g_2 . This is illustrated in Example 3.1




Figure 3: Accompanying Example 3.1, the PPV for both groups is 0.77. However, the threshold for g_1 (dark blue) is higher than the threshold for g_2 (orange), even though g_2 is likely the more privileged group.

and Figure 3. This means that a group blind threshold post-processor yields larger PPV on g_2 , since large scores are given more weight in g_2 . So, to equalize the PPV between the two groups, we will classify more low scores as positive in g_2 than g_1 . This effectively means that our threshold on group g_2 is *more lenient* than our threshold on g_1 , which seems blatantly unfair, since g_2 is the privileged group in the first place!

Example 3.1 (Socially Unsatisfying Example). *Fix groups g_1 and g_2 , and we fix the AP of the soft classifier \hat{S} as follows. Let $\text{Supp}(\hat{P}_{g_i})$ be $\{0, 0.01, 0.02, \dots, 1\}$ for $i = 1$ and 2 , let $\hat{P}_{g_1}(s) \propto a - s$ for appropriately selected constant $a > 0$ and let $\hat{P}_{g_2}(s) \propto s$. Group g_2 has a higher base rate and may have social advantages over group g_1 .*

Let $\hat{D}_{(\tau, \mathcal{R})}$ be a non-trivial post-processor. If $\hat{D}_{(\tau, \mathcal{R})}$ were group blind, then by Lemma 3.1, since $\hat{P}_{g_2, \hat{D}_{(\tau, \mathcal{R})}}$ stochastically dominates $\hat{P}_{g_1, \hat{D}_{(\tau, \mathcal{R})}}$, the PPV on g_2 must be larger than the PPV of g_1 .

To equalize PPV, by Claim 3.1, we must have either $\tau(g_2) < \tau(g_1)$, or $\tau(g_2) = \tau(g_1)$ and $\mathcal{R}(g_2) < \mathcal{R}(g_1)$. The disadvantaged group is now held to a higher standard than the privileged group to maintain equality of PPV. Figure 3 illustrates example thresholds that equalize the PPV.

The only property needed for Example 3.1 is that the AP of one (privileged) group stochastically dominates the AP of another (unprivileged) group. We suspect that this will occur in many settings. Indeed, the COMPAS scores we analyze in Section 5 have this property, when considering Caucasians as the privileged group and African-Americans the unprivileged group. As in the above example, using group-aware thresholds to equalize PPV between the groups on COMPAS data results in a more permissive decision rule for Caucasians, demonstrating a problem with this approach (see also Figure 7).

Furthermore, thresholding cannot in general equalize both PPV and NPV simultaneously, even for nice APs and using non-group blind thresholds.

Proposition 3.7. *Fix groups g_1 and g_2 . There exists a soft classifier \hat{S} with a nice AP \hat{P} such that no threshold post-processor can simultaneously equalize PPV and NPV between groups g_1 and g_2 .*

Before proving the statement, we first show that the base rates of g_1 and g_2 can be written as convex combinations of the PPV and $1 - \text{NPV}$ on the respective groups:

Claim 3.3. *Fix any group $g \in \mathcal{G}$, and let the hard classifier \hat{Y} be non-trivial. Then the base rate of g can be written as a convex combination of $\text{PPV}_{g,\hat{Y}}$ and $1 - \text{NPV}_{g,\hat{Y}}$*

Proof.

$$\begin{aligned} \text{BR}_g &= \Pr[Y(X_g) = 1] \\ &= \Pr[Y(X_g) = 1 \mid \hat{Y}(X_g) = 1] \Pr[\hat{Y}(X_g) = 1] \\ &\quad + \Pr[Y(X_g) = 1 \mid \hat{Y}(X_g) = 0] \Pr[\hat{Y}(X_g) = 0] \\ &= \text{PPV}_{g,\hat{Y}} \cdot \theta + (1 - \text{NPV}_{g,\hat{Y}}) \cdot (1 - \theta) \end{aligned}$$

where $\theta := \Pr[\hat{Y}(X_g) = 1]$. □

A simple intuition for the proof of Proposition 3.7 is as follows. Suppose we have two groups, and the soft classifier is almost perfect on one group - for all but a small fraction of people, it gives the correct binary score, and gives the remaining people score 0.5. On the other group, it is the opposite - almost every person is given score 0.5, and there are only a few people given their ground truth score. The corresponding APs of each group have equal supports, and therefore are “nice.” However, it is clear that any threshold post-processor on the first group will have extremely high PPV and NPV, while any threshold post-processor on the second group will have to make decision on where to round the people in the 0.5 bucket, and will correspondingly either have low PPV or NPV.

This example is somewhat unsatisfying, as the AP satisfies niceness by a technicality (for example, it is extremely close to a AP with disjoint supports between groups). The proof generalizes the above example to a case where scores for one group are slightly more “correlated” with the ground truth labels than scores for the other group. Appendix A.2 shows the corresponding result when the range of the soft classifier is infinite.

Proof of Proposition 3.7. Let $\text{Supp}(\hat{\mathcal{P}}_{g_i})$ be $\{0.1, 0.2, \dots, 0.9\}$ for $i = 1$ and 2. Let $\hat{\mathcal{P}}_{g_1}$ be uniform over scores, and let $\hat{\mathcal{P}}_{g_2}(s) = -a(s - 1/2)^2 + b$ for some appropriately chosen constants $a, b \geq 0$ such that $\hat{\mathcal{P}}_{g_2}$ is a valid probability distribution.

We claim that there is no threshold post-processor that can equalize PPV and NPV simultaneously for these two groups. First, note that the base rates for the two groups are equal to $1/2$ by design, due to the symmetric nature of the AP on each group. Just like in the proof of Proposition 3.5, we use the notation $\hat{\mathcal{P}}_{g_1, \geq(\tau, \mathcal{R})}$ to denote the conditional AP supported on scores $\geq \tau$ that results from starting with the AP $\hat{\mathcal{P}}_{g_1}$ over all scores in group g_1 and conditioning on the scores that $\hat{D}_{(\tau, \mathcal{R})}$ sends to 1, and define $\hat{\mathcal{P}}_{g_2, \geq(\tau, \mathcal{R})}$ similarly. Let $\hat{\mathcal{P}}_{g_i, \leq(\tau, \mathcal{R})}$ denote the conditional AP starting from group g_i and conditioning on the scores that $\hat{D}_{(\tau, \mathcal{R})}$ sends to 0.

We proceed by a case analysis on the location of the threshold for group g_1 - that is, on $\tau(g_1)$.

Case 1: $\tau(g_1) \geq 1/2$. First, suppose additionally that $\hat{D}_{(\tau, \mathcal{R})}$ is group blind over g_1, g_2 , then $\hat{\mathcal{P}}_{g_2, (\tau, \mathcal{R})}$ is strictly stochastically dominated by $\hat{\mathcal{P}}_{g_1, (\tau, \mathcal{R})}$ so the PPV is lower on g_2 than g_1 (Lemma 3.1). Therefore, to equalize the PPV between the two groups, the threshold on group g_2 must be to the

“right” of the threshold on g_1 (that is, either $\tau(g_2) > \tau(g_1)$ or $\tau(g_2) = \tau(g_1)$ and $\mathcal{R}(g_2) < \mathcal{R}(g_1)$), which follows from Claim 3.1.

However, we claim that this setting of thresholds makes the NPV on group g_2 lower than the NPV on group g_1 . It suffices to show that, for any constant threshold post-processor classifier $\hat{D}_{(\tau, \mathcal{R})}$ with $\tau(g_1) = \tau(g_2) \geq 1/2$, the NPV on group g_2 is lower than the NPV on group g_1 . This implies that the same statement holds when either $\tau(g_2) > \tau(g_1)$ or $\tau(g_2) = \tau(g_1)$ and $\mathcal{R}(g_2) < \mathcal{R}(g_1)$ as well, due to the monotonicity property of NPV (Claim 3.1). This suffices to prove Proposition 3.7 for the case when $\tau(g_1) \geq 1/2$: we need to set $\tau(g_2), \mathcal{R}(g_2)$ such that either $\tau(g_2) > \tau(g_1)$ or $\tau(g_2) = \tau(g_1)$ and $\mathcal{R}(g_2) < \mathcal{R}(g_1)$ in order to equalize PPV, but this leaves the NPV on group g_2 lower than the NPV on group g_1 .

Now we proceed to show that the NPV is lower on g_2 when the threshold post-processor is constant. Fix a constant threshold post-processor $\hat{D}_{(\tau, \mathcal{R})}$ with $\tau \geq 1/2$. We have already established that this means that $\text{PPV}_{g_2, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}} < \text{PPV}_{g_1, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}}$.

Now, since the base rates of g_1 and g_2 are equal, this implies that

$$1 - \text{NPV}_{g_2, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}} > 1 - \text{NPV}_{g_1, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}}$$

Rearranging, this shows that $\text{NPV}_{g_2, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}} < \text{NPV}_{g_1, \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}}$, finishing the proof of this case.

Case 2: $\tau(g_1) < 1/2$. The argument in this case is symmetrical to the previous case, where we switch the roles of PPV and NPV in the argument. We sketch it for the sake of completeness.

Fix a constant threshold post-processor such that $\tau(g_1) = \tau(g_2) < 1/2$. By design, $\hat{\mathcal{P}}_{g_2, \leq (\tau, \mathcal{R})}$ strictly stochastically dominates $\hat{\mathcal{P}}_{g_1, \leq (\tau, \mathcal{R})}$. Hence, the NPV on group g_2 is smaller than the NPV on group g_1 (this follows from an analogous version of Lemma 3.1 for NPV instead of PPV). This means that the threshold post-processor cannot be constant to equalize NPVs - it must be moved such that either $\tau(g_2) < \tau(g_1)$ or $\tau(g_2) = \tau(g_1)$ and $\mathcal{R}(g_2) > \mathcal{R}(g_1)$ (Claim 3.1).

However, by the same convex combination argument as in the previous case, we get that any such constant threshold post-processor must make the PPV on g_2 strictly smaller than the PPV on g_1 . By the monotonicity of PPV, this means that any threshold post-processor with either $\tau(g_2) < \tau(g_1)$ or $\tau(g_2) = \tau(g_1)$ and $\mathcal{R}(g_2) > \mathcal{R}(g_1)$ must also make the PPV on g_2 smaller than the PPV on g_1 , finishing the proof. \square

3.5 Equalizing APs

While thresholding is a conceptually simple approach to post-processing a soft classifier, its power is limited. We now consider a very different approach using soft post-processors to equalize the APs across groups of a soft classifier. The intuition is that if the APs are equal across groups, then any hard post-processor that is *group blind* should result in equal PPV, NPV, FPR, and FNR. We formalize this intuition in Claim 3.4.

Let \hat{S} be a soft classifier and for each group $g \in \mathcal{G}$, let $\hat{\mathcal{P}}_g$ be the AP of \hat{S} for group g . For a soft post-processor \hat{D}^{soft} , let $\hat{S}' = \hat{D}^{\text{soft}} \circ \hat{S}$ and let $\hat{\mathcal{P}}'_g$ be the corresponding AP for group g .

Our goal is to find a soft post-processor \hat{D}^{soft} such that \hat{S}' is groupwise calibrated, and $\hat{\mathcal{P}}'_g = \hat{\mathcal{P}}'_{g'}$ for all $g, g' \in \mathcal{G}$. In this section, we describe only one approach to constructing \hat{D}^{soft} which we call *mass averaging*.

The approach of equalizing APs has a fundamental weakness: if $\hat{\mathcal{P}}'_g = \hat{\mathcal{P}}'_{g'}$ and both are calibrated, then $\text{BR}_g = \text{BR}_{g'}$. This severely limits applicability of this approach. However, this limitation will be removed in Section 4.2 by allowing deferrals.

Claim 3.4. *If the APs are equal for two groups, then PPV, NPV, FPR, and FNR are equalized by any hard post-processor \hat{D} satisfying group blindness.*

The group-blindness requirement in the claim is necessary: consider the (not group blind) post-processor that outputs 0 on one group and 1 on the other; PPV will not be equalized.

Proof of Claim: 3.4. We prove only that PPV is equalized; the remaining properties may be proved similarly. Let $\hat{Y} = \hat{D} \circ \hat{S}$ be a hard classifier \hat{Y} that is a group blind post-processor \hat{D} composed with a calibrated soft classifier \hat{S} . All probabilities below are over $X_g \sim \mathcal{X}_g$, and the coins of \hat{S} and \hat{D} .

$$\begin{aligned} \text{PPV}_{\hat{Y},g} &= \Pr[Y(X_g) = 1 \mid \hat{Y}(X_g, g) = 1] \\ &= \frac{\Pr[Y(X) = 1]}{\Pr[\hat{Y}(X_g, g) = 1]} \\ &\quad \cdot \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \left(\frac{\Pr[Y(X_g) = 1 \mid \hat{S}(X_g) = s] \cdot \Pr[\hat{S}(X_g) = s]}{\Pr[Y(X_g) = 1]} \right. \\ &\quad \left. \cdot \Pr[\hat{Y}(X_g, g) = 1 \mid \hat{S}(X_g) = s, Y(X_g) = 1] \right) \end{aligned}$$

Each factor in this product is equal across groups by the assumptions. Namely, $\Pr[Y(X_g) = 1]$ is equalized by calibration and equalized APs; $\Pr[\hat{Y}(X_g, g) = 1]$ by group blindness and equalized APs; $\Pr[\hat{Y}(X_g, g) = 1 \mid \hat{S}(X_g) = s, Y(X_g) = 1]$ by group blindness; $\Pr[Y(X_g) = 1 \mid \hat{S}(X_g) = s]$ by calibration; and finally $\Pr[\hat{S}(X_g) = s]$ by equalized APs. \square

3.5.1 Mass Averaging

The mass-averaging technique is best illustrated with an example. Suppose that $\hat{\mathcal{P}}_{g_1}$ is uniform over $\{0, 0.5, 1\}$, and $\hat{\mathcal{P}}_{g_2}$ is uniform over $\{0, 1\}$. It is easy to define a soft post-processor \hat{D}^{soft} which equalizes these two APs. On g_1 , we leave the score unchanged: $\hat{D}^{\text{soft}}(s, g_1) = s$. On g_2 , we compute the output as

$$\hat{D}^{\text{soft}}(s, g_2) = \begin{cases} s & \text{w.p. } 2/3 \\ 0.5 & \text{w.p. } 1/3 \end{cases}.$$

The APs for groups g_1 and g_2 of the resulting soft classifier $\hat{S}' = \hat{D}^{\text{soft}} \circ \hat{S}$ are equal, and are equal to $\hat{\mathcal{P}}_{g_1}$.

In the example, the probability mass is being redistributed by averaging the scores. This can be equivalently viewed as adding noise to the scores and then recalibrating the scores, something discussed in [CPF⁺17].

More generally, a mass-averaging post processor \hat{D}^{soft} assigns to each possible pair (s, g) a distribution over possible output scores s' . Such a \hat{D}^{soft} is fully specified by $k \cdot k' \cdot |\mathcal{G}|$ parameters, where k is the number of possible values of s and k' is the number of possible values of s' . Given a soft classifier \hat{S} and a mass-averaging post processor \hat{D}^{soft} , the constraint that the resulting APs

are equalized across groups is linear in these parameters. Such classifiers, therefore, may be found by a linear program. We do not explore the choice of mass-averaging post-processors further.

4 Post-Processing Calibrated Classifiers with Deferrals

In the first part of the paper, we considered the problem of post-processing calibrated soft classifiers, which output a score $s \in [0, 1]$, into fair hard classifiers, which output a decision in $\hat{y} \in \{0, 1\}$, subject to a number of group fairness conditions. In the remainder of this work, we reconsider this problem, but with one important change: we allow classifiers to “refuse to decide” by outputting the special symbol \perp . We call such classifiers *deferring* classifiers, borrowing the nomenclature from [MPZ17]. The output \perp is the deferring classifier’s way of refusing to make a decision and deferring to a downstream decision maker. For example, a risk assessment tool might aid a parole board to make a decision by categorizing an individual as high risk or low risk, or it might output \perp —providing no advice and deferring to the judgment of the board.

We now modify our notation appropriately. Instances x are still associated with a true type $Y(x) \in \{0, 1\}$ and a group $G(x) \in \mathcal{G}$. A deferring hard classifier \hat{Y} is a randomized function $\hat{Y} : \mathcal{X} \rightarrow \{0, 1, \perp\}$. A deferring soft classifier is a randomized function $\hat{S} : \mathcal{X} \rightarrow [0, 1] \cup \{\perp\}$. A deferring hard (resp. soft) post-processor is a randomized function $\hat{D} : [0, 1] \cup \{\perp\} \times \mathcal{G} \rightarrow \{0, 1, \perp\}$ (resp. $\hat{D}^{\text{soft}} : [0, 1] \cup \{\perp\} \times \mathcal{G} \rightarrow [0, 1] \cup \{\perp\}$) that takes as input the output of a deferring soft and post-processes it into a deferring hard (resp. soft) classifier. We also introduce new versions of the FPR and FNR, conditioned on not deferring.

Definition 4.1. The *conditional false positive rate* and *conditional false negative rate* of a deferring hard classifier \hat{Y} for a group g are, respectively:

$$\begin{aligned} \text{cFPR}_{\hat{Y},g} &= \Pr[\hat{Y}(X_g) = 1 \mid Y(X_g) = 0, \hat{Y}(X_g) \neq \perp] \\ \text{cFNR}_{\hat{Y},g} &= \Pr[\hat{Y}(X_g) = 0 \mid Y(X_g) = 1, \hat{Y}(X_g) \neq \perp]. \end{aligned}$$

We additionally consider a version of the accuracy profile conditioned on not deferring, which we call the *conditional AP*. For non-deferring soft classifiers, Definitions 4.2 and 2.4 coincide.

Definition 4.2. The *conditional AP* $\hat{\mathcal{P}}_g$ of a classifier \hat{S} for a group g is the PMF of $\hat{S}(X_g)$, conditioned on not outputting \perp . That is, for $s \in [0, 1]$, $\hat{\mathcal{P}}_g(s) = \Pr[\hat{S}(X_g) = s \mid \hat{S}(X_g) \neq \perp]$. Note that the conditional AP is undefined if $\Pr[\hat{S}(X_g) \neq \perp] = 0$.

Abusing notation, we denote by $\hat{\mathcal{P}}$ the collection $\{\hat{\mathcal{P}}_g\}_{g \in \mathcal{G}}$, and call it the *conditional AP* of \hat{S} .

The conditional error rates are applicable generally, but they can be difficult to interpret. The consequences of using the conditional FPR and FNR are discussed further in Section 6 along with a discussion of different deferral models. They are also amenable to the consideration of additional goals which we will briefly address. For example, one could seek to minimize the total deferral rate, equalize the deferral rate among groups, or prefer deferrals on positive instances.

4.1 Thresholding with deferrals

We return now to the problem of post-processing of calibrated soft classifiers, but now with the extra power of deferring on some inputs. We revisit the two approaches discussed in Section 3: thresholding and equalizing APs.

Proposition 3.3 stated PPV and NPV cannot both be equalized across groups in general when using only a single threshold per group. By using two thresholds per groups and deferring on some inputs, PPV and NPV can always be equalized across groups.

We post-process using two thresholds per group as follows: return 0 when s is lower than the first threshold, return \perp between the thresholds, and return 1 above the second threshold, as shown in Figure 4. This buys us more degrees of freedom when equalizing binary constraints, and it has the useful property that we say \perp on the instances where we are the least confident about the predicted type.

We adapt our notation as follows:

Definition 4.3 (Deferring Threshold Post-Processor). A deferring threshold post-processor $\hat{D}_{(\tau_0, \tau_1, \mathcal{R}_0, \mathcal{R}_1)}$ assigns to each group g two thresholds $\tau_0(g), \tau_1(g) \in \text{Supp}(\hat{\mathcal{P}}_g)$, and two probabilities $\mathcal{R}_0(g), \mathcal{R}_1(g) \in [0, 1]$, with the following requirements:

1. for all $g \in \mathcal{G}$, $\tau_0(g) \leq \tau_1(g)$
2. for all $g \in \mathcal{G}$ for which $\tau_0(g) = \tau_1(g)$, $\mathcal{R}_1(g) + \mathcal{R}_0(g) \leq 1$. This corresponds to the case where the two thresholds are the same, and therefore individuals with that score must be mapped to 1 with probability $\mathcal{R}_1(g)$, and to 0 with probability $\mathcal{R}_0(g)$, with the remainder mapped to \perp .

The corresponding threshold post-processor is defined as follows:

$$\hat{D}_{(\tau_0, \tau_1, \mathcal{R}_0, \mathcal{R}_1)}(s, g) = \begin{cases} 1 & s > \tau_1(g) \\ 0 & s < \tau_0(g) \\ \perp & \tau_0(g) < s < \tau_1(g) \\ 1 \text{ w.p. } \mathcal{R}_1(g), \text{ else } \perp & s = \tau_1(g) \\ 0 \text{ w.p. } \mathcal{R}_0(g), \text{ else } \perp & s = \tau_0(g) \\ 1 \text{ w.p. } \mathcal{R}_1(g), 0 \text{ w.p. } \mathcal{R}_0(g), \text{ else } \perp & s = \tau_0(g) = \tau_1(g) \\ \perp & s = \perp \end{cases}$$

Using two thresholds allows the equalization of both PPV and NPV across groups in general, whereas without deferrals we could only equalize one or the other. We first demonstrate the existence of post-processors that are fairly limited, analogously to those defined in Proposition 3.4.

Proposition 4.1. *Let \hat{S} be a soft classifier with a nice AP for a set of groups \mathcal{G} . Then every threshold post-processor $\hat{D}_{(\tau_0, \tau_1, \mathcal{R}_0, \mathcal{R}_1)}$ satisfying the following properties equalizes both the PPV and NPV for all groups in \mathcal{G} of the composed classifier $\hat{Y} = \hat{D}_{(\tau_0, \tau_1, \mathcal{R}_0, \mathcal{R}_1)} \circ \hat{S}$:*

1. $\tau_0(g) = \min(\text{Supp}(\hat{\mathcal{P}}_g))$ for all g
2. $\mathcal{R}_0(g) > 0$ for all g .
3. $\tau_1(g) = \max(\text{Supp}(\hat{\mathcal{P}}_g))$ for all g
4. $\mathcal{R}_1(g) > 0$ for all g .




Figure 4: For threshold post-processors with deferrals, defer between the thresholds.

Notice that these classifiers cannot be trivial, because we defined \mathcal{R}_0 and \mathcal{R}_1 in a way that prohibits the possibility that the composed classifier never returns 0 or 1. For the cases where the range of soft classifier outputs is infinite and there is no max or min element, these classifiers do not exist.

Proof of Proposition 4.1. The reasoning is similar to the non-deferral case for equality of PPV alone. The thresholds only allow one score s to map to 0 and one score to map to 1. Thus, PPV for both groups is equal to the largest score in the support and NPV for both groups is equal to 1 minus the smallest score in the support. \square

Now, much like in Proposition 3.6, which showed the existence of meaningful non-trivial threshold post-processors that equalized PPV across groups, we show the existence of meaningful, non-trivial *deferring* threshold post-processors that equalize PPV and NPV across groups.

Proposition 4.2. *Let \hat{S} be a soft classifier with a nice AP that is groupwise calibrated for a set of groups \mathcal{G} . Suppose that $|\text{Supp}(\hat{\mathcal{P}}_g)| \geq 2$ for all $g \in \mathcal{G}$. Then there exists a non-trivial threshold post-processor $\hat{D}_{(\tau_0, \tau_1, \mathcal{R}_0, \mathcal{R}_1)}$ that is not one of those defined in Proposition 4.1, such that the hard classifier $\hat{Y} = \hat{D}_{(\tau_0, \tau_1, \mathcal{R}_0, \mathcal{R}_1)} \circ \hat{S}$ equalizes PPV_g and NPV_g for all $g \in \mathcal{G}$.*

The main idea of the proof of this proposition is to use Proposition 3.6 twice: once for getting thresholds to equalize the PPV, and once for thresholds to equalize the NPV. These thresholds may be invalid because there may be a group g for which $\tau_0(g) > \tau_1(g)$. We use Claims 3.1 and 3.2 to allow ourselves to push the PPV thresholds toward 1 and the NPV thresholds toward 0 until they no longer overlap, while still maintaining equalization of PPV and NPV for the other groups.

Proof of Proposition 4.2. Recall by Claim 3.1 that the PPV of \hat{S} on a group g monotonically increases as $\tau_1(g)$ increases and, if $\tau_1(g)$ is constant, as $\mathcal{R}_1(g)$ increases. By Claim 3.1, NPV monotonically increases as $\tau_0(g)$ decreases, and, if $\tau_0(g)$ is constant, as $\mathcal{R}_0(g)$ increases. Recall by Claim 3.2 that if the total variance distance between conditional APs (conditioned on being post-processed to a result of 1) for different groups is at most ϵ , then the PPV difference for these groups is bounded by $O(\epsilon)$. Thus, PPV_g is continuous and monotonically increasing with regard to $(\tau_1(g), \mathcal{R}_1(g))$. Similarly, NPV_g is continuous and monotonically increasing with $(-\tau_0(g), \mathcal{R}_0(g))$.

We know by Proposition 3.6 that there exists a non-group blind threshold rule (without deferrals) that equalizes the PPV among the groups. By the analogous Proposition A.4, there exists a (different) non-group-blind threshold rule that equalizes the NPV among the groups. For both of these, we know that they are not the classifiers from Proposition 4.1.

If the thresholds meet the conditions of being a deferring post-processor listed in Definition 4.3, then the statement is proven. If they do not meet the conditions because the thresholds “overlap,” we repeat the following procedure until the conditions are met:

1. Let g be a group for which the conditions are not met, i.e. either $\tau_0(g) > \tau_1(g)$, or $\tau_0(g) = \tau_1(g)$ and $\mathcal{R}_0(g) + \mathcal{R}_1(g) > 1$.
2. If $\tau_0(g) > \tau_1(g)$, define $t' = \frac{\tau_0(g) + \tau_1(g)}{2}$. Let $t = \arg \min_{s \in \text{Supp}(\hat{p}_g)} |s - t'|$. Notice that $t \leq \tau_0(g)$ and $t \geq \tau_1(g)$, but because by assumption $\tau_0(g) > \tau_1(g)$, t cannot be equal to both thresholds. Set the new value for both thresholds to t : $\tau_0(g) = \tau_1(g) = t(g)$.
3. If $\mathcal{R}_0(g) + \mathcal{R}_1(g) > 1$, then do the following:
 - (a) If $\tau_0(g)$ remained unaltered in the previous step, then keep $\mathcal{R}_0(g)$ the same, and set $\mathcal{R}_1(g) = 1 - \mathcal{R}_0(g)$.
 - (b) If $\tau_1(g)$ remained unaltered in the previous step, then set $\mathcal{R}_1(g)$ the same and set $\mathcal{R}_0(g) = 1 - \mathcal{R}_1(g)$.
 - (c) If neither of these is true, then let $r = \frac{\mathcal{R}_0(g)}{\mathcal{R}_0(g) + \mathcal{R}_1(g)}$. Set $\mathcal{R}_0(g) = r$ and $\mathcal{R}_1(g) = 1 - r$.
4. These thresholds are no longer overlapping, but they altered PPV_g and NPV_g . Notice that, by the monotonicity properties described above, the threshold rules were changed in ways that can only increase PPV_g or NPV_g :
 - (a) $\tau_1(g)$ has increased or remained constant
 - (b) if $\tau_1(g)$ remained constant, then $\mathcal{R}_1(g)$ also remained constant
 - (c) $\tau_0(g)$ has decreased or remained constant
 - (d) if $\tau_0(g)$ remained constant, then $\mathcal{R}_0(g)$ remained constant

The new PPV and NPV for g may now be higher than those of the other groups.

5. For all other groups $g' \neq g$, by the Intermediate Value Theorem and the continuity of NPV, there exists some $(\tau_0(g'), \mathcal{R}(g'))$ that sets $\text{NPV}_{g'} = \text{NPV}_g$, and by the monotonicity of NPV, this threshold is lower than the old one. Similarly, there exists some $(\tau_1(g'), \mathcal{R}(g'))$ that sets $\text{PPV}_{g'} = \text{PPV}_g$ and it is higher than the old one. By the monotonicity of PPV and NPV, we know that this process will not cause non-overlapping thresholds to become overlapping.

The ultimate effect of these steps was to reduce the number of overlapping thresholds by at least one. We can repeat this process up to $2|\mathcal{G}|$ times until none of the thresholds overlap.

Notice that this classifier is not one of the ones from Proposition 4.1 - if we did not have to correct for “overlapping,” then this is true by assumption, and if we did do the correction process, then $\tau_0(g) = \tau_1(g)$ for at least one g , and by assumption we had $|\text{Supp}(\hat{p})| \geq 2$.

Thus, we have created valid a post-processor $\hat{D}_{(\tau_0, \tau_1, \mathcal{R}_0, \mathcal{R}_1)}$ that equalizes PPV and NPV for all groups simultaneously and is not one of the ones in Proposition 4.1, proving the claim. \square

The following example demonstrates that it is sometimes possible to equalize PPV, NPV, FPR, and FNR using deferrals, but without equalizing the APs themselves:

Example 4.1 (Equalizing PPV, NPV, cFPR, and cFNR with Thresholds). *This example is presented with continuous support $[0, 1]$ for simplicity. Consider two APs, one for group g_1 and one for g_2 . Let the AP for g_1 be uniform (with density given by the line $\hat{P}(s) = 1$), and let the AP for group g_2 have density given by the parabola $\hat{P}(s) = 6s(1 - s)$, as shown in Figure 5.*

Consider the post-processor $\hat{D}_{(\tau_0, \tau_1)}^{\text{soft}}$.⁴ Let $\tau_0(g_1) = \tau_0(g_1) = 0.5$, let $\tau_0(g_2) = \frac{1}{6}(5 - \sqrt{7})$ and let $\tau_1(g_2) = 1 - \frac{1}{6}(5 - \sqrt{7})$ as shown in Figure 5.

The PPV and NPV of both groups is $\frac{3}{4}$, and the cFPR and cFNR of both is $\frac{1}{4}$, thus equalizing all four values.

This example is somewhat unsatisfactory because the base rates are equal in the two groups. We did not find a similar example without equal base rates.




Figure 5: This threshold post-processor equalizes PPV, NPV, cFPR, and cFNR as described in Example 4.1.

4.2 Equalizing APs with deferrals

As with Claim 3.4, equalizing the conditional APs between groups renders trivial the task of downstream decision-making subject to equality of PPV, NPV, cFPR, and cFNR. Importantly, unlike in Section 3.5, equalizing the conditional APs between groups does not require the groups to have equal base rates, greatly increasing the applicability of this approach.

Claim 4.1. *If the conditional APs are equal for two groups, then PPV, NPV, cFPR, and cFNR are equalized (or simultaneously undefined) by any hard deferring post-processor \hat{D} satisfying (1) group blindness and (2) $\hat{D}(\perp, g) = \perp$ ($\forall g$).*

The additional condition—that \hat{D} defers on input \perp —is necessary: if \hat{D} output 1 on all inputs (even on \perp), then PPV would remain unequal as long as the base rates differed. The proof is similar to the proof of Claim 3.4.

⁴In the case where the distributions are continuous, \mathcal{R}_0 and \mathcal{R}_1 are meaningless because $\Pr[s = \tau_0(g)] = \Pr[s = \tau_1(g)] = 0 \forall s$.

Proof of Claim 4.1. We prove only that PPV is equalized; the remaining properties may be proved similarly. Let $\hat{Y} = \hat{D} \circ \hat{S}$. All probabilities below are over $X_g \sim \mathcal{X}_g$, and the coins of \hat{S} and \hat{D} .

$$\begin{aligned} \text{PPV}_{\hat{Y},g} &= \Pr[Y(X_g) = 1 \mid \hat{Y}(X_g, g) = 1] \\ &= \frac{\Pr[Y(X_g) = 1 \mid \hat{S}(X_g) \neq \perp]}{\Pr[\hat{Y}(X_g, g) = 1 \mid \hat{S}(X_g) \neq \perp]} \\ &\quad \cdot \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \left(\frac{\Pr[Y(X_g) = 1 \mid \hat{S}(X_g) = s] \cdot \Pr[\hat{S}(X_g) = s \mid \hat{S}(X_g) \neq \perp]}{\Pr[Y(X_g) = 1 \mid \hat{S}(X_g) \neq \perp]} \right. \\ &\quad \left. \cdot \Pr[\hat{D}(s, g) = 1 \mid \hat{S}(X_g) = s, Y(X_g) = 1] \right) \end{aligned}$$

Each factor in this product is equal across groups by the assumptions. Namely, $\Pr[Y(X_g) = 1 \mid \hat{S}(X_g) \neq \perp]$ is equalized by calibration and equalized APs; $\Pr[\hat{Y}(X_g, g) = 1 \mid \hat{S}(X_g) \neq \perp]$ by group blindness and equalized APs; $\Pr[\hat{D}(s, g) = 1 \mid \hat{S}(X_g) = s, Y(X_g) = 1]$ by group blindness; $\Pr[Y(X_g) = 1 \mid \hat{S}(X_g) = s]$ by calibration; and finally $\Pr[\hat{S}(X_g) = s \mid \hat{S}(X_g) \neq \perp]$ by equalized APs. \square

Deferrals are a powerful tool for manipulating, and thereby equalizing, conditional APs. Consider a function $Q : (s, g) \mapsto [0, 1]$

$$\hat{D}_Q^{\text{soft}}(s, g) = \begin{cases} \perp & \text{if } s = \perp \\ \perp \text{ w.p. } Q(s, g), \text{ else } s & \text{otherwise} \end{cases}$$

If \hat{S} is a calibrated classifier, the soft deferring classifier $\hat{S}' := \hat{D}_Q^{\text{soft}} \circ \hat{S}$ is still calibrated. For a group g , let $\hat{\mathcal{P}}_g$ be the AP of \hat{S} and $\hat{\mathcal{P}}'_g$ be the AP of \hat{S}' . There is a simple graphical intuition for the shape of $\hat{\mathcal{P}}'_g$, as shown in Figure 6. More formally,

$$\hat{\mathcal{P}}'_g(s) = \frac{\hat{\mathcal{P}}_g(s)(1 - Q(s, g))}{1 - \Delta} \quad (7)$$

where $\Delta := \Pr[\hat{S}'(X_g) = \perp \mid \hat{S}(X_g) \neq \perp] = \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \hat{\mathcal{P}}_g(s)Q(s, g)$.

By appropriate choice of Q , any conditional AP can be transformed into almost any other conditional AP.

Theorem 4.1. *Let $\hat{\mathcal{P}}_g$ be a conditional AP of a soft classifier \hat{S} on group g , and let $\hat{\mathcal{P}}^*$ be any probability mass function such that $\text{Supp}(\hat{\mathcal{P}}^*) \subseteq \text{Supp}(\hat{\mathcal{P}}_g)$. Then there exists Q for which the calibrated AP $\hat{\mathcal{P}}'_g$ of $\hat{D}_Q^{\text{soft}} \circ \hat{S}$ is equal to $\hat{\mathcal{P}}^*$.*

Proof of Theorem 4.1. Let $\Delta = 1 - \min_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \frac{\hat{\mathcal{P}}_g(s)}{\hat{\mathcal{P}}^*(s)}$. For all $s \in \text{Supp}(\hat{\mathcal{P}}_g)$, let

$$Q(s, g) = 1 - \frac{\hat{\mathcal{P}}^*(s)}{\hat{\mathcal{P}}_g(s)} \cdot (1 - \Delta).$$

Observe that

$$\begin{aligned} \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \hat{\mathcal{P}}_g(s) Q(s, g) &= \sum_{s \in \text{Supp}(\hat{\mathcal{P}}_g)} \left(\hat{\mathcal{P}}_g(s) - (1 - \Delta) \hat{\mathcal{P}}^*(s) \right) \\ &= \Delta \end{aligned}$$

and hence Δ is defined as in Equation 7, where we used the fact that $\sum \hat{\mathcal{P}}_g(s) = \sum \hat{\mathcal{P}}^*(s) = 1$ in the last line. Plugging into the earlier formula for $\hat{\mathcal{P}}'_g$ (Equation 7) completes the proof. \square

Together, Theorem 4.1 and Claim 4.1 suggest a general framework for using deferrals to post-process a soft, possibly deferring classifier \hat{S} which is groupwise calibrated into a hard deferring classifier which simultaneously equalizes PPV, NPV, cFPR, and cFNR across groups, as follows.

For each $g \in \mathcal{G}$, let $\hat{\mathcal{P}}_g$ be the conditional AP of \hat{S} for group g . Let $\hat{\mathcal{P}}^*$ be any conditional AP such that $\text{Supp}(\hat{\mathcal{P}}^*) \subseteq \cap_{g \in \mathcal{G}} \text{Supp}(\hat{\mathcal{P}}_g)$. Use Theorem 4.1 to equalize the conditional APs for all groups $g \in \mathcal{G}$. Then use any hard post-processor \hat{D} satisfying the requirements of Claim 4.1 to make the ultimate deferring hard classifier. This method is shown in Figure 6.

This framework allows for enormous flexibility in the choice of both $\hat{\mathcal{P}}^*$ and \hat{D} , even when considering just two groups g_1 and g_2 . In Figure 10, we illustrate the first step of the framework on a COMPAS dataset using $\min\{\hat{\mathcal{P}}_{g_1}, \hat{\mathcal{P}}_{g_2}\}$ as $\hat{\mathcal{P}}^*$, where g_1 is African-Americans and g_2 is Caucasians. In Figures 8 and 9 in Section 5, we also use $\hat{\mathcal{P}}_{g_1}$ and $\hat{\mathcal{P}}_{g_2}$ as $\hat{\mathcal{P}}^*$.




Figure 6: Choosing deferrals cleverly allows transforming one AP into another (conditional) AP. In this example, the solid orange line is the original AP $\hat{\mathcal{P}}_g = \Pr[\hat{S}(X_g) = s]$. By deferring at the rates indicated by the shaded region, the resulting conditional AP $\hat{\mathcal{P}}'_g = \Pr[\hat{S}(X_g) = s \mid \hat{S}(X_g) \neq \perp]$ is represented by the dark blue line. The area of the shaded region is Δ .

One can design $\hat{\mathcal{P}}^*$ to achieve additional goals. For example, the choice $\hat{\mathcal{P}}^* = \min\{\hat{\mathcal{P}}_{g_1}, \hat{\mathcal{P}}_{g_2}\}$ results in *equal deferral rate* across each group (equal to the total variation distance between the two initial conditional APs). The framework can be further expanded by combining deferrals with other methods for manipulating conditional APs, including the mass-averaging discussed in Section 3.5. A better understanding of these techniques is left for future work.

5 Experiments on COMPAS Data

We test our methodology on the Broward County data made publicly available by ProPublica [ALMK16]. This data set contains the recidivism risk decile scores given by the COMPAS tool, 2-year recidivism outcomes, and a number of demographic and crime-related variables on individuals who were scored in 2013 and 2014. We restrict our attention to the subset of defendants whose race is recorded as African-American or Caucasian. These will form the two groups with respect to which we wish to examine different fairness criteria. After applying the same data pre-processing and filtering as reported in the ProPublica analysis, we are left with a data set on $n = 5278$ individuals, of whom 3175 are African-American and 2103 are Caucasian.

Indeed, it has been shown that the COMPAS scoring mechanism is an approximately calibrated soft classifier with 10 possible outcomes. We note here that the distribution of the COMPAS scores differs significantly across the two groups. In particular, the scores for Caucasians are more evenly distributed as opposed to the skewed distribution seen with African-Americans.

5.1 Thresholding with Deferrals

We first ran our two-threshold post-processing mechanism (Section 4.1) and obtained a binary decision algorithm with deferrals which equalizes both PPV and NPV across Caucasians and African-Americans (See Figure 7). For simplicity we avoid using randomization for members within a particular decile score and instead settle for approximate equalization of PPV and NPV. We observe that the percent of deferrals in total is smaller than 20% of the decisions to be made which shows that a fairly large number of the defendants can be classified without having to defer to a downstream decision maker.




Figure 7: Two thresholds are applied to each AP for the COMPAS data from 2016, (approximately) equalizing PPV and NPV. In the left plot we show the thresholds for the African American group, and in the right plot, we show the thresholds for the Caucasian group.

The thresholds suffer from the issue highlighted in Example 3.1, demonstrating that blindly equalizing PPV and NPV using thresholds can be problematic. Namely, the resulting deferring

hard classifier is much stricter for African-Americans than for Caucasians.

Next we look at our post-processing mechanisms to equalize all four quantities PPV, NPR, FPR, and NPR using deferrals (Section 4.2). As was noted earlier in the paper, equalizing the APs of the two groups post-deferral achieves the goal of equalizing all four of the above quantities. We implement two methods for doing so.

5.2 Converting one AP into Another

In the first method, decisions are deferred only on one group so as to convert its AP into that of the other group. First, we consider deferring only on Caucasians to convert their AP into that of African-Americans (Figure 9); next, decisions are deferred only for African Americans (Figure 8).




Figure 8: Our conditional AP equalization method applied to COMPAS data from 2016. We use deferrals to create a conditional AP for African Americans that matches the AP for Caucasians.




Figure 9: We create a conditional AP for Caucasians using the AP for African Americans. Notice the difference in the rates and distributions of deferrals between the two methods.

5.3 Equalizing APs

Alternately we have a second method where decisions are deferred for an equal fraction of Caucasians and of African Americans (Figure 10). This fraction is precisely equal to the statistical (*total variation*) distance between the distributions of scores produced by the soft classifier on the two groups.

We observe several striking phenomena on the COMPAS data set. First, using the method of deferring only on African-Americans we defer on roughly 36% of the total decisions. This number goes down to roughly 25% when we defer only on Caucasians. This seems to suggest as a general heuristic to try and use deferrals on the group with smaller size. The total deferral fraction is also roughly 25% when we defer on an equal fraction of Caucasians and African-Americans.

Second, for all three methods that equalize the score distributions, deferrals happen more on the “extremes”, namely on individuals with respect to which the classifier had relatively high confidence (either close to 0 or close to 1). This stands in sharp contrast to how the two-threshold method (Figure 7) distributes its deferrals—they occur only in the middle of the distribution (namely for




Figure 10: A version of our conditional AP equalization method applied to COMPAS data from 2016. The AP for African Americans and the AP for Caucasians are converted into a conditional AP that has the same distribution as the pointwise minimum of the two APs. Notice that the total deferral rate is equalized between the two groups (this is equal to precisely the total variation distance between the two APs), but the distribution of deferrals across scores is not.

elements for which the classifier is “unsure”). More work is needed to better understand the full space of deferral strategies.

6 Models of Deferring

Whether or not a classifier is thought of as promoting fairness depends on the context; this is true for both deferring and non-deferring classifiers. In addition to the myriad considerations present for non-deferring classifiers, deferring classifiers and downstream decision makers introduce some additional axes for consideration.

Cost to the individual: Even though it is not intended to be a final decision, a deferral may impose burdensome costs to an individual being classified. It may mean that a defendant remains in jail while additional hearings are scheduled, that invasive and expensive medical tests are ordered, or that continued investigation engenders social stigma. These costs may not be borne equally by all individuals, and may depend on their group membership, their true type, or other factors. For example, a delay in granting a loan to an applicant may overly burden poorer applicants, even those very likely to repay.

Cost to the decider: Allowing deferrals might make the decision process more cost-effective: Given that in most cases making a determination is cheap, one may now invest more in the deferred cases. For instance, a team of trained moderators might be hired to manually review content

on which an automated content filter defers, or an expensive investigation might be required to adjudicate insurance claims that are not cut-and-dry.

Accuracy of downstream decision One reason to defer is to introduce a delay that will allow for a more accurate decision. Thus the usefulness of allowing a classifier to defer depends on the accuracy of the downstream decision maker. Additional medical tests might allow for highly accurate diagnoses. But a judge deciding bail will be prone to a variety of errors and biases.

“Fairness” of downstream decision (and of composed classifier) Similar to the above, the fairness of the downstream decision maker (however one wants to interpret that) will impact our interpretation of the deferring classifier. Here one should take into account also the “procedural” aspect of the two-step evaluation; here it is important that the downstream classifier will be deemed as “more fair” and “more knowledgeable” than the first stage. Exploring fairness criteria for systems of deferring classifiers and downstream decision-makers, e.g. as done in [BKN⁺17] did for non-deferring classifiers, is an interesting direction for future work.

Frequency of decisions In many settings, the deferring classifier is a fast, automated test (e.g., automated risk assessment) while the downstream decision maker is a slow, manual process (e.g., parole board). However, we anticipate situations in which there may be repeated deferring classifiers chained together which comprise the complete decision making pipeline. For example, a doctor might have a sequence of diagnostic tests at her disposal as needed, or a bank might allow many rounds of appeal for loan applications, but with lengthy delays. Some applications might even permit hundreds or thousands of near-continuous deferring classifiers. As an example, consider a live video streaming platform that passively monitors streams for inappropriate content in real time. The automated passive monitor might decide the content is inappropriate, and shut it down; appropriate, and continue passive monitoring; or suspicious (by deferring), and begin active monitoring by devoting more computing resources or bringing in a human moderator.

6.1 Technical implications of deferral model

The contextual considerations discussed above directly impact the appropriate application of a deferring classifier and its goals. An obvious goal is to minimize the overall rate of deferrals while maintaining the best possible FPR, FNR, PPV, and NPV for the classifier conditioned on not deferring, and without considering the distribution of deferrals. However, one might desire very different properties from the distribution of deferrals in different contexts. The deferrals may be distributed differently among individuals with different true type, group membership, or soft-classifier scores, while the burden imposed by deferrals and errors may differ greatly between different populations.

In a medical diagnosis scenario, a false negative (i.e., failing to diagnose a disease) may have serious consequences, and deferring to run additional non-invasive and inexpensive additional tests may be generally acceptable. On the other hand, an insurance provider may prefer to minimize expensive investigations by paying out more false claims.

The context may also affect the way one defines the deferral analogues of FPR and FNR. While calibration, PPV, and NPV apply directly to deferring classifiers, it is not clear how best to generalize the definitions of error rates. For example, consider false positive rate: by Definition 2.5, the false

positive rate of a non-deferring hard classifier \hat{Y} for a group g is $\text{FPR}_{\hat{Y},g} = \Pr[\hat{Y}(X_g) = 1 \mid Y(X_g) = 0]$.

The approach we take in Section 4 is to condition on not deferring (Definition 4.1). A deferring classifier \hat{Y} that output 1 on half of true negative instances (within a g) would have conditional false positive rate as low as 0.5 (if it never output \perp on true negatives) or as high as 1 (if it never output 0 on true negatives). The conditional false positive rate is agnostic towards the downstream decision maker. It codifies no value judgements as to whether a deferral is desirable or undesirable as an individual nor whether deferrals ultimately result in accurate or inaccurate decisions. This is, itself, a value judgement.

A second approach is to leave the original definition unchanged. The same deferring hard classifier as above would have unconditional false positive rate 0.5. This would be true regardless of whether \hat{Y} output 0 or \perp on the other half of true negative instances. We call this the *unconditional false positive rate*. The unconditional false positive rate effectively categorizes deferrals as correct outputs. This may be appropriate if the downstream decision maker has very high accuracy. If, for example, a doctor orders an additional, more accurate diagnostic test in response to a deferral, the unconditional false positive rate might be appropriate.

Finally, a third approach is to base our measure of inaccuracy on true negatives instead of false positives, a reverse of the above.

Just as in the case of non-deferring classifiers, the relationships among these contrasting group statistics, their meaningfulness in different settings, and their application in different settings are not well understood and deserve further study.

References

- [ALMK16] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias: Theres software used across the country to predict future criminals. and its biased against blacks. *propublica*, may 23, 2016, 2016.
- [BKN⁺17] Amanda Bower, Sarah N Kitchen, Laura Niss, Martin J Strauss, Alexander Vargas, and Suresh Venkatasubramanian. Fair pipelines. *arXiv preprint arXiv:1707.00391*, 2017.
- [Cho17] Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *CoRR*, abs/1703.00056, 2017.
- [CPF⁺17] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. *CoRR*, abs/1701.08230, 2017.
- [DHP⁺11] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. Fairness through awareness. *CoRR*, abs/1104.3913, 2011.
- [DI18] Cynthia Dwork and Christina Ilvento. Fairness under composition. *CoRR*, abs/1806.06122, 2018.
- [HKRR17] Úrsula Hébert-Johnson, Michael P. Kim, Omer Reingold, and Guy N. Rothblum. Calibration for the (computationally-identifiable) masses. *CoRR*, abs/1711.08513, 2017.
- [HPS16] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems 29: Annual Conference on*

Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain, pages 3315–3323, 2016.

- [KMR16] Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. *CoRR*, abs/1609.05807, 2016.
- [KNRW18] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2564–2572, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- [MP17] Andrew Morgan and Rafael Pass. Paradoxes in fair computer-aided decision making. *CoRR*, abs/1711.11066, 2017.
- [MPZ17] David Madras, Toniann Pitassi, and Richard S. Zemel. Predict responsibly: Increasing fairness by learning to defer. *CoRR*, abs/1711.06664, 2017.
- [PRW⁺17] 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.

A Results for NPV and Continuous, Full Support APs

A.1 Results for Negative Predictive Value (NPV)

In Section 3, we proved limitations on the ability of post-processors to equalize PPV given a distribution on calibrated scores with finite support. For completeness, in this section, we give the statements of the analogous limitations for equalizing NPV and for continuous probability density functions with full support $[0, 1]$.

We start with the analogous statement of Proposition 3.3 for NPV instead of PPV.

Proposition A.1. *Fix two disjoint groups g_1 and g_2 with respective base rates BR_1 and BR_2 such that $BR_1 \neq BR_2$. Then there exists a soft-valued classifier \hat{S} that is groupwise calibrated, but for which there is no post-processor $\hat{D} : [0, 1] \times \mathcal{G} \rightarrow \{0, 1\}$ such that $\hat{D} \circ \hat{S}$ equalizes NPV, unless $\Pr[\hat{D}(BR_i, g_i) = 0] = 0$ for $i = 1$ or 2 .*

The nontriviality condition ensures that the NPV is well-defined on both groups (which can be compared to the nontriviality condition in Proposition 3.3, which ensures that the PPV is well-defined on both groups). The proof is essentially identical to the proof of Proposition 3.3: the fraction of predicted 0's in group g_i that are true 0's is $1 - BR_i$, as the post-processor has no other information by which to make its decision, and hence the NPVs are unequal due to the differing base rates.

We now proceed to give the NPV analogs of our results on threshold post-processors in Section 3. We start with the analogous statement for Proposition 3.4 - that there is a class of simple group-blind threshold post-processors that equalizes the NPV across groups.

Proposition A.2. *For every nice groupwise calibrated soft classifier \hat{S} and for every group-blind threshold post-processor $\hat{D}_{(\tau, \mathcal{R})}$ such that $\tau(g) = \min(\text{Supp}(\hat{P}_g))$ and $\mathcal{R}(g) < 1$ for all g , the composed classifier $\hat{D} \circ \hat{S}$ equalizes NPVs across all groups.*

The existence of the threshold post-processors in Proposition A.2 follows from the assumed finiteness of the range of the soft classifier. In the case where the range of the soft classifier is infinite, such post-processors may not exist (as there may be no minimum element of the support). The proof is once again analogous to the proof of Proposition 3.4: these classifiers only ever output 0 on the minimum element of the support of \hat{P}_g , and hence the NPV is simply 1 minus the smallest element of the support for each group.

However, much like the case for PPV, other group-blind threshold post-processors cannot possibly equalize NPV.

Proposition A.3. *There exists a groupwise-calibrated soft classifier with a nice AP for which no non-trivial group blind threshold post-processor, other than the ones in Proposition 3.4, can equalize NPV across groups.*

The example of the groupwise-calibrated soft classifier is the exact same one that shows the statement for PPV. Indeed, we can see this in the following way. We can make the example in the proof of Proposition A.3 have equal base rates across groups, in which case it follows from Proposition 3.5 due to Claim 3.3.

When we turn to threshold post-processors that are not group-blind, we again get analogous results for NPV.

Proposition A.4. *Let \mathcal{G} be a set of groups. For any soft classifier \hat{S} with a nice AP $\hat{\mathcal{P}}$ such that \hat{S} is groupwise-calibrated over \mathcal{G} and $|\text{Supp}(\hat{\mathcal{P}}_g)| \geq 2$ for all $g \in \mathcal{G}$, there exists a (non-group-blind), non-trivial threshold post-processor $\hat{D}_{(\tau, \mathcal{R})}$ that is not one of the group blind post-processors in Proposition 3.4, such that the hard classifier $\hat{Y} = \hat{D}_{(\tau, \mathcal{R})} \circ \hat{S}$ equalizes NPV across \mathcal{G} .*

This holds even if we require that the NPV of all the groups is an arbitrary value in $(s_{\min}, \min_i \text{BR}_{g_i})$, where $\min_i \text{BR}_{g_i}$ is the minimum base rate among the groups and s_{\min} is the minimum score in the support of $\hat{\mathcal{P}}_{g_i}$.⁵

Again the proof of this follows via the same kind of continuity argument as we used to prove Proposition 3.6. By definition, each group has base rate at least $\min_i \text{BR}_{g_i}$, and so if the post-processor always says 0 on some group i , then $\text{NPV}_{g_i} = \text{BR}_{g_i} \geq \min_i \text{BR}_{g_i}$. Hence, for each group, we can start with the always-0 classifier and slide down the threshold until the desired NPV is reached.

Finally, the socially unsatisfying example also generalizes to NPV. A privileged group will have higher scores than a disadvantaged group in general, and hence if they are given the same threshold, the NPV will be lower on the privileged group. To rectify this, the threshold for the disadvantaged group will have to move higher, to decrease the NPV. But then, the disadvantaged group is being subjected to a harsher standard.

Finally we note that Claim 3.2 also can be written with NPV instead of PPV, where the proof follows from using the characterization of NPV given in Proposition 3.1:

Claim A.1 (Continuity of NPV). *Fix a soft classifier \hat{S} and a corresponding AP $\hat{\mathcal{P}}$, as well as a group g . Let \hat{D}_1 and \hat{D}_2 be two post-processing algorithms. Let $\hat{\mathcal{P}}_{g, \hat{D}_1}$ be the expected conditional AP that results from starting with the AP $\hat{\mathcal{P}}_g$ over scores in group g and then conditioning on the scores that \hat{D}_1 sends to 0, and define $\hat{\mathcal{P}}_{g, \hat{D}_2}$ similarly. If $d_{\text{TV}}(\hat{\mathcal{P}}_{g, \hat{D}_1}, \hat{\mathcal{P}}_{g, \hat{D}_2}) < \epsilon$, then $|\text{NPV}_{g, \hat{D}_1 \circ \hat{S}} - \text{NPV}_{g, \hat{D}_2 \circ \hat{S}}| < O(\epsilon)$.*

We omit the proof of Claim A.1, which resembles the proof of Claim 3.2 and follows from Proposition 3.1.

A.2 Results for Continuous, Full Support APs

In this section, we briefly address how to extend our results on thresholds in Section 3 to the setting where every AP $\hat{\mathcal{P}}$ is a continuous probability distribution with $\text{Supp}(\hat{\mathcal{P}}_g) = [0, 1]$ for all $g \in G$ - that is, the support equals the entire interval $[0, 1]$ for each group $g \in G$. Note that this automatically makes $\hat{\mathcal{P}}$ a “nice” AP, and hence rules out the general counterexample we came up with in Proposition 3.3. For the purposes of this section, call such an AP that 1) $\hat{\mathcal{P}}_g$ is a continuous probability density function for every $g \in G$ and 2) $\text{Supp}(\hat{\mathcal{P}}_g) = [0, 1]$ for all $g \in G$ a *very nice* AP. As the name suggests, we can extend the remaining results in Section 3 to the setting of very nice AP. We give the results for equalizing PPV as done in Section 3: extending these results to equalizing NPV in the setting of continuous, full support AP can be accomplished by combining the statements here with the modifications described in Section A.1.

First, we note that a threshold post-processor can be described much more easily in the continuous setting than in the setting where the AP has finite support on each group. Indeed, in the setting where $\hat{\mathcal{P}}_g$ is a continuous density function for all $g \in G$, the post-processor can truly be a

⁵For the case where the support of $\hat{\mathcal{P}}_{g_i}$ is infinite, s_{\min} should be the infimum of scores.

threshold, with no question of how to classify the score that is exactly equal to the threshold τ . This is because the score τ has probability 0 under the density $\hat{\mathcal{P}}_g$.

Hence, Proposition 3.4 has no true analog in this setting. This follows because the maximum element of the support in this case is 1, and a threshold at $\tau = 1$ sends every score (outside of a measure 0 set) to 0.

This allows us to strengthen Proposition 3.5 accordingly.

Proposition A.5. *There exists a groupwise-calibrated soft classifier with a very nice AP for which no non-trivial group blind threshold post-processor can equalize PPV across groups.*

This follows from a nearly identical stochastic domination argument to the one used for Proposition 3.5 - in fact, the natural generalization of the distributions given for the proof of Proposition 3.5 to the continuous and full-support setting can be used in this proof.

A non group blind threshold can still always equalize PPV for very nice APs.

Proposition A.6. *Let \mathcal{G} be a set of groups. For any soft classifier \hat{S} with a very nice AP $\hat{\mathcal{P}}$ such that \hat{S} is groupwise-calibrated over \mathcal{G} , then there exists a non group blind, non-trivial threshold post-processor such that the hard classifier equalizes PPV across \mathcal{G} .*

This holds even if we require that the PPV of all the groups is equal to an arbitrary value in $(\max_i \text{BR}_{g_i}, 1)$, where $\max_i \text{BR}_{g_i}$ is the maximum base rate among the groups $g_i \in \mathcal{G}$.

This follows from the same continuity approach as the proof of Proposition 3.6, by sliding the threshold continuously down from 1 until the PPV reaches the desired value $v \in (\max_i \text{BR}_{g_i}, 1)$.

The socially unsatisfying example generalizes in the natural way - Example 3.1 consists of one group having a monotonically increasing PMF and another one having a monotonically decreasing PMF. We can skip the discretization step in the definition of these PMFs and have them be continuous PDFs, and the example still goes through.

We cannot equalize PPV and NPV simultaneously in general, just like in the finite support case (Proposition 3.7).

Proposition A.7. *Fix groups g_1 and g_2 . There exists a soft classifier \hat{S} with a very nice AP $\hat{\mathcal{P}}$ such that no threshold post-processor can simultaneously equalize PPV and NPV between groups g_1 and g_2 .*

The proof we give of Proposition 3.7 for finite support naturally generalizes to this case - in fact, we can simply use the same proof but without discretizing the probability distributions. The necessary lemmas about monotonicity of PPV and being able to express the base rate as a convex combination of PPV and NPV still hold.

It is unsurprising that the result in Section 4 on using thresholds *with deferrals* to equalize PPV and NPV also goes through for very nice AP.

Proposition A.8. *Let \hat{S} be groupwise calibrated for the n groups g_1, \dots, g_n , and suppose that \hat{S} has a very nice AP. Then there exists a nontrivial threshold decision rule such that the hard classifier $\hat{Y} = \hat{D} \circ \hat{S}$ equalizes PPV and NPV for \mathcal{G} .*

Again, the explanation is very similar to the one for Proposition A.6. PPV and NPV change continuously when we slide the respective thresholds, and unlike the case without deferrals, we can change the PPV *without* changing the NPV, by keeping the “0” threshold still and sliding the “1” threshold (and deferring in the middle). Hence, we can simply continuously slide the two thresholds on each group until they reached the desired values.