Blind Pareto Fairness and Subgroup Robustness

Natalia Martinez^{*1} Martin Bertran^{*1} Afroditi Papadaki² Miguel Rodrigues² Guillermo Sapiro¹

Abstract

Much of the work in the field of group fairness addresses disparities between predefined groups based on protected features such as gender, age. and race, which need to be available at train, and often also at test, time. These approaches are static and retrospective, since algorithms designed to protect groups identified a priori cannot anticipate and protect the needs of different at-risk groups in the future. In this work we analyze the space of solutions for worst-case fairness beyond demographics, and propose Blind Pareto Fairness (BPF), a method that leverages no-regret dynamics to recover a fair minimax classifier that reduces worst-case risk of any potential subgroup of sufficient size, and guarantees that the remaining population receives the best possible level of service. BPF addresses fairness beyond demographics, that is, it does not rely on predefined notions of at-risk groups, neither at train nor at test time. Our experimental results show that the proposed framework improves worst-case risk in multiple standard datasets, while simultaneously providing better levels of service for the remaining population, in comparison to competing methods.

1. Introduction

A large body of literature has shown that machine learning (ML) algorithms trained to maximize average performance on existing datasets may present discriminatory behaviour across pre-defined demographic groups (Barocas & Selbst, 2016; Hajian et al., 2016), meaning that segments of the overall population are measurably under-served by the ML model. This has sparked interest in the study on why these disparities arise, and on how they can be addressed (Mitchell

Proceedings of the 37th International Conference on Machine Learning, Vienna, Austria, PMLR 108, 2020. Copyright 2020 by the author(s).

et al., 2018; Chouldechova & Roth, 2018; Barocas et al., 2019). One popular notion is group fairness, where the algorithm has access to a set of predefined demographic groups during training, and the goal is to learn a model that satisfies a certain notion of fairness across these groups (e.g., statistical parity, equality of opportunity) (Dwork et al., 2012; Hardt et al., 2016); this is usually achieved by adding a constraint to the standard optimization objective. It has been shown that optimality may be in conflict with some notions of fairness (e.g., the optimal risk is different across groups) (Kaplow & Shavell, 1999; Chen et al., 2018), and perfect fairness can, in general, only be achieved by degrading the performance on the benefited groups without improving the disadvantaged ones. This conflicts with notions of no-harm fairness such as in (Ustun et al., 2019), which are appropriate where quality of service is paramount. Notions such as minimax fairness, commonly known as Rawlsian max-min fairness from an utility maximization perspective (Rawls, 2001; 2009), combined with Pareto efficiency, naturally address this no-harm concern (Martinez et al., 2020; Diana et al., 2020).

Recent works study fairness in ML when no information about the protected demographics is available, for example, due to privacy or legal regulations (Kallus et al., 2019). This is an important research direction and has been identified as a major industry concern (Veale & Binns, 2017; Holstein et al., 2019), since many applications and datasets in ML currently lack demographic records. We therefore study the problem of building minimax Pareto fair algorithms beyond demographics, meaning that not only we lack group membership records but also have no prior knowledge about the demographics to be considered (e.g., any subset of the population can be a valid protected group, see computationally identifiable groups (Hébert-Johnson et al., 2018)). This has the advantage of making the model robust to any potential demographic even if they are unknown at the time of design, or change through time; it is also efficient, since the model offers the best level of service to all the remaining (i.e., non-critical) population.

Main Contributions. We analyze *subgroup robustness*, where a model is minimax fair w.r.t. any group of sufficient size, regardless of any preconceived notion of protected groups; we also adhere to the notion of *no-harm* fairness by

^{*}Equal contribution ¹Duke University, Durham, NC, USA ²University College London, London, UK. Correspondence to: Natalia Martinez <natalia.martinez@duke.edu>, Martin Bertran <martin.bertran@duke.edu>.

requiring our minimax model to be Pareto efficient (Mas-Colell et al., 1995) by providing the best level of service to the remaining population. A model with these characteristics has a performance guarantee even for unidentified protected classes.

We show that being *subgroup robust* w.r.t. an unknown number of groups, where no individual group is smaller than a certain size, is mathematically equivalent in terms of worst group performance to solving a simplified two-group problem, where the population is divided into high and low risk groups, thereby providing a clear means to design "universal" minimax fair ML models. We further show the critical role of the minimum group size by proving that, for standard classification losses (cross-entropy and Brier score), there is a limit to the smallest group size we can consider before the solution degenerates to a trivial, uniform classifier. We additionally study the cost of blind subgroup robustness when compared to learning a model that is minimax fair w.r.t. predefined demographics.

We then propose *Blind Pareto Fairness* (BPF), a simple learning procedure that leverages recent methods in noregret dynamics (Chen et al., 2017) to solve *subgroup robustness* subject to a user-defined minimum subgroup size. Our method is provably convergent and can be used on classification and regression tasks. We experimentally evaluate our method on a variety of standard ML datasets and show that it effectively reduces worst-case risk and compares favourably with previous works in the area. Although our work is motivated by fairness, subgroup robustness has applications beyond this important problem, see for example (Sohoni et al., 2020a; Duchi et al., 2020).

2. Related Work

A body of work has addressed fairness without explicit demographics by using proxy variables to impute the protected population labels (Elliott et al., 2008; Gupta et al., 2018; Zhang, 2018). These methods contrast with our assumptions by relying on a preconceived notion on what the protected demographics are (i.e., the protected demographics are known, but unobserved), since prior knowledge is needed to design useful proxy variables. Moreover, it has been reported that these approaches can exacerbate disparities by introducing undesired bias (Chen et al., 2019; Kallus et al., 2019); aiming to be fair by inferring protected attributes may be in conflict with privacy or anonymity concerns. These works might need re-training if new protected classes are identified, since a model trained under these conditions may be considerably harmful on an unknown population. This phenomena further supports the value of blind subgroup robustness.

Individual fairness (Dwork et al., 2012) provides guarantees

beyond protected attributes, but requires predefined similarity functions which may be hard or infeasible to design for real-world tasks. The works of (Hébert-Johnson et al., 2017; Kearns et al., 2018) address fairness w.r.t. subgroups based solely on input features, and while these works greatly extend the scope of the protected demographics, they still rely on labeled protected features for guidance. The work of (Sohoni et al., 2020b) partitions the input space to address robust accuracy. We note that partitions, based only on the input space of the model do not modify the solution of risk-based Pareto optimal models, since the optimal classifier for any input value remains unchanged (i.e., there is no conflict between objectives for any value of the input space, see Theorem 4.1 in (Martinez et al., 2020)). In our work we consider subgroups based on both outcome and all input features, which broadens the scope to all conceivable subgroups based on the information available to the trainer. For many risk-based measures of utility, such as crossentropy, Brier score, or ℓ_2 regression loss, the optimal classifier can be expressed as a function of the conditional output probability p(Y|X), X being the input (features) and Y the output. In particular, if we only consider groups that introduce covariate shift (i.e., p(X|A) varies across different values of the group membership A) but do not change the conditional target distribution (p(Y|X,A) = p(Y|X)for all A), then the set of Pareto classifiers only contain one element and the Pareto curve degenerates to the utopia point. By specifically taking outcomes into account in our partition function, we allow for robustness to perturbations on the conditional distribution p(Y|X,A).

There are two recent approaches that are the closest to our objective (protecting unknown and unobserved demographics). One is distributionally robust optimization (DRO) (Hashimoto et al., 2018; Duchi et al., 2020), where the goal is to achieve minimax fairness for unknown populations of sufficient size. Similar to our work, they minimize the risk of the worst-case group for the worst-case group partition, they use results from distributional robustness that focus the attention of the model exclusively on the high-risk samples (i.e., their model reduces the tail of the risk distribution). However, they do not explicitly account for Pareto efficiency, meaning that their solution may be sub-optimal on the population segment that lies below their high-risk threshold, doing unnecessary harm. The other recent method that tackles the minimax objective is adversarially reweighted learning (ARL) (Lahoti et al., 2020), where the model is trained to reduce a positive linear combination of the sample errors, these weighting coefficients are proposed by an adversary (implemented as a neural network), with the goal of maximizing the weighted empirical error. This method focuses on computationally identifiable subgroups, meaning that

¹In this work we consider "partition" and "subgroup" interchangeable.

they can be characterized by a function $f: \mathcal{X} \times \mathcal{Y} \to [0,1]$, (Hébert-Johnson et al., 2018). However, they do not provide an optimality guarantee on the adversary, nor do they pose a constraint on the computationally identifiable subgroups. Our theoretical results indicate that adding an easily interpretable group size constraint on this subgroup is necessary so that the worst-case partition does not yield a trival, uniform classifier for the optimal adversary; this observation is validated in experimental results.

3. Problem Formulation

3.1. Minimax Fairness

We first consider the supervised group fairness classification scenario (Barocas et al., 2019), where we have access to a dataset $\mathcal{D} = \{(x_i, y_i, a_i)\}_{i=1}^n \sim p(X, Y, A)^{\otimes n}$ containing n i.i.d. triplets. Here $X \in \mathcal{X}$ denotes the input features, $Y \in \mathcal{Y}$ the categorical target variable, and $A \in \mathcal{A}$ group membership. We consider a classifier $h \in \mathcal{H}$ belonging to an hypothesis class \mathcal{H} whose goal is to predict Y from X, $h: \mathcal{X} \to \Delta^{|\mathcal{Y}|-1}$; note that h(X) can take any value in the simplex and is readily interpretable as a distribution over labels Y. Given a loss function $\ell: \Delta^{|\mathcal{Y}|-1} \times \Delta^{|\mathcal{Y}|-1} \to \mathbb{R}^+$, fairness is considered in the context of a Multi-Objective Optimization Problem (MOOP), where the objective is to learn a classifier that minimizes the conditional group risks $\mathbf{r}(h) = \{r_a(h)\}_{a \in \mathcal{A}}$,

$$\min_{h \in \mathcal{H}} (r_1(h), ..., r_{|\mathcal{A}|}(h)),
r_a(h) = E_{X,Y|A=a}[\ell(h(X), Y)].$$
(1)

The solution to this MOOP may not be unique (e.g., the optimal classifier of different groups differs), and therefore there is a set of optimal (Pareto) solutions that can be achieved. It is possible that none of these Pareto solutions satisfy some group fairness criteria (e.g., equality of risk), meaning that achieving perfect fairness comes at the cost of optimality (Kaplow & Shavell, 1999; Bertsimas et al., 2011). In this work we do not to compromise optimality, meaning that we do not degrade the performance of a low-risk group if it does not directly benefit another, and consider a minimax fairness approach (Rawls, 2001; 2009), where the goal is to find a properly Pareto optimal classifier that minimizes the worst-case group risk,

$$\min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{a \in \mathcal{A}} r_a(h). \tag{2}$$

 $\mathcal{H}_{P_{\mathcal{A}}}$ represents the set of properly Pareto optimal classifiers in \mathcal{H} given a group set \mathcal{A} as defined next.

Definition 3.1. (Geoffrion, 1968) An hypothesis $h^* \in \mathcal{H}$ is properly Pareto optimal if h^* is Pareto efficient, meanining that $\not \exists h' \in \mathcal{H} : \mathbf{r}(h') \prec \mathbf{r}(h), ^2$ and there exists a finite

$${}^{2}\mathbf{r}(h') \prec \mathbf{r}(h)$$
 if $r_a(h') \leq r_a(h) \forall a \land \exists a' : r_{a'}(h') < r_{a'}(h)$

 $M>0: \forall i,h\in\mathcal{A}\times\mathcal{H}$ satisfying $r_i(h)< r_i(h^*)\ \exists j\in\mathcal{A}$ such that $r_j(h^*)< r_j(h)$ and $\big(r_i(h^*)-r_i(h)\big)/\big(r_j(h)-r_j(h^*)\big)\leq M$. Given a partition set \mathcal{A} we denote the set of properly Pareto hypothesis in \mathcal{H} as \mathcal{H}_{P_A} .

Note that Definition 3.1 establishes that there is no other model in the hypothesis class with unbounded risk tradeoffs between groups whose associated group risks are uniformly better for all groups. If the loss function is convex w.r.t. the model h, and the hypothesis class \mathcal{H} is a convex set,³ a linear weighting problem on the conditional group risks $(\min_{h \in \mathcal{H}} \sum \mu_a r_a(h); \sum \mu_a = 1, \mu_a > 0)$ characterizes all of the properly Pareto solutions (Geoffrion, 1968). Solving Problem 2 over the properly Pareto set is equivalent to finding the weighting coefficients such that a classifier with the minimum worst-case group risk is the solution (Geoffrion, 1968; Chen et al., 2017; Martinez et al., 2020).

3.2. Blind Pareto Fairness

In this work we consider a more challenging problem, namely $Blind\ Pareto\ Fairness$ (BPF), where the group variable A and the conditional distribution p(A|X,Y) are completely unknown (not just unobserved), even at training time. Here the goal is to learn a model that has the best performance on the worst-group risk of the worst partition density p(A|X,Y) ("sensitive" group assignment), subject to a group size constraint $(p(A=a) \geq \rho, \forall a)$. We formulate the following new problem,

$$R^* = \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y)} \max_{a \in \mathcal{A}} r_a(h).$$

$$s.t.p(A) \succeq \rho^4$$
(3)

Here R^* is the minimum worst group error achieved for the worst partition density with known number of partitions $|\mathcal{A}|$. The partition size constraint plays a key role on the solution to Problem 3, as we will show below. Problem 3 is undetermined in the sense that it admits several worst partition densities and classifiers for $|\mathcal{A}| > 2$. Fortunately, it is possible to show that the minimum worst group error R^* in Problem 3 is the same as the one achieved if we were to consider an alternative formulation where a variable $A \in \{0,1\}$ represents the worst-group risk membership, this is shown in Lemma 3.1. This makes the study of the binary problem attractive when we wish to minimize the number of assumptions we make on the protected groups. Here the objective becomes

³Meaning that for any $h, h' \in \mathcal{H}$ and $\lambda \in [0, 1]$, exists $h_{\lambda} \in \mathcal{H}: h_{\lambda}(x) = \lambda h(x) + (1 - \lambda)h'(x) \forall x \in \mathcal{X}.$

 $^{{}^4}p(A) \succeq \rho \text{ if } p(A=a) \ge \rho \ \forall a \in \mathcal{A}$

⁵Attempting to be minimax optimal w.r.t. ρ smaller than some value will result in a random classifier, the minimax risk of the partition is directly dependent on ρ .

$$R^* = \begin{array}{ccc} \min & \max & r_a(h), \\ h \in \mathcal{H}_{P_A} & a \in \{0, 1\} \\ & p(A|X, Y) \\ & s.t. \ p(A) \succeq \rho \end{array}$$
(4)

with $h^*, p^*(A|X, Y)$ achieving this solution. We overload the notation \mathcal{H}_{P_A} in the context of Problem 4 to refer to the Pareto set for a binary group distribution. Figure 1 shows an example of the risks for the worst and best partitions at different sizes ρ achieved with a method that optimizes for the worst case partition, like DRO or our proposed BPF (the latter shows better performance on the remaining partition owing to the Pareto constraint), versus deploying a baseline model that minimizes the empirical risk. Lemma 3.1 shows that the minimum worst risk R^* is the same for problems 3 and 4, hence, we focus our analysis on the latter throughout the text. There are two main advantages of working with the binary problem, the first is that finding the worst partition $p(A \mid X, Y)$ for a given h is straightforward when |A| =2. The second is that, in general, we may not know the number of groups we wish to be fair to, and this equivalence shows that it is sufficient to specify the minimum size a group must have before it is considered for the purposes of minimax fairness. Moreover, restricting the minimum size of the partitions to be considered is an interpretable way of constraining the adversary.

Lemma 3.1. Given an hypothesis class \mathcal{H} and a finite alphabet, $\mathcal{A}: |\mathcal{A}| \geq 2$, problems 3 and 4 have the same minimum worst-group risk solution R^* if $\rho \leq \frac{1}{|\mathcal{A}|}$.

A question that arises from Problem 4 is how the optimal classifier and partition function depend on the partition size. In Lemma 3.2, we show the existence of a critical size ρ^* for standard classification losses (cross-entropy and Brier score) whereby solving Problem 4 for partitions smaller than ρ^* leads to a uniformly random classifier. This result shows that attempting to be minimax fair w.r.t. arbitrarily small group sizes yields a trivial classifier with limited practical utility. Therefore, if one where to consider an adversary without any capacity restriction this would be it's optimal solution, something that we corroborate empirically.

Lemma 3.2. Given Problem 4 with $p(Y|X) > 0 \ \forall X, Y,^6$ and let the classification loss be cross-entropy or Brier score. Let

$$\bar{h}(X):\bar{h}_i(X)=\frac{1}{|\mathcal{Y}|}\forall X,\forall i\in\{0,...,|\mathcal{Y}|-1\},$$

be the uniform classifier, and let $\bar{h} \in \mathcal{H}$. There exists a critical partition size

$$\rho^* = |\mathcal{Y}| E_X[\min_y p(y \mid X)] \le 1$$

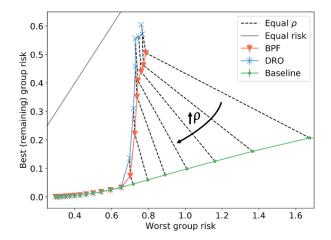


Figure 1. Worst and best crossentropy risks achieved by DRO, BPF and a baseline classifier for varying minimum size partitions (ρ) on the UCI Adult dataset. The tradeoffs shown for DRO and BPF correspond to model pairs that were optimized for a specific ρ value (i.e., different points in the tradeoff curve correspond to different BPF and DRO classifiers). Lower worst group risks correspond to larger group sizes (ρ); results for the same ρ value are connected with a dashed line. We observe that BPF is able to achieve the same worst-case group performance that DRO achieves, but with better results on the non-critical partition owing to its Pareto optimality constraint, this is especially apparent on small group sizes. The baseline classifier suffers significantly larger worst group errors for small partition sizes.

such that solutions to Problem 4, $\forall \rho \leq \rho^*$, are $h^* = \bar{h}$ and $R^* = \bar{R} = \begin{cases} \log |\mathcal{Y}| & \text{if } \ell = \ell_{CE} \\ \frac{|\mathcal{Y}|-1}{|\mathcal{Y}|} & \text{if } \ell = \ell_{BS} \end{cases}$.

That is, any partitions smaller than ρ^* yield the uniform classifier with constant risk \bar{R} .

It is straightforward to prove that R^* is non-increasing with ρ (see Supplementary Material A.1). A natural question that arises is what is the additional cost in optimality we pay if we apply subgroup robustness instead of optimizing for a known partition. Lemma 3.3 provides an upper bound for the cost of blind fairness, showing that it is at most the difference between R^* and the risk of the baseline model, and can be zero (no cost) if the known group happens to be the worst case partition for the dataset. Moreover, the upper bound decreases with larger group size, and is in no scenario larger than the difference between the risk of the uniform classifier and the baseline classifier for BE and CE losses.

Lemma 3.3. Given a distribution p(X,Y) and any predefined partition group p(A'|X,Y) with $A' \in \mathcal{A}'$, $|\mathcal{A}'|$ finite. Let $\hat{h}, \hat{R} = \{\arg\} \min_{h \in \mathcal{H}} \max_{a' \in \mathcal{A}'} r_{a'}(h)$ be the minimax fair solution for this partition and its corresponding minimax risk. Let h^* and R^* be the classifier and risks that solve Problem

⁶This restriction can be lifted and a similar result holds, see Supplementary Material for details.

4 with $\rho = \min_{a'} p(a')$. Then the price of minimax fairness can be upper bounded by

$$\max_{a' \in \mathcal{A}'} r_{a'}(h^*) - \hat{R} \le R^* - \min_{h \in \mathcal{H}} r(h).^7$$
 (5)

In the following section, we provide a practical algorithm that asymptotically, solves Problem 4 and yields a classifier that both minimizes the worst-group risk and is also Paretoefficient w.r.t. the remaining population. (All proofs are presented in the Supplementary Material A.1.)

4. Optimization

In order to develop our optimization approach, we begin by showing that for group sizes $\rho \leq \frac{1}{2}$, we can drop the innermost max operator in Problem 4, and we only need to consider the risk for the a=1 partition on a partition of size exactly equal to ρ (i.e., $p(A=1)=\rho$). This is presented in the following lemma, proved in the Supplementary Material.

Lemma 4.1. Given Problem 4 with minimum group size $\rho \leq \frac{1}{2}$, the following problems are value equivalent:

$$R^{\mathbf{I}} = \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y)} \max_{a \in \{0,1\}} r_a(h),$$

$$s.t.p(A) \succeq \rho$$

$$R^{\mathbf{II}} = \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y)} r_1(h),$$

$$s.t. p(A=1) = \rho$$

$$R^{\mathbf{I}} = R^{\mathbf{II}}.$$
(6)

We note that the second problem in Eq. 6 is in itself a valuable optimization problem for $\rho>1/2$, since it equates to efficiently minimizing the risk of the at-risk majority. The Pareto optimality constraint in this scenario is easy to enforce if the base loss function $\ell(h(x),y)$ is both bounded (i.e., $\ell(h(x),y) \leq C \ \forall x,y,h \in \mathcal{X} \times \mathcal{Y} \times \mathcal{H}$), and strictly convex w.r.t. model h. The result is shown in the following lemma.

Lemma 4.2. Given the problem on the right hand side of Eq. 6, a convex hypothesis class \mathcal{H} , and a bounded loss function $0 \le \ell(h(x), y) \le C \ \forall x, y, h \in \mathcal{X} \times \mathcal{Y} \times \mathcal{H}$ that is strictly convex w.r.t its first input h(x), the following problems are equivalent:

$$\mathcal{H}^{I} = \{\arg\} \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y)} r_{1}(h)$$

$$s.t. \ p(A=1) = \rho$$

$$\mathcal{H}^{II} = \{\arg\} \min_{h \in \mathcal{H}} \sup_{h \in \mathcal{H}} \sup_{p(A|X,Y)} r_{1}(h),$$

$$p(A|X,Y)$$

$$s.t. \ p(A=1) = \rho$$

$$p(A=1|X,Y) > 0, \ \forall X, Y$$

$$\mathcal{H}^{I} = \mathcal{H}^{II},$$
(7)

where we explicitly add {arg} to both sides of the equivalence to indicate that these problems are also equivalent in terms of the models h that achieve these minimax solutions. The BS loss satisfies both conditions in Lemma 4.2, CE loss also satisfies these conditions if we restrict the hypothesis set such that $h \in \mathcal{H} \subseteq \{h: h_i(x) \geq \zeta > 0 \forall i \in [\mathcal{Y}], x \in \mathcal{X}\}$ (i.e., the classifier assigns a minimum label probability for all values). The ℓ_2 regression loss over a bounded set also satisfies these conditions. To deal with the supremum constraint on the distribution in practice, we slightly limit adversary capacity and ensure $p(A=1|X,Y) \geq \epsilon > 0 \,\forall X,Y \in \mathcal{X} \times \mathcal{Y}$.

We solve Problem 7 using no-regret dynamics (Freund & Schapire, 1999), the solution is the Nash equilibrium of a two-player zero-sum game, where one player, the adversary, iteratively proposes partition distributions p(A|X,Y), the modeler then responds near optimally with a model h, and incurs loss $r_1(h)$. Based on the history of losses, the adversary iteratively refines its proposed partition function into the worst-case partition.

To solve the above problem with parameter $\epsilon>0$, we leverage the theoretical results presented in (Chen et al., 2017) for improper robust optimization of infinite loss sets with oracles. We first present the results in terms of a finite dataset with n samples; assume that both players have access to $\{x_i,y_i\}_{i=1}^n \sim P(X,Y)^{\otimes n}$, and let $t\in\{0,\ldots,T\}$ indicate the current round of the zero-sum game. In each round t, the modeler produces a classifier h^t and the adversary proposes an empirical distribution of p(A|X,Y), denoted as $\alpha^t=\{\alpha_i^t\}_{i=1}^n\in\mathcal{U}_{\epsilon,\rho}$ such that $\mathcal{U}_{\epsilon,\rho}=\{\alpha:\alpha_i\in[\epsilon,1],\sum_i\frac{\alpha_i}{n}=\rho\}$, where ρ is the minimum partition size. The empirical risk (cost) of round t is $L^t=L(h^t,\alpha^t)$, with

$$L(h, \boldsymbol{\alpha}) = \frac{\sum_{i=1}^{n} \alpha_i \ell(h(x_i), y_i)}{n\rho}.$$
 (8)

In order to find the Nash equilibrium of this game, we use projected gradient ascent on the adversary, while the modeler uses approximate best response with a Bayesian oracle $h^t = M(\alpha^t)$. In particular, we use a variant proposed in (Chen et al., 2017) for robust non-convex optimization. Algorithm 1 shows the proposed approach.

The proposed Algorithm 1 is an instantiation of Algorithm 3 in (Chen et al., 2017) for oracle efficient improper robust optimization with infinite loss sets. To implement the projection operator $\prod_{\mathcal{U}_{\epsilon,\rho}}(\cdot)$, we observe that the set $\mathcal{U}_{\epsilon,\rho}$ can be described as the intersection of the $[\epsilon,1]$ hypercube and the $\langle \alpha, 1 \rangle = n\rho$ plane, where both of these sets have known and exact projection functions individually. Dykstra's projection algorithm (Boyle & Dykstra, 1986) is a provably convergent best approximation method that can implement this operation. This projection update converges in linear

 $^{^{7}}r(h) = E_{X,Y}[\ell(h(X),Y)].$

⁸The algorithm is iterative, we prove convergence to the optimal solution with the number of iterations.

Algorithm 1 Blind Pareto Fairness

Require: Inputs: Dataset $\{(x_i,y_i)\}_{i=1}^n$, partition size ρ **Require:** Hyper-parameters: rounds T, parameter η , adversary boundary coefficient $\epsilon>0$, γ -approximate Bayesian solver $\mathrm{M}(\cdot)\simeq\arg\min_{h\in\mathcal{H}}L(h,\cdot)$

Initialize $\alpha^0 = \hat{\alpha} = \{\rho\}_{i=1}^n$

Initialize classifier and loss

$$h^0 = \mathbf{M}(\hat{\boldsymbol{\alpha}}), L^0 = L(h^0, \hat{\boldsymbol{\alpha}})$$

for round $t = 1, \dots, T$ do

Adversary updates partition function by projected gradient ascent:

$$\alpha^{t} \leftarrow \alpha^{t-1} + \eta \nabla_{\alpha} L(h^{t}, \hat{\alpha}) = \alpha^{t-1} + \eta \frac{\ell(h^{t}, y)}{n\rho}$$
$$\hat{\alpha} \leftarrow \prod_{\mathcal{U}_{\epsilon, \rho}} (\alpha^{t}), \mathcal{U}_{\epsilon, \rho} = \{\alpha : \alpha_{i} \in [\epsilon, 1], \sum_{i} \frac{\alpha_{i}}{n} = \rho\}$$

Solver approximately solves for the current partition: $h^t \leftarrow M(\hat{\alpha})$

end for

Return: Classifier h^T

time and only requires access to the last risk evaluation on each sample during the training stage, making it a scalable and lightweight addition to the standard supervised training scenario. We can then immediately leverage the results in Theorem 7 in (Chen et al., 2017) to show that the algorithm converges (proof presented in the Supplementary Material).

Lemma 4.3. Consider the setting of Algorithm 1, with parameter $\epsilon > 0$, and $\eta = \max_{\boldsymbol{\alpha} \in \mathcal{U}_{\epsilon,\rho}} \frac{||\boldsymbol{\alpha}||_2}{\sqrt{2T}} \leq \sqrt{\frac{n\rho}{2T}}$ with $\mathcal{U}_{\epsilon,\rho} = \{\boldsymbol{\alpha} : \alpha_i \in [\epsilon,1], \sum_i \frac{\alpha_i}{n} = \rho\}$, and L a 1-Lipschitz function w.r.t. $\boldsymbol{\alpha}$, let P be a uniform distribution over the set of models $\{h^1,\ldots,h^T\}$, and let R^* be the minimax solution to the loss presented in Eq. 8. Then we have

$$\max_{\boldsymbol{\alpha} \in \mathcal{U}_{\epsilon,\rho}} \mathbb{E}_{h \sim P} L(h, \boldsymbol{\alpha}) \leq \gamma R^* + \sqrt{\frac{2n\rho}{T}}.$$

As in (Chen et al., 2017), we use h^t instead of the ensemble $\{h^1, \ldots, h^T\}$. We use stochastic gradient descent (SGD) as our γ -approximate Bayesian oracle, in practice, we alternate a single epoch of SGD with the adversary update for simplicity. We note that the 1-Lipschitz constraint can be relaxed to any G-Lipschitz function by working through the no regrets guarantees for projected gradient descent of G-Lipschitz functions in the proof provided in (Chen et al., 2017).

Figure 2 shows how the performance of the recovered classifiers (trained for a given partition size ρ) is optimal for its own partition size, but sub-optimal for other ρ values. These curves give a better picture on how risks are being traded off across samples.

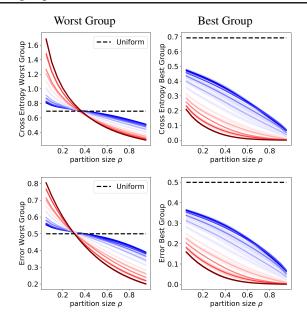


Figure 2. Crossentropy and error rates of BPF classifiers for varying (evaluation) partition sizes ρ evaluated on the UCI Adult dataset. Each individual curve denotes the performance of a unique BPF classifier, trained for a particular ρ , across a range of evaluation partition sizes. The uniform classifier is shown for reference. We observe how performance is traded off between low and high risk groups for varying partition sizes, in particular, smaller training partition sizes yield uniformly worse performance on the best group, no matter the evaluation partition size. Conversely, optimality on the worst group is dependent on how matched the train and test partition sizes are.

5. Experimental Results

We experimentally validate our methods and theoretical results on a variety of standard datasets, we compare performance against DRO (Hashimoto et al., 2018), ARL (Lahoti et al., 2020), and a baseline classifier (empirical risk minimization). We show the trade-offs of each method on their worst group and the remaining population. As presented below (see Figure 3), the baseline method performs best on the low-risk population, but it suffers from large, fat tails in terms of loss distribution. We also show how both DRO and BPF empirically achieve the theoretical results laid in Lemma 3.2, with BPF having better results on the low-risk population than DRO, owing to the Pareto optimality constraint. Moreover, if the adversary's network on the ARL framework is given enough capacity, it degrades to the uniform (trivial) solution presented in Lemma 3.2 since it does not control for other restrictions on the partitions learned (e.g., group size). We also show that the performance of ARL can vary with adversarial network capacity (e.g., depth and width of the network). However, translating this to the effective size of the worst case partition being optimized is

not easy to interpret or evaluate beforehand.

Datasets. We used four standard fairness datasets for comparison. The UCI Adult dataset (Dua & Graff, 2017) which contains 48,000 records of individual's annual income as well as 13 other attributes, including race, gender, relationship status, and education level. The target task is income prediction (binary, indicating above or below 50K). The Law School dataset (Wightman, 1998) contains law school admission data used to predict successful bar exam candidates; in our examples we limit ourselves to UGPA and LSAT scores as input covariates. The COMPAS dataset (Barenstein, 2019) which contains the criminal history, serving time, and demographic information such as sex, age, and race of convicted criminals. The goal is prediction of recidivism per individual.9 Lastly we used the MIMIC-III dataset, which consists of clinical records collected from adult ICU patients at the Beth Israel Deaconess Medical Center (Johnson et al., 2016). The objective is predicting patient mortality from clinical notes. We analyze clinical notes acquired during the first 48 hours of ICU admission following the pre-processing methodology in (Chen et al., 2018), ICU stays under 48 hours and discharge notes are excluded from the analysis. Tf-idf statistics on the 1,000 most frequent words in clinical notes are taken as input features.

Setup and Results. We train BPF for 16 minimum group sizes $\rho = \{0.05, 0.10, \dots, 0.8\}$, we report cross-entropy loss and error rate, 10 on the worst partition of the dataset (i.e., average over the worst $100 \times \rho\%$ samples based on cross-entropy loss), the values for the remaining low risk group is also reported to evaluate optimality. DRO models were trained on 20 equispaced values of their threshold parameter $(\eta \in [0, 1])$. For ARL, we tried four configurations for their adversarial network (adversary with 1 or 2 hidden layers with 256 or 512 units each), we additionally evaluated the same setup when the ARL adversary has access to the learned features of the classifier network, this setup still falls within the computationally identifiable scenario in their work, but offers a more challenging adversary, we denote these latter experiments as (ARL HC). The classifier architecture for BPF, ARL, and DRO was standardized to a single-layer MLP with 512 hidden units. In all cases we use cross-entropy loss and same input data. Results correspond to the best hyper-parameter for each group size; mean and standard deviations are computed using 5-fold crossvalidation. Further implementation details are provided in Supplementary Material.

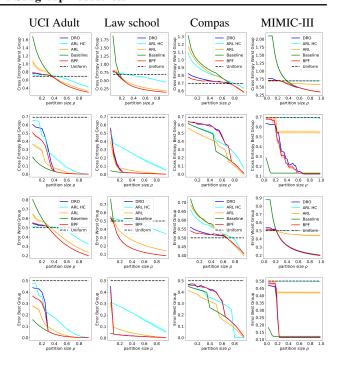


Figure 3. Cross-entropy (CE) and error rate (Error) metrics on best and worst groups as a function of group size for BPF, DRO, ARL, and baseline classifiers; results for very high capacity adversarial networks for ARL (ARL HC) are also shown, random classifier shown for reference. Results are provided for UCI adult, law school, COMPAS, and MIMIC-III datasets. Cross-entropy of both ARL and baseline classifiers for the worst group are very large for small group size, DRO and BPF both approximate the theoretical result shown in Lemma 3.2. The main experimental difference between DRO and the proposed BPF is that BPF exhibits better results on the best group partition than DRO for the same level of worst group performance, owing to the Pareto restriction on the BPF classifier resulting in no-unnecessary-harm for any group (see also Figure 1). Error results largely mimic the observations on the cross-entropy metric.

Figure 3 shows the performance of the best and worst groups across partition sizes. Both DRO and BPF recover results close to the random classifier for the smaller group sizes, which aligns with the results shown in Lemma 3.2, that is, below a certain partition size (e.g., $\rho \simeq 0.3$ for adult dataset) the average cross-entropy of the worst group is the risk of the uniform classifier ($\log 2$). We observe that the performance of ARL seems to be dataset dependent, but is generally able to reduce worst-case risks w.r.t. the baseline classifier for small partition sizes. The high capacity ARL (ARL HC) behaves as expected in most cases, producing results much more closely aligned with the uniform classifier, this supports our theoretical results stating that an unconstrained adversary should converge to the uniform classifier. Altogether, BPF obtains better results for the low risk group than DRO for the same worst group risk, owing

⁹This dataset is the source of extensive and very legitimate controversy in the fairness community, and is here used for benchmarking only.

 $^{^{10}}$ Error rate is computed on the randomized classifier $\hat{Y} \sim h(X)$.

to the former's optimality constraint. The magnitude of this improvement is dataset dependent.

Although none of the compared models address disparities along predefined populations, we can nonetheless observe how each classifier performs on these groups. Table 1 shows accuracy conditioned on different demographics for each competing method on the Adult dataset. We observe that the different methods achieve results close to the uniform classifier for small partition sizes as expected. In many cases, the results for BPF are better than ARL and DRO for each protected attribute (at same ρ value). We also observe that on several minorities, the BPF model provides the best utility values out of all the competing methods, BPF is also the best model at preserving worst group performance.

Table 2 shows how target labels and predefined sensitive groups are represented in the high risk group identified by BPF. We observe that, for low partition sizes, outcomes are balanced across groups (in concordance with Lemma 3.2). As the partition size increases, the composition of the high risk group becomes more similar to the base distribution. Similar results to tables 1 and 2 are provided in Supplementary Material A.2 for the remaining datasets.

| Method/ρ | White | Black | Asian-PacI | Other |
|------------------------------|-------------------------|-------------------------|-------------------------|-------------------------|
| Prop(%) | 85.6% | 9.6% | 2.9% | 1.8% |
| ARL.15 DRO .15 BPF .15 | 52.1% 55.5% 62.3% | 51.9% 57.0% 67.5% | 52.1% 55.5% 62.0% | 53.1% 57.1% 65.5% |
| ARL.25 DRO .25 BPF .25 | 60.3% $57.4%$ $65.0%$ | 65.9% $59.4%$ $69.7%$ | 59.5% $57.3%$ $64.1%$ | 63.8% $59.5%$ $68.5%$ |
| ARL .4 DRO .4 BPF .4 | 68.6% 80.2% 80.2% | 77.5% 88.0% 88.0% | 68.6% 79.1% 79.1% | 74.3% 86.9% 86.9% |

Table 1. Accuracy across demographic partitions (groups given no special consideration by the algorithms) in the Adult dataset for ARL, DRO and BPF models for varying partition sizes.

6. Discussion

In this work we formulate and analyze subgroup robustness, particularly in the context of fairness without demographics or labels. Our goal is to recover a model that minimizes the risk of the worst-case partition of the input data subject to a minimum size constraint, while we additionally constrain this model to be Pareto efficient w.r.t. the low-risk population as well. This means that we are optimizing for the worst unknown subgroup without causing unnecessary harm on the rest of the data. We show that it is possible to protect high risk groups without explicit knowledge of their number or structure, only the size of the smallest one, and that there is a minimum partition size under which the ran-

| Group | Prop(%) | BPF .15 | BPF .25 | BPF .4 | | | |
|-------------------------------------------------|---------|---------|---------|--------|--|--|--|
| Proportion on Worst Partition, Ethnicity/Income | | | | | | | |
| White/0 | 64.2 | 43.8 | 44.8 | 53.5 | | | |
| White/1 | 21.4 | 45.3 | 44.9 | 35.6 | | | |
| Black/0 | 8.5 | 2.7 | 3.0 | 4.6 | | | |
| Black/1 | 1.1 | 2.9 | 2.5 | 1.8 | | | |
| Asian-PacI/0 | 2.1 | 1.8 | 1.7 | 1.9 | | | |
| Asian-PacI/1 | 0.8 | 2.0 | 1.8 | 1.4 | | | |
| Other/0 | 1.5 | 0.4 | 0.4 | 0.8 | | | |
| Other/1 | 0.3 | 1.0 | 0.8 | 0.5 | | | |

Table 2. Demographic composition of worst groups as a function of minimum partition size on the Adult dataset. BPF homogenizes outcomes across partitions and protected attributes. For larger group sizes, the demographics of the partition approach that of the baseline population.

dom classifier is the only minimax option for cross-entropy and Brier score losses.

We propose BPF, an algorithm that provably converges to a properly Pareto minimax solution, it requires minimal modifications to the standard learning pipeline of a standard model, and can scale easily to large datasets. Our results on a variety of standard fairness datasets show that this approach reduces worst-case risk as expected, and produces better models than competing methods for the low-risk population, thereby avoiding unnecessary harm. It (computationally) identifies high risk samples and is easy to interpret since the user can control the optimal adversary through the use of a target worst partition size.

If a policymaker has a desired risk tradeoff instead of a target group size, we can search for the smallest partition size achieving this tradeoff using the proposed BPF; this now guarantees that the recovered model can satisfy this risk tradeoff for the worst possible partition up to size ρ , and for any smaller partition size there exists a partition such that this tradeoff is violated.

Future work includes incorporating additional domainspecific constraints on the worst partition and developing an algorithm that combines BPF with knowledge about some subgroups that must be protected as well.

References

Barenstein, M. Propublica's compas data revisited. *arXiv* preprint arXiv:1906.04711, 2019.

Barocas, S. and Selbst, A. D. Big data's disparate impact. *Calif. L. Rev.*, 104:671, 2016.

Barocas, S., Hardt, M., and Narayanan, A. *Fairness and Machine Learning*. fairmlbook.org, 2019. http://www.fairmlbook.org.

- Bertsimas, D., Farias, V. F., and Trichakis, N. The price of fairness. *Operations Research*, 59(1):17–31, 2011.
- Boyle, J. P. and Dykstra, R. L. A method for finding projections onto the intersection of convex sets in hilbert spaces. In *Advances in order restricted statistical inference*, pp. 28–47. Springer, 1986.
- Chen, I., Johansson, F. D., and Sontag, D. Why is my classifier discriminatory? In *Advances in Neural Information Processing Systems*, pp. 3539–3550, 2018.
- Chen, J., Kallus, N., Mao, X., Svacha, G., and Udell, M. Fairness under unawareness: Assessing disparity when protected class is unobserved. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 339–348, 2019.
- Chen, R. S., Lucier, B., Singer, Y., and Syrgkanis, V. Robust optimization for non-convex objectives. In *Advances in Neural Information Processing Systems*, pp. 4705–4714, 2017.
- Chouldechova, A. and Roth, A. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810*, 2018.
- Diana, E., Gill, W., Kearns, M., Kenthapadi, K., and Roth, A. Convergent algorithms for (relaxed) minimax fairness. *arXiv preprint arXiv:2011.03108*, 2020.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL http://archive.ics.uci.edu/ml.
- Duchi, J., Hashimoto, T., and Namkoong, H. Distributionally robust losses for latent covariate mixtures. *arXiv* preprint arXiv:2007.13982, 2020.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pp. 214–226, 2012.
- Elliott, M. N., Fremont, A., Morrison, P. A., Pantoja, P., and Lurie, N. A new method for estimating race/ethnicity and associated disparities where administrative records lack self-reported race/ethnicity. *Health Services Research*, 43 (5p1):1722–1736, 2008.
- Freund, Y. and Schapire, R. E. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.
- Geoffrion, A. M. Proper efficiency and the theory of vector maximization. *Journal of Mathematical Analysis and Applications*, 22(3):618–630, 1968.
- Gupta, M., Cotter, A., Fard, M. M., and Wang, S. Proxy fairness. *arXiv preprint arXiv:1806.11212*, 2018.

- Hajian, S., Bonchi, F., and Castillo, C. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2125–2126, 2016.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Advances in Neural Informa*tion Processing Systems, pp. 3315–3323, 2016.
- Hashimoto, T. B., Srivastava, M., Namkoong, H., and Liang, P. Fairness without demographics in repeated loss minimization. arXiv preprint arXiv:1806.08010, 2018.
- Hébert-Johnson, U., Kim, M. P., Reingold, O., and Rothblum, G. N. Calibration for the (computationally-identifiable) masses. *arXiv preprint arXiv:1711.08513*, 2017.
- Hébert-Johnson, U., Kim, M., Reingold, O., and Rothblum, G. Multicalibration: Calibration for the (computationallyidentifiable) masses. In *International Conference on Machine Learning*, pp. 1939–1948. PMLR, 2018.
- Holstein, K., Wortman Vaughan, J., Daumé III, H., Dudik, M., and Wallach, H. Improving fairness in machine learning systems: What do industry practitioners need? In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, pp. 1–16, 2019.
- Johnson, A. E., Pollard, T. J., Shen, L., Li-wei, H. L., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., and Mark, R. G. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3:160035, 2016.
- Kallus, N., Mao, X., and Zhou, A. Assessing algorithmic fairness with unobserved protected class using data combination. *arXiv preprint arXiv:1906.00285*, 2019.
- Kaplow, L. and Shavell, S. The conflict between notions of fairness and the Pareto principle. *American Law and Economics Review*, 1(1):63–77, 1999.
- Kearns, M., Neel, S., Roth, A., and Wu, Z. S. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*, pp. 2564–2572, 2018.
- Lahoti, P., Beutel, A., Chen, J., Lee, K., Prost, F., Thain, N., Wang, X., and Chi, E. H. Fairness without demographics through adversarially reweighted learning. *arXiv preprint arXiv:2006.13114*, 2020.
- Martinez, N., Bertran, M., and Sapiro, G. Minimax Pareto fairness: A multi objective perspective. In *International Conference on Machine Learning*, 2020.

- Mas-Colell, A., Whinston, M. D., Green, J. R., et al. *Microeconomic Theory*, volume 1. Oxford University Press New York, 1995.
- Mitchell, S., Potash, E., Barocas, S., D'Amour, A., and Lum, K. Prediction-based decisions and fairness: A catalogue of choices, assumptions, and definitions. *arXiv preprint arXiv:1811.07867*, 2018.
- Rawls, J. Justice as Fairness: A Restatement. Harvard University Press, 2001.
- Rawls, J. A Theory of Justice. Harvard University Press, 2009.
- Sohoni, N., Dunnmon, J., Angus, G., Gu, A., , and Re, C. Addressing hidden stratification: Fine-grained robustness in coarse-grained classification problems, 2020a. http://hazyresearch.stanford.edu/hidden-stratification, July2020.
- Sohoni, N. S., Dunnmon, J. A., Angus, G., Gu, A., and Ré, C. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. *To Appear in Neural Information Processing Systems*, 2020b. http://stanford.edu/~nims/no_subclass_left_behind.pdf.
- Ustun, B., Liu, Y., and Parkes, D. Fairness without harm: Decoupled classifiers with preference guarantees. In *International Conference on Machine Learning*, pp. 6373–6382, 2019.
- Veale, M. and Binns, R. Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data & Society*, 4(2):2053951717743530, 2017.
- Wightman, L. F. LSAC National Longitudinal Bar Passage Study. LSAC Research Report Series. 1998.
- Zhang, Y. Assessing fair lending risks using race/ethnicity proxies. *Management Science*, 64(1):178–197, 2018.

A. Supplementary Material

A.1. Proofs

Lemma 3.1. Given an hypothesis class \mathcal{H} and a finite alphabet $\mathcal{A}: |\mathcal{A}| \geq 2$, under mild conditions, problems 3 and 4 have the same minimum worst-group risk solution R^* if $\rho \leq \frac{1}{|\mathcal{A}|}$.

Proof. For any $h \in \mathcal{H}$, let $l_h = \ell(h(X), Y)$ be the random variable associated with the loss distribution of h induced by the randomness of X, Y. Let $\hat{\ell}_{h,\rho} = F_{l_h}^{-1}(1-\rho)$ be the $100*(1-\rho)\%$ percentile of l_h , where $F_{l_h}^{-1}(\alpha) = \inf\{l \in \mathbb{R} : P(l_h \leq l) \geq \alpha\}$ is the inverse cdf of l_h . It is easy to observe that any distribution $p(A \mid X, Y), A \in \mathcal{A}$, that satisfies

$$p(A = a' \mid X, Y) = \begin{cases} 1 & \text{if } \ell(h(X), Y) > \hat{\ell}_{h, \rho}, \\ \alpha(X, Y) \in [0, 1] & \text{if } \ell(h(X), Y) = \hat{\ell}_{h, \rho}, \\ 0 & \text{if } \ell(h(X), Y) < \hat{\ell}_{h, \rho}, \end{cases}$$

$$p(A = a) \ge \rho, \ \forall a \in \mathcal{A},$$

$$p(A = a') = \rho, \ a' \in \mathcal{A}.$$

$$(9)$$

is a solution to

$$\max_{p(A|X,Y)} \max_{a \in \mathcal{A}} r_a(h),$$

s.t. $p(A) \succeq \rho$,

attaining the maximum risk at $r_{a'}(h)$. Here $\alpha(X,Y) \in [0,1]$ is any tie-breaking assignment such that $p(A=a') = \rho$ and $p(A=a) \geq \rho$. That is, the worst-case partition greedily assigns A=a' to all high loss samples until the budget $p(A=a') = \rho$ is satisfied, the tie-breaker assignment $\alpha(X,Y)$ simply indicates that for loss values exactly equal to $\hat{\ell}_{h,\rho}$, we can make any assignment we wish to as long as $p(A=a') = \rho$.

Furthermore, by applying the same reasoning as above, we observe that the simplified distribution $\hat{p}(A \mid X, Y), A \in \{0, 1\}, \ \hat{p}(A = 1 \mid X, Y) = p(A = a' \mid X, Y)$ is also a solution to

$$\max_{a \in \{0,1\}} r_a(h),
p(A|X,Y)
s.t. p(A) \succeq \rho,$$

with both achieving the same maximum risk. At this point we have proved the following equivalence:

$$\min_{h \in \mathcal{H}} \max_{p(A|X,Y)} \max_{a \in \mathcal{A}} r_a(h) = \min_{h \in \mathcal{H}} \max_{a \in \{0,1\}} r_a(h).$$

$$s.t. \ p(A) \succeq \rho, \qquad p(A|X,Y)$$

$$s.t. \ p(A) \succeq \rho,$$

If we further assume that \mathcal{H} does not contain models with unbounded risk trade-offs between groups, we want to prove that, **in terms of worst case risk**, minimizing over $h \in \mathcal{H}$ is equivalent to minimizing over its respective properly Pareto classifiers sets, $h \in \mathcal{H}_{P_A}$ for the left side of the equation and $h \in \mathcal{H}_{P_{A}=\{0,1\}}$ for the right side.

Looking at the left side equation, we note that for all $\bar{h} \in \mathcal{H}$ and $\bar{p}(A|X,Y): p(A) \succeq \rho$, we have a corresponding risk vector $\{r_a(\bar{h})\}_{a\in\mathcal{A}}$. Let $a'=\arg\max_a r_a(\bar{h})$ be the worst group; by the properties of Pareto optimality, we know that there exists a model \hat{h} such that

$$\hat{h} \in \mathcal{H}_{P_{\mathcal{A}}} : r_{a'}(\hat{h}) = r_{a'}(\bar{h}), r_{a}(\hat{h}) \le r_{a}(\bar{h}) \ \forall a \in \mathcal{A} \setminus \{a'\}.$$

That is, there exists a Pareto efficient model that achieves the same risk on a' but less or equal risk in all other coordinates (note that if $\bar{h} \in \mathcal{H}_{P_A}$ then $\bar{h} = \hat{h}$). Applying this property we observe that

$$\min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y)} \max_{a \in \mathcal{A}} r_a(h) = \min_{h \in \mathcal{H}} \max_{p(A|X,Y)} \max_{a \in \mathcal{A}} r_a(h).$$

$$s.t. \ p(A) \succeq \rho,$$

$$s.t. \ p(A) \succeq \rho,$$

Using similar reasoning, we have that

$$\min_{h \in \mathcal{H}_{P_{\mathcal{A}} = \{0,1\}}} \max_{a \in \{0,1\}} r_a(h) = \min_{h \in \mathcal{H}} \max_{a \in \{0,1\}} r_a(h),$$

$$p(A|X,Y) \qquad p(A|X,Y)$$

$$s.t. \ p(A) \succeq \rho, \qquad s.t. \ p(A) \succeq \rho,$$

and thus,

$$\min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y)} \max_{a \in \mathcal{A}} r_a(h) = \min_{h \in \mathcal{H}_{P_{\mathcal{A}=\{0,1\}}}} \max_{a \in \{0,1\}} r_a(h),$$

$$s.t. \ p(A) \succeq \rho, \qquad p(A|X,Y)$$

$$s.t. \ p(A) \succeq \rho,$$

We want to restate that the equalities are valid in terms of worst case risk, there may be minimax models $h \in \mathcal{H}$ that do not belong to the Pareto set $h \in \mathcal{H}_{P_A}$

Lemma 3.2. Given Problem 4 with $p(Y|X) > 0 \ \forall X, Y$, and let the classification loss be cross-entropy or Brier score. Let $\bar{h}(X) : \bar{h}_i(X) = \frac{1}{|\mathcal{Y}|} \forall X, \forall i \in \{0,...,|\mathcal{Y}|-1\}$ be the uniform classifier, and let $\bar{h} \in \mathcal{H}$.

There exists a critical partition size

$$\rho^* = |\mathcal{Y}| E_X[\min_y p(y \mid X)] \le 1$$

such that the solutions to Problem 4, $\forall \rho \leq \rho^*$, are

$$h^* = \bar{h},$$

$$R^* = \bar{R} = \begin{cases} \log |\mathcal{Y}| & \text{if } \ell = \ell_{CE} \\ \frac{|\mathcal{Y}| - 1}{|\mathcal{Y}|} & \text{if } \ell = \ell_{BS} \end{cases}.$$

That is, the solutions to all partitions with size smaller than ρ^* yield the uniform classifier with constant risk R.

Proof. This proof is done in three steps, first we provide an upper bound of the solution of Problem 4, we then show that we can design a (potentially nonexistent) partition density that achieves this upper bound, and finally, we derive conditions under which the previously identified partition is guaranteed to exist.

We first prove that for any distributions p(X, Y, A), it follows that

$$\min_{h \in \mathcal{H}} \max_{a \in \{0,1\}} r_a(h) \le \bar{R},$$

meaning that the solution to Problem 4 is upper bounded by the risk associated with the uniform classifier for cross-entropy and Brier score losses.

This is done by considering that for any distribution p(X,Y,A), the conditional risk of the uniform classifier $\bar{h}(X)$: $\bar{h}_i(X) = \frac{1}{|\mathcal{Y}|} \forall X, \forall i \in \{0,...,|\mathcal{Y}|-1\}$ is

$$E_{X,Y|A}[\ell(\bar{h}(X),Y)] = \bar{R} = \begin{cases} \log |\mathcal{Y}| & \text{if } \ell = \ell_{CE} \\ \frac{|\mathcal{Y}|-1}{|\mathcal{Y}|} & \text{if } \ell = \ell_{BS} \end{cases}, \ \forall p(X,Y,A),$$

Since $\bar{h} \in \mathcal{H}$, we have that $\min_{h \in \mathcal{H}} \max_{a \in \{0,1\}} r_a(h) \leq \bar{R} \ \forall p(X,Y,A)$.

Then we show that if we can design $p(A|X,Y): p(Y|X,A=1)=\frac{1}{|\mathcal{Y}|}\forall X,Y$ we have that, under this distribution, $\bar{h},\bar{R}=\{\arg\}\min_{h\in\mathcal{H}}\max_{a\in\{0,1\}}r_a(h)$, which is the upper bound identified above.

For this assume that we have $p(Y|X,A=1)=\frac{1}{|\mathcal{Y}|}\forall X,Y$, it then follows that $\min_{h\in\mathcal{H}}r_1(h)=r_1(\bar{h})=\bar{R}$ since \bar{h} is, by design, the optimal classifier for group a=1 and \bar{R} its best achievable risk. Then $r_{a=1}(h)\geq \bar{R}$ $\forall h$ and since $r_{a=1}(\bar{h})=r_{a=0}(\bar{h})$ it follows that

$$\bar{h}, \bar{R} = \{\arg\} \min_{h \in \mathcal{H}} \max_{a \in \{0,1\}} r_a(h).$$

Finally, we derive a necessary and sufficient condition for the existence of $p(A|X,Y): p(Y|X,A=1) = \frac{1}{|\mathcal{Y}|} \forall X,Y$. Since we need

$$p(A=1|X,Y) = \frac{1}{|Y|} \frac{p(A=1|X)}{p(Y|X)}$$

to be a well-defined distribution, the only degree of freedom available is $p(A=1 \mid X)$. Note that

$$p(Y|A=0,X) = \frac{p(Y|X)|\mathcal{Y}| - p(A=1|X)}{1 - p(A=1|X)} \frac{1}{|\mathcal{Y}|} \ge 0 \ \forall X, Y,$$

therefore $p(A=1|X) \leq |\mathcal{Y}|p(Y|X), \ \forall Y, X \rightarrow p(A=1|X) \leq |\mathcal{Y}| \min_{y \in \mathcal{Y}} p(Y=y|X)$ and therefore

$$p(A=1) \le E_X[|\mathcal{Y}| \min_{y \in \mathcal{Y}} p(y|X)] = \rho^*.$$

We also note that $\min_{y \in \mathcal{Y}} p(y|X) \leq \frac{1}{|\mathcal{Y}|}$, therefore $\rho^* \leq 1$

Note that the Lemma above can drop the hypothesis $p(Y|X) > 0 \ \forall X, Y$ by defining a new semi-uniform classifier $\bar{h}(X): \bar{h}_i(X) = \frac{\mathbb{1}(i \in Y(X))}{|\mathcal{Y}(X)|} \forall X, \forall i \in \{0,...,|\mathcal{Y}|-1\}$, where $\mathcal{Y}(X)$ indicates the subset of labels y such that p(y|X) > 0. The proof proceeds similarly, with the resulting partition size $E_X[|\mathcal{Y}(X)| \min_{y \in \mathcal{Y}(X)} p(y|X)] = \rho^*$.

Lemma 3.3. Given a distribution p(X,Y) and any predefined partition group p(A'|X,Y) with $A' \in \mathcal{A}'$, $|\mathcal{A}'|$ finite. Let $\hat{h}, \hat{R} = \{\arg\} \min_{h \in \mathcal{H}} \max_{a' \in \mathcal{A}'} r_{a'}(h)$ be the minimax fair solution for this partition and its corresponding minimax risk. Let h^* and R^* be the classifier and risk that solve Problem 4 with $\rho = \min_{a'} p(a')$. Then the price of minimax fairness can be upper bounded by

$$\max_{a' \in \mathcal{A}'} r_{a'}(h^*) - \hat{R} \le R^* - \min_{h \in \mathcal{H}} r(h). \tag{10}$$

Proof. Observe that $\forall A$, and for any distribution $p(A \mid X, Y), A \in A$ and $\forall h' \in \mathcal{H}$ we have

$$\min_{h \in \mathcal{H}} r(h) \le r(h') = \sum_{a \in \mathcal{A}} p(a) r_a(h') \le \max_{a \in \mathcal{A}} r_a(h').$$

We also have

$$h^*, p^*(A|X, Y), R^* = \{\arg\} \min_{h \in \mathcal{H}_{P_A}} \max_{a \in \{0, 1\}} r_a(h).$$
$$p(A|X, Y)$$
$$s.t. p(A) \succeq \rho$$

Which, together with Lemma 3.1, implies

$$\max_{a' \in \mathcal{A}'} r_{a'}(h^*) \le R^* \le \max_{a^* \in \{0,1\}} r_{a^*}(h').$$

We combine the two and show

$$\max_{a' \in \mathcal{A}'} r_{a'}(h^*) - \hat{R} \leq R^* - \hat{R}$$
$$\leq R^* - \min_{h \in \mathcal{H}} r(h).$$

Lemma 4.1. Given Problem 4 with minimum group size $\rho \leq \frac{1}{2}$, the following problems are value equivalent:

$$\begin{split} R^{\mathrm{I}} &= \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y)} \max_{a \in \{0,1\}} r_a(h), \\ s.t.p(A) \succeq \rho \\ R^{\mathrm{II}} &= \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y)} r_1(h), \\ s.t. \ p(A=1) = \rho \\ R^{\mathrm{I}} &= R^{\mathrm{II}}. \end{split}$$

Proof. Following the arguments in the proof of Lemma 3.1 we observe that, for any $h \in \mathcal{H}$ and $\mathcal{A} = \{0, 1\}$, we can consider the partition proposed in Equation 9 with a' = 1, which is a risk maximizing distribution for that particular h. This distribution satisfies $\max_{a \in \{0, 1\}} r_a(h) = r_1(h)$, and also satisfies $p(A = 1) = \rho$. Following the same reasoning as in the proof of Lemma 3.1, we can translate this equivalence in terms of worst case risk from the set $h \in \mathcal{H}$ to the set $h \in \mathcal{H}_{P_A}$.

Lemma 4.2. Given the problem on the right hand side of Eq. 6, a convex hypothesis class \mathcal{H} , and a bounded loss function $0 \le \ell(h(x), y) \le C \ \forall x, y, h \in \mathcal{X} \times \mathcal{Y} \times \mathcal{H}$ that is strictly convex w.r.t its first input h(x), under mild conditions the following problems are equivalent:

$$\begin{split} \mathcal{H}^{\mathrm{I}} &= \{ \arg \} \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} &\max_{p(A|X,Y)} r_{1}(h) \\ &s.t. \ p(A=1) = \rho \\ \mathcal{H}^{\mathrm{II}} &= \{ \arg \} \min_{h \in \mathcal{H}} &\sup_{p(A|X,Y)} r_{1}(h), \\ &p(A|X,Y) \\ &s.t. \ p(A=1) = \rho \\ &p(A=1|X,Y) > 0, \ \forall X,Y \end{split}$$

$$\mathcal{H}^{\mathrm{I}} &= \mathcal{H}^{\mathrm{II}}, \end{split}$$

Proof. We present this proof in two steps. First, we show that, under the hypothesis class \mathcal{H}_{P_A} , we can change the maximum over the set of distributions $P(A|X,Y):P(A=1)=\rho$ for the supremum over the set of distributions $P(A|X,Y):P(A=1)=\rho$, P(A=1|X,Y)>0 $\forall X,Y$. That is,

$$\begin{aligned} \{\arg\} \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} & \max_{p(A|X,Y)} & r_1(h) &= \{\arg\} \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} & \sup_{p(A|X,Y)} & r_1(h). \\ s.t. & p(A=1) = \rho & s.t. & p(A=1) = \rho \\ & p(A=1|X,Y) > 0 \ \forall X, Y \end{aligned}$$

To prove this we start by defining the set of distributions complying with the restriction on the left hand side as

$$Q_{\rho,\geq} = \{p(A|X,Y): \int p(A=1|x,y)p(x,y) = \rho, p(A=1|X,Y) \geq 0 \ \forall X,Y \in \mathcal{X} \times \mathcal{Y}\},$$

and the distribution subset on the right hand side as

$$Q_{\rho,>} = \{ p(A|X,Y) : \int p(A=1|x,y)p(x,y) = \rho, p(A=1|X,Y) > 0 \ \forall X,Y \in \mathcal{X} \times \mathcal{Y} \}.$$

П

We can then observe that, for any model h, and distributions $\hat{p}(A|X,Y) \in Q_{\rho,\geq}$ and $\bar{p}(A|X,Y) \in Q_{\rho,>}$, the distribution $p_{\lambda}(A|X,Y) = \lambda \bar{p}(A|X,Y) + (1-\lambda)\hat{p}(A|X,Y)$ satisfies $p_{\lambda}(A|X,Y) \in Q_{\rho,>} \ \forall \lambda \in (0,1]$. Furthermore, we have, by linearity of expectation

$$r_{1}(h; p_{\lambda}(A|X, Y)) = \lambda r_{1}(h; \bar{p}(A|X, Y)) + (1 - \lambda)r_{1}(h; \hat{p}(A|X, Y))$$

$$\leq \lambda C + (1 - \lambda)r_{1}(h; \hat{p}(A|X, Y)),$$

$$r_{1}(h; p_{\lambda}(A|X, Y)) \geq (1 - \lambda)r_{1}(h; \hat{p}(A|X, Y))$$

where we used explicit notation to indicate what distribution we are using to take expectation and the fact that the loss is upper bounded by C and lower bounded by 0. Therefore we conclude

$$\lim_{\lambda \to 0^+} p_{\lambda}(A|X,Y) = \hat{p}(A|X,Y)$$

and

$$\lim_{\lambda \to 0^+} r_1(h; p_{\lambda}(A|X, Y)) = r_1(h; \hat{p}(A|X, Y)).$$

Similarly for r_0

$$\lim_{\lambda \to 0^+} r_0(h; p_{\lambda}(A|X, Y)) = r_0(h; \hat{p}(A|X, Y)).$$

Since this transformation preserves the entire risk vector $r_0(h)$, $r_1(h)$, and the results hold for any $h \in \mathcal{H}$ and $\hat{p}(A|X,Y) \in Q_{\rho,\geq}$, we can conclude

$$\{\arg\} \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \max_{p(A|X,Y) \in Q_{\rho,\geq}} r_1(h) = \{\arg\} \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \sup_{p(A|X,Y) \in Q_{\rho,>}} r_1(h).$$

Secondly, we show that, under these conditions, and further assuming that \mathcal{H} does not contain models with unbounded risk trade-offs between groups, minimizing the supremum over $h \in \mathcal{H}_{P_A}$ is the same as minimizing over $h \in \mathcal{H}$. That is,

$$\{\arg\} \min_{h \in \mathcal{H}_{P_{\mathcal{A}}}} \sup_{p(A|X,Y)} r_1(h) = \{\arg\} \min_{h \in \mathcal{H}} \sup_{p(A|X,Y)} r_1(h).$$

$$s.t. \ p(A=1) = \rho$$

$$p(A=1|X,Y) > 0 \ \forall X,Y$$

$$s.t. \ p(A=1) = \rho$$

$$p(A=1|X,Y) > 0 \ \forall X,Y$$

We observe that, if ℓ is a strictly convex function w.r.t h, and $p(A|X,Y) \in Q_{\rho,>}$, we can write the following statements.

Let $\hat{h}, \bar{h} \in \arg\min_{h \in \mathcal{H}} R_1(h; p(A|X,Y))$ such that $\hat{h}(x) \neq \bar{h}(x)$ if and only if x in some set $\bar{\mathcal{X}} \subseteq \mathcal{X}$, and let $h_{\lambda} = \lambda \hat{h} + (1 - \lambda)\bar{h} \in \mathcal{H} \ \forall \lambda \in [0,1]$. By the strict convexity of ℓ we have

$$\ell(h_{\lambda}(X), Y) = \lambda \ell(\hat{h}(X), Y) + (1 - \lambda)\ell(\bar{h}(X), Y) \ \forall X, Y \in \mathcal{X} \setminus \bar{\mathcal{X}} \times \mathcal{Y},$$

$$\ell(h_{\lambda}(X), Y) < \lambda \ell(\hat{h}(X), Y) + (1 - \lambda)\ell(\bar{h}(X), Y) \ \forall X, Y \in \bar{\mathcal{X}} \times \mathcal{Y}.$$

Since for any $h \in \mathcal{H}$ we can write

$$r_1(h) = \int_{x \in \bar{\mathcal{X}}} \int_{y \in \mathcal{Y}} \frac{p(x, y)p(a = 1|X, Y)}{\rho} \ell(h(X), Y) dx dy + \int_{x \in \mathcal{X} \setminus \bar{\mathcal{X}}} \int_{y \in \mathcal{Y}} \frac{p(x, y)p(a = 1|X, Y)}{\rho} \ell(h(X), Y) dx dy,$$

and we need $r_1(h_\lambda) \geq R_1(\bar{h}) = r_1(\hat{h})$, using the inequalities from the strict convexity of ℓ we note that $\bar{\mathcal{X}}$ must satisfy

$$\int_{x \in \bar{\mathcal{X}}} \int_{y \in \mathcal{Y}} \frac{p(x, y)p(a = 1|X, Y)}{\rho} dx dy = 0,$$

or, equivalently, since $\{x, y : p(A = 1 | x, y) > 0\} = \mathcal{X} \times \mathcal{Y}$ by hypothesis

$$\int_{x \in \bar{\mathcal{X}}} \int_{y \in \mathcal{Y}} p(x, y) dx dy = 0.$$

From this we conclude that \hat{h} and \bar{h} can differ only in a zero-measure set, and thus $r_0(\hat{h}) = r_0(\bar{h})$, which implies that $\hat{h}, \bar{h} \in \arg\min_{\mathcal{H}_{P_A}} r_1(h; p(A|X,Y))$ for any $p(A|X,Y) \in Q_{\rho,>}$.

Lemma 4.3. Consider the setting of Algorithm 1, with parameter $\epsilon > 0$, and $\eta = \max_{\alpha \in \mathcal{U}_{\epsilon,\rho}} \frac{||\alpha||_2}{\sqrt{2T}} \le \sqrt{\frac{n\rho}{2T}}$ with $\mathcal{U}_{\epsilon,\rho} = \{\alpha : \alpha_i \in [\epsilon,1], \sum_i \frac{\alpha_i}{n} = \rho\}$, and L a 1-Lipschitz function w.r.t. α , let P be a uniform distribution over the set of models $\{h^1,\ldots,h^T\}$, and let R^* be the minimax solution to the loss presented in Eq. 8. Then we have

$$\max_{\alpha \in \mathcal{U}_{\epsilon,\rho}} \mathbb{E}_{h \sim P} L(h, \alpha) \leq \gamma R^* + \sqrt{\frac{2n\rho}{T}}.$$

Proof. We observe that loss function $L(h,\alpha)$ is concave (linear) w.r.t. α , and the set $\mathcal{U}_{\epsilon,\rho}$ is convex, with maximum norm $\max_{\alpha \in \mathcal{U}_{\epsilon,\rho}} ||\alpha||_2 \leq \sqrt{n\rho}$. For each $\epsilon > 0$ we are therefore able to use Theorem 7 in (Chen et al., 2017) to state

$$\max_{\alpha \in \mathcal{U}_{\epsilon,\rho}} \mathbb{E}_{h \sim P} L(h,\alpha) \leq \gamma \min_{h \in \mathcal{H}} \max_{\alpha \in \mathcal{U}_{\epsilon,\rho}} L(h,\alpha) + \max_{\alpha \in \mathcal{U}_{\epsilon,\rho}} ||\alpha||_2 \sqrt{\frac{2}{T}},$$

$$\leq \gamma \min_{h \in \mathcal{H}} \max_{\alpha \in \mathcal{U}_{\epsilon,\rho}} L(h,\alpha) + \sqrt{\frac{2n\rho}{T}}.$$

A.2. Additional Results

Similar to Table 1, tables 3, 5 and 7 compare the performance of the competing methods (Baseline, ARL, DRO and BPF) on a predefined demographic. For the law school dataset we considered gender and outcome; race and outcome was considered for the Compas and MIMIC-III datasets. Table 6 show the demographic composition of worst groups based on the mentioned populations. It is worth noting that for these particular predefined groups there is no major difference between DRO and BPF, moreover, in the case of Compas they do not seem to be deviating from the uniform classifier despite increasing the partition size, which could be due to a high level of noise.

| Group | Prop(%) | ARL.15 | DRO .15 | BPF .15 ARL.25 | DRO .25 | BPF .25 ARL.4 | DRO .4 | BPF .4 |
|----------------------------|----------------|----------------|----------------|--------------------------------|---------------|--------------------------------|----------------|----------------|
| White,Male White,Female | 48.2% 35.5% | 68.9% 70.5% | 85.3% 85.2% | 86.8% 68.9% 87.0% 70.5% | 94.1% $94.4%$ | 94.1% 68.9% 94.5% 70.5% | 94.1% 94.4% | 94.1% 94.4% |
| Nonwhite,Male | 7.8% | 59.5% | 76.6% | 77.5% 59.5% | 81.9% | 82.0% 59.5% | 81.9% | 81.9% |
| Nonwhite,Female | 8.5% | 58.8% | 75.9% | 76.7% 58.8% | 80.9% | 80.9% 58.8% | 80.9% | 80.9% |
| Worst-case group | - | 48.0% | 64.5% | 64.5% 49.6% | 75.7% | 75.7% 52.2% | 83.0% | 83.0% |

Table 3. Accuracy on law school dataset across gender and ethnicity partitions (groups given no special consideration by the algorithms), worst-case group performance is shown for reference. Results shown for ARL, DRO and BPF models for varying partition sizes.

Blind Pareto Fairness and Subgroup Robustness

| Group | Prop(%) ARL.15 | DRO .15 | BPF .15 A | ARL.25 | DRO .25 | BPF .25 | ARL.4 | DRO .4 | BPF .4 |
|------------------|------------------|---------|-------------|--------|---------|---------|-------|--------|--------|
| Black | 8.5% 50.0% | 52.8% | 02.070 | 50.0% | 62.6% | | 50.0% | 74.7% | 73.2% |
| White | 79.3% 50.0% | 52.8% | 51.2% 5 | 50.0% | 62.6% | 60.6% | 50.0% | 74.5% | 73.1% |
| Asian | 2.7% 50.0% | 52.7% | 51.2% 5 | 50.0% | 62.2% | 60.2% | 50.0% | 73.8% | 72.4% |
| Hispanic | 3.5% $50.0%$ | 53.1% | 51.4% 5 | 50.0% | 63.9% | 61.7% | 50.0% | 77.3% | 75.6% |
| Other | 6.0% 50.0% | 52.8% | 51.2% 5 | 50.0% | 62.3% | 60.4% | 50.0% | 74.1% | 72.6% |
| Worst-case group | - 50.0% | 48.5% | 48.7% 5 | 50.0% | 55.9% | 58.9% | 50.0% | 68.2% | 69.1% |

Table 4. Accuracy across gender and ethnicity partitions (groups given no special consideration by the algorithms) in the MIMIC-III dataset for ARL, DRO and BPF models for varying partition sizes. Worst-case group performance is shown for reference.

| Group | Prop(%) ARL.15 | DRO .15 | BPF .15 ARL.25 | DRO .25 | BPF .25 ARL.4 | DRO .4 | BPF .4 |
|----------------------------------------------------|----------------------------------------------------------------|----------------------------------|-----------------------------------------------------------------------------------|----------------------------------|-----------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| African-American Caucasian Hispanic Other | 51.8% 55.9% 34.3% 57.0% 8.3% 55.7% 5.6% 57.3% | 52.1% 52.4% 52.0% 53.3% | 53.0% 55.9% 54.5% 57.0% 53.8% 55.7% 55.2% 57.3% | 52.1% 52.4% 52.0% 53.3% | 53.0% 55.9% 54.5% 57.0% 53.8% 55.7% 55.2% 57.3% | 52.1% 52.4% 52.0% 53.3% | 53.0% 54.5% 53.8% 55.2% |
| Worst-case group | - 36.7% | 41.5% | 41.0% 40.2% | 44.0% | 43.4% 43.7% | 45.2% | 44.7% |

Table 5. Accuracy on Compas dataset across ethnicity partitions (groups given no special consideration by the algorithms), worst-case group performance is shown for reference. Results are shown for ARL, DRO and BPF models for varying partition sizes.

A.3. Experimental Details

To train the BPF classifier with standard SGD, we use the minibatch version of Eq.8

$$L(h, \boldsymbol{\alpha}) = \frac{1}{n_b} \frac{\sum_{i=1}^{n_b} \alpha_i \ell(h(x_i), y_i)}{\rho},$$

where n_b is the batch size, this loss is simply the minibatch equivalent of the importance-weighted loss

$$\mathbb{E}_{i \sim i \mid A=1}[\ell(h(x_i), y_i)] = \mathbb{E}_{i \sim U[\frac{\alpha_1}{N\rho}, \dots, \frac{\alpha_N}{N\rho}]}[\ell(h(x_i), y_i)] = \mathbb{E}_{i \sim U[1, \dots, N]}[\frac{\alpha_i}{\rho}\ell(h(x_i), y_i)]$$

The training parameters used for all methods presented in the main paper are summarized in the following table

All methods are implemented using the same codebase on PyTorch. Experiments are run on a single GeForce RTX 2080 Ti and take less than 1 hour wall time each. The code implementation for BPF will be provided.

| Group/Outcome | prop(%) | BPF .15 | BPF .30 | BPF .40 | | | |
|--------------------|---------|---------|---------|---------|--|--|--|
| Law school | | | | | | | |
| White,Male/0 | 1.5% | 4.5% | 2.6% | 1.5% | | | |
| White, Male/1 | 46.6% | 36.9% | 44.0% | 46.6% | | | |
| White,Female/0 | 1.2% | 3.5% | 2.0% | 1.2% | | | |
| White,Female/1 | 34.3% | 27.6% | 29.7% | 34.3% | | | |
| Nonwhite, Male/0 | 1.1% | 3.2% | 1.8% | 1.1% | | | |
| Nonwhite, Male/1 | 6.7% | 9.2% | 8.4% | 6.7% | | | |
| Nonwhite, Female/0 | 1.3% | 3.8% | 2.2% | 1.3% | | | |
| Nonwhite,Female/1 | 7.2% | 11.4% | 9.3% | 7.2% | | | |
| Compas | | | | | | | |
| African-American/0 | 24.7% | 24.7% | 24.7% | 24.7% | | | |
| African-American/1 | 27.1% | 27.1% | 27.1% | 27.1% | | | |
| Caucasian/0 | 20.9% | 20.9% | 20.9% | 20.9% | | | |
| Caucasian/1 | 13.4% | 13.4% | 13.4% | 13.4% | | | |
| Hispanic/0 | 5.2% | 5.2% | 5.2% | 5.2% | | | |
| Hispanic/1 | 3.1% | 3.1% | 3.1% | 3.1% | | | |
| Other/0 | 3.6% | 3.6% | 3.6% | 3.6% | | | |
| Other/1 | 2.0% | 2.0% | 2.0% | 2.0% | | | |
| | MIM | IC-III | | | | | |
| Black/0 | 7.6% | 5.2% | 5.2% | 6.3% | | | |
| Black/1 | 1.0% | 3.4% | 3.2% | 2.2% | | | |
| White/0 | 70.0% | 48.4% | 48.4% | 58.7% | | | |
| White/1 | 9.2% | 31.0% | 31.1% | 20.7% | | | |
| Asian/0 | 2.4% | 1.7% | 1.7% | 2.0% | | | |
| Asian/1 | 0.3% | 1.1% | 1.1% | 0.8% | | | |
| Hispanic/0 | 3.2% | 2.2% | 2.2% | 2.7% | | | |
| Hispanic/1 | 0.3% | 0.9% | 0.9% | 0.6% | | | |
| Other/0 | 5.3% | 3.6% | 3.6% | 4.4% | | | |
| Other/1 | 0.7% | 2.4% | 2.5% | 1.6% | | | |

Table 6. Demographic composition of worst groups as a function of minimum partition size on the law school and Compas dataset. BPF homogenizes outcomes across partitions and protected attributes.

| Method | DRO | ARL | ARL HC | BPF |
|----------------------|--------------------------------|-------------------------------|-------------------------------|-------------------------|
| Learning rate | $5\exp{-5}$ | $5\exp{-5}$ | $5\exp{-5}$ | $5\exp{-5}$ |
| Batch size | 256 | 256 | 256 | 256 |
| Training Loss | cross entropy | cross entropy | cross entropy | cross entropy |
| Network architecture | (512,) | (512,) (512,512,), (512,), | (512,) (512,512,), (512,), | (512,) |
| Custom parameter | $\eta = \{0, 1/20, \dots, 1\}$ | (256,256,), (256,) | (256,256,), (256,) | $\rho = \{0.05,, 0.8\}$ |

Table 7. Summary of training parameters per method for all presented experiments. Network architecture is the number and size of the hidden layers of the classifer; for ARL and ARL HC, the adversarial network architecture is presented in the same format on the custom parameter row.