

# FLEA: Provably Robust Fair Multisource Learning from Unreliable Training Data

Anonymous authors

Paper under double-blind review

## Abstract

Fairness-aware learning aims at constructing classifiers that not only make accurate predictions, but also do not discriminate against specific groups. It is a fast-growing area of machine learning with far-reaching societal impact. However, existing fair learning methods are vulnerable to accidental or malicious artifacts in the training data, which can cause them to unknowingly produce unfair classifiers. In this work we address the problem of fair learning from unreliable training data in the robust multisource setting, where the available training data comes from multiple sources, a fraction of which might not be representative of the true data distribution. We introduce FLEA, a filtering-based algorithm that identifies and suppresses those data sources that would have a negative impact on fairness or accuracy if they were used for training. As such, FLEA is not a replacement of prior fairness-aware learning methods but rather an augmentation that makes any of them robust against unreliable training data. We show the effectiveness of our approach by a diverse range of experiments on multiple datasets. Additionally, we prove formally that –given enough data– FLEA protects the learner against corruptions as long as the fraction of affected data sources is less than half.

## 1 Introduction

Machine learning systems have started to permeate many aspects of our everyday life, such as finance (e.g. credit scoring), employment (e.g. judging job applications) or even judiciary (e.g. recidivism prediction). In the wake of this trend, other aspects besides prediction accuracy become important to consider. One crucial aspect is *(group) fairness*, which aims at preventing learned classifiers from acting in a discriminatory way. To achieve this goal, fairness-aware learning methods adjust the classifier parameters in order to fulfill an appropriate measure of fairness. This strategy is highly successful, but only under idealized conditions of clean i.i.d.-sampled data. Unfortunately, *fairness-aware learning methods are not robust against unintentional errors or intentional manipulations of the training data*.

In this work, we propose a new algorithm, FLEA (Fair LEarning against Adversaries) that overcomes this problem in the setting where the training data is not one monolithic block, but rather comes as a collection of multiple data sources. This is, in fact, a common scenario. For instance, organizations that specialize in large-scale data mining, such as large hospital chains or political analysis firms, may receive data that is collected separately from multiple physical locations or data vendors, of which not every one may be completely trustworthy, and so robustness concerns arise.

FLEA adds a filtering step that effectively identifies and suppresses data sources that could have a negative impact on the classifier fairness or accuracy, so long as they constitute less than half of all sources. To accomplish this, we introduce a new dissimilarity measure, *disparity*, that measures the maximum achievable difference in classifier fairness between two data sources. We combine this with the existing *discrepancy* measure, which plays an analogous role for the classifier accuracy, and the *disbalance*, which measures changes to the group composition of the training data. We show both empirically and theoretically that a combination




Figure 1: Illustration of robust multisource learning with FLEA: (top) We are given multiple sources, some of which might contain noisy or manipulated data. (left) Ordinary multisource learning pools the data from all sources, which can cause the resulting classifier to be inaccurate and/or unfair, even if fairness-aware training is employed. (right) FLEA filters the data before pooling, thereby suppressing likely corrupted sources. This allows fairness-aware training to succeed.

of these three measures provides a sufficient criterion for detecting harmful data, as long as the fraction of harmful sources is less than half.<sup>1</sup>

While previous method for robust fairness-aware learning were only able to protect against specific data issues, such as random label flips, FLEA ensures that even a worst-case adversary is unable to negatively affect the training process: either the changes to the data are minor and will not hurt learning, or they are large enough so that the affected data sources are identified and removed. Our theoretical analysis provides finite sample guarantees and certifies the ability of FLEA to learn classifiers with optimal fairness and accuracy in the infinite sample size limit. Our extensive experimental evaluation demonstrates FLEA’s practical usefulness in suppressing the effect of corrupted data when learning fair models, even in cases where previous robust methods fail.

## 2 Preliminaries and related work

### 2.1 Fair classification

Throughout this work, we adopt a standard classification setting in which the task is to predict a binary label  $y \in \mathcal{Y} = \{0, 1\}$  for any  $x \in \mathcal{X}$ . For a fixed data distribution  $p(x, y) \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$ , the classic goal of learning is to find a prediction function  $f : \mathcal{X} \rightarrow \mathcal{Y}$  with high accuracy, i.e. small *risk*,  $\mathcal{R}_p(f) = \mathbb{E}_{p(x, y)} \llbracket y \neq f(x) \rrbracket$ , where  $\llbracket P \rrbracket = 1$  if a predicate  $P$  is true and  $\llbracket P \rrbracket = 0$  otherwise.

With the recent trend to consider not only the accuracy but also the *fairness* of a classifier, a number of statistical measures have been proposed to formalize this notion. In this work, we focus on the most common and simplest one, *demographic parity (DP)* (Calders et al., 2009). It postulates that the probability of a positive classifier decision should be equal for all subgroups of the population. Formally, we assume that each example  $(x, y)$  also possesses a *protected attribute*,  $a \in \mathcal{A}$ , which indicates its membership in a specific subgroup of the population. For example,  $a$  could indicate *race*, *gender* or a *disability*. For simplicity of exposition, we treat the protected attribute as binary-valued, but extensions to multi-valued attributes are straightforward by summing over all pairwise terms. Note that  $a$  might be a component or a function of  $x$ , in which case it is available at prediction time, or it might be contextual information, in which case it would only be available for the learning algorithm at training time, but not for the resulting classifier at prediction time. We cover both aspects by treating  $a$  as an additional random variable, and write the underlying joint data distribution as  $p(x, y, a)$ .

<sup>1</sup>The case where half or more sources are harmful is impossible to solve in general, see e.g. Charikar et al. (2017)

For a classifier  $f : \mathcal{X} \rightarrow \{0, 1\}$ , the *demographic parity violation*,  $\Gamma_p$ , and the empirical counterpart,  $\Gamma_S$  for a dataset  $S \subset \mathcal{X} \times \mathcal{Y} \times \mathcal{A}$  are defined as (Calders et al., 2009; Dwork et al., 2012),

$$\Gamma_p(f) = \left| \mathbb{E}_{p(x|a=1)} f(x) - \mathbb{E}_{p(x|a=0)} f(x) \right|, \quad \Gamma_S(f) = \left| \frac{1}{n^{a=1}} \sum_{x \in S^{a=1}} f(x) - \frac{1}{n^{a=0}} \sum_{x \in S^{a=0}} f(x) \right| \quad (1)$$

where  $S^{a=z} = \{(x, y, a) \in S : a = z\}$  and  $n^{a=z} = |S^{a=z}|$  for  $z \in \{0, 1\}$ . Smaller values of the demographic parity violation indicate more fair classifiers. Analogous quantities can be defined for related fairness measures, such as *equality of opportunity* or *equalized odds* (Hardt et al., 2016; Zafar et al., 2017a). A detailed description of these and many others choices can be found in Barocas et al. (2019).

**Fairness-aware learning** In the last years, a plethora of algorithms have been developed that are able to learn classifiers that are not only accurate but also fair, see, for example Mehrabi et al. (2021a) for an overview. They mostly rely on one or multiple of four core mechanisms. *Preprocessing methods* (Kamiran & Calders, 2012; Calmon et al., 2017; Wang et al., 2019; Celis et al., 2020) change the training data to remove a potential bias. This is often simple and effective, but comes with the danger of reduced accuracy, since the data distribution at training time will not reflect the distribution at prediction time anymore. *Postprocessing methods* (Hardt et al., 2016; Woodworth et al., 2017; Chzheng et al., 2020) adjust the acceptance thresholds of a previously trained classifier for each protected group, so that the desired fairness criterion is met. This is a simple, reliable and often effective method, but it requires the protected attribute to be available at prediction time. *Penalty-based methods* (Kamishima et al., 2012; Zemel et al., 2013; Zafar et al., 2017b; Donini et al., 2018; Mandal et al., 2020; Chuang & Mroueh, 2021) add a regularizer or constraints to the learning objective that penalize or prevent parameter choices that lead to unfair decisions. *Adversarial methods* (Beutel et al., 2017; Wadsworth et al., 2018; Zhang et al., 2018; Lahoti et al., 2020) train an adversary in parallel to the classifier that tries to predict the protected attribute from the model outputs; if this cannot be done better than chance level, fairness is achieved.

Many other methods have been proposed, e.g. based on distributionally robust optimization (Rezaei et al., 2020) or tailored to a specific family of classifiers or optimization procedures (Cho et al., 2020; Tan et al., 2020). They all share, however, the property that accurate information about the data distribution and the protected attribute is needed at training time.

If the training data is not representative of the actual data distribution, e.g. it is noisy, biased, or has been manipulated, then fairness-enforcing mechanisms fall short (Kallus et al., 2020; Mehrabi et al., 2021b). Partial solutions have been proposed, e.g., when only the protected attribute or only the label is noisy (Lamy et al., 2019; Wang et al., 2020; Celis et al., 2021b;a; Mehrotra & Celis, 2021; Roh et al., 2021). However, as shown in Konstantinov & Lampert (2022), full protection against malicious manipulations of the training data is provably impossible when learning from a single dataset.

## 2.2 Robust multisource learning

**Learning from multiple sources** The multisource learning setting formalizes the increasingly frequent situation in which the training data is not collected as a single batch, but from multiple data sources (Ting & Low, 1997; Russakovsky et al., 2015). For fairness-aware learning, this means we are given  $N$  datasets,  $S_1, \dots, S_N \subset \mathcal{X} \times \mathcal{Y} \times \mathcal{A}$ . Each  $S_i = \{(x_1^{(i)}, y_1^{(i)}, a_1^{(i)}), \dots, (x_{n_i}^{(i)}, y_{n_i}^{(i)}, a_{n_i}^{(i)})\}$  contains i.i.d. samples from a data distribution  $p_i(x, y, a)$ . Given these datasets, the learning algorithm has the goal of selecting a prediction function  $f$  from a hypothesis class  $\mathcal{H}$  that has as-small-as-possible risk (expected prediction error) and unfairness (e.g. demographic parity violation) with respect to the unknown distribution at prediction time,  $p(x, y, a)$ , (also called *target distribution*). The classical setting of  $p_1 = p_2 = \dots = p_N = p$ , we call *homogeneous multisource learning*. Otherwise, we call the setting *heterogeneous*.

In a *clean data scenario*, when all data distribution are the same or very similar to each other, then there is no drawback to simply merging all sources and training on the resulting large dataset. However, merging all data is not the best strategy when some of the data sources are *unrepresentative*, i.e. their data distribution differs a lot from the target one. Such data can occur accidentally, for instance due to biases in the data collection or annotation process. In some cases, such issues can be overcome by domain adaptation techniques (Ben-David

et al., 2010; Crammer et al., 2008; Natarajan et al., 2013). Unrepresentative sources can also be the result of intentional manipulations, which are typically harder to detect and compensate for (Feng et al., 2019; Fowl et al., 2021). In fact, the datasets might not be samples from any probability distribution in that case, but adversarially constructed.

**Robust multisource learning** In this work we aim to cover as wide a range of possible problems with some of the data sources as possible. Therefore, we study the multisource learning problem in the presence of an *adversary*<sup>2</sup>. In this setting, the adversary observes an original collection of  $N$  datasets,  $\tilde{S}_1, \dots, \tilde{S}_N$ , where each  $\tilde{S}_i$  contains i.i.d. samples from a data distribution  $p_i(x, y, a)$ . Next, the adversary manipulates the data in an arbitrary (deterministic or randomized) way with the only restriction that for a fixed subset of indices,  $G \subset \{1, \dots, N\}$ , the data source remains unaffected. That is,  $S_i = \tilde{S}_i$  for all  $i \in G$ , and  $S_i$  is arbitrary for  $i \notin G$ . The subset  $G$  is unknown to the algorithm, of course. The adversary model places no restrictions on the corruptions, and thus subsumes many scenarios that have to otherwise be studied in isolation. In particular, both data-quality issues, such as sampling bias, data entry errors or label noise, as well as malicious manipulations, such as class erasure or data poisoning, are covered as special cases.

Multisource learning with protection against potential manipulations is known as *robust multisource learning* (Erfani et al., 2017). In order to detect harmful sources, a natural approach is to compare all pairs of datasets with an appropriate distance measure and then use the pairwise distances to filter out sources that are far from the others. Key to the success of such an approach is using the right definition of distance. On the one hand one must be able to estimate the measure from finite sample sets in a statistically efficient way. Many common information-theoretic measures, such as *Kullback-Leibler divergence* (Kullback & Leibler, 1951), *total variation* (Tsybakov, 2009) or *Wasserstein distance* (Villani, 2009), do not fulfill this criterion. On the other hand, the measure must be sensitive enough such that if two sources appear similar then training on either of them must yield similar classifiers. Classical two-sample tests, such as Student’s  $t$ -test (Student, 1908) or MMD (Gretton et al., 2012), fail to guarantee this.

In the context of multisource learning a measure that combines both useful properties is the (*empirical*) *discrepancy distance* (Kifer et al., 2004; Mohri & Medina, 2012). For two datasets,  $S_1, S_2$ , and a hypothesis set  $\mathcal{H} \subset \{h : \mathcal{X} \rightarrow \mathcal{Y}\}$ , it measures the maximal amount by which their estimates of the classification accuracy can differ:

$$\text{disc}(S_1, S_2) = \sup_{h \in \mathcal{H}} \left| \mathcal{R}_{S_1}(h) - \mathcal{R}_{S_2}(h) \right|, \quad (2)$$

where  $\mathcal{R}_S(h) = \frac{1}{|S|} \sum_{(x,y) \in S} \mathbb{I}[y \neq h(x)]$  is the *empirical risk* of  $h$  on  $S$ . In Konstantinov et al. (2020) the discrepancy is used as a distance measure to identify and suppress data source that might harm the classifier’s accuracy. However, the associated algorithm is mostly of theoretical interest: it only suppresses those sources of which it is certain that they are malignant using thresholds that are derived from its generalization bound. As a consequence, it requires training sets that are too large to be practical. Similarly, Jain & Orlitsky (2020b) provide an analysis of the learning-theoretic limitations of robust multisource learning. Konstantinov & Lampert (2019) also use the discrepancy measure for detecting harmful data sources, but the proposed algorithm requires access to a reference set that is guaranteed to be free of data manipulations. In Qiao & Valiant (2018); Chen et al. (2019); Jain & Orlitsky (2020a) robust multisource learning is addressed using tools from robust statistics, but only in the context of discrete density estimation.

All of the above works are tailored to the task of ensuring high accuracy of the learned classifiers or estimators, but they are not sensitive to issues of fairness. To our awareness, the only prior work that considers achieving fairness in a multisource learning setting and in the presence of data corruption is the one of Li et al. (2021). However, this paper focuses on *personalized federated learning* and on a fairness objective which postulates that models’ performances should be relatively similar across edge devices. In contrast, we study a setup where a *single global model is trained* and aim to ensure that this model does not act discriminatory against members of protected subgroups.

<sup>2</sup>*Adversary* is the common computer science term for a process whose aim it is to prevent a system from operating as intended. Our adversaries manipulate the training data and should not be confused with adversaries in *adversarial machine learning*, such as *adversarial examples* (Goodfellow et al., 2015), or *generative adversarial networks* (Goodfellow et al., 2014).

### 3 Fair multisource learning

The goal of this work is to develop a method that allows fairness-aware learning, even if some of the available data sources are unrepresentative of the true training distribution. For this, we introduce FLEA, a filtering-based algorithm that identifies and suppresses those data sources that would negatively impact the fairness of the trained classifier. Its main innovation is the *disparity measure* for comparing datasets in terms of their fairness estimates.

**Definition 1** (Empirical Disparity). For two datasets  $S_1, S_2 \subset \mathcal{X} \times \mathcal{Y} \times \mathcal{A}$ , their *empirical disparity* with respect to a hypothesis class  $\mathcal{H}$  is

$$\text{disp}(S_1, S_2) = \sup_{h \in \mathcal{H}} \left| \Gamma_{S_1}(h) - \Gamma_{S_2}(h) \right|. \quad (3)$$

where  $\Gamma_S : \mathcal{H} \rightarrow \mathbb{R}$  is an empirical (un)fairness measure, such as the *demographic parity violation* (1).

The *disparity* measures the maximal amount by which the estimated fairness of a classifier in  $\mathcal{H}$  can differ between using  $S_1$  or  $S_2$  as the basis of the estimate. A small disparity value implies that if we construct a classifier that is fair with respect to  $S_1$ , then it will also be fair with respect to  $S_2$ .

Definition 1 is inspired by the *empirical discrepancy* (2). Just as low discrepancy implies that a classifier learned on one dataset will have comparable accuracy as one learned on the other, low disparity implies that the two classifiers will have comparable fairness. FLEA relies on the discrepancy as well as the disparity, because ensuring fairness alone does not suffice (e.g. a constant classifier is perfectly fair). As a third relevant quantity we introduce the (*empirical*) *disbalance*.

$$\text{disb}(S_1, S_2) = \left| \frac{|S_1^{a=1}|}{|S_1|} - \frac{|S_2^{a=1}|}{|S_2|} \right|. \quad (4)$$

The disbalance compares the relative sizes of the protected groups of two datasets. Its inclusion is a technical requirement to be able to also formally prove that demographic parity fairness remains unaffected by corruptions.

In combination, disparity, discrepancy, and disbalance form an effective criterion for detecting dataset manipulations. This is most apparent in the homogeneous setting: if two datasets of sufficient size are sampled i.i.d. from distributions close to the target one, then by the law of large numbers we can expect all three measures to be small. If one of the datasets is sampled like this (called *clean* from now on) but the other is manipulated, then there are two possibilities. It is still possible that all three values are small. In this case, equations (2)–(4) ensure that neither accuracy nor fairness would be negatively affected, and we call such manipulations *benign*. If at least one of the values is large, training on such a manipulated dataset could have undesirable consequences. Such manipulations we will call *malignant*. Finally, when comparing two manipulated datasets, discrepancy, disparity, and disbalance can each have arbitrary values.

In the inhomogeneous setting, a path of similar reasoning applies, though the measures for clean sources will not approach exactly zero due to the difference in their data distributions.

#### 3.1 FLEA: Fair learning against adversaries

We now introduce the FLEA algorithm, which is able to learn fair classifiers even if up to half of the datasets are noisy, biased or have been manipulated. Similar to classic outlier rejection techniques (Barnett & Lewis, 1984) and statistical two-sample tests (Corder & Foreman, 2014), the main algorithm (Algorithm 1) takes a filtering approach. Given the available data sources and additional parameters, it calls a subroutine that identifies a subset of clean or benign sources, merges the training data from these, and trains a (presumably fairness-aware) learning algorithm on the resulting dataset.

FLEA’s crucial component is the filtering subroutine. This estimates the pairwise disparity, discrepancy and disbalance between all pairs of data sources and combines them into a matrix of dissimilarity scores (short: *D*-scores). As discussed above, large values indicate that at least one of the two compared sources must be *malignant*. It is not a priori clear, though, how to use this information. On the one hand, we do not know which of the two datasets is malignant or if both are.

On the other hand, malignant sources can also occur in pairs with small  $D$ -score, when both datasets were manipulated in similar ways. Finally, even the  $D$ -scores between two clean or benign sources will have non-zero values, which depend on a number of factors, in particular the data distributions and the hypothesis class.

FLEA overcomes this problem by using tools from robust statistics. For any dataset  $S_i$ , it computes a value  $q_i$  (called  $q$ -value) as the  $\beta$ -quantile of the  $D$ -scores to all other datasets, where  $\beta$  is a hyperparameter we discuss below. It then computes the  $\beta$ -quantile of all such values and selects those datasets with  $q$ -values up to this threshold.

To see that this procedure has the desired effect of filtering out malignant datasets, we first look at the case in which the sources are homogeneous and  $\beta = \frac{K}{N}$ , where  $K = |G| > \frac{N}{2}$  is the number of clean data sources.

For any clean dataset  $S_i$ , by assumption there are at least  $K - 1$  other clean sources with which it is compared. We can expect the  $D$ -scores of these pairs are small, and, of course, that  $D_{ii} = 0$ . Because  $\beta = \frac{K}{N}$ , the  $\beta$ -quantile,  $q_i$ , is simply the  $K$ th-smallest of  $S_i$ 's  $D$ -scores. Consequently,  $q_i$  will be at least as small as the result of comparing two clean sources. For benign sources, the same reasoning applies, since their  $D$ -scores are indistinguishable from clean ones. For a malignant  $S_i$ , at least  $K$  of the  $D$ -scores will be large, namely the ones where  $S_i$  is compared to a clean source. Hence, there can be at most  $N - K$  small  $D$ -scores for  $S_i$ . Because  $\beta N = K$  and  $K > N - K$ , the  $\beta$ -quantile  $q_i$  will be at least as large as comparing a clean dataset to a malignant one.

Choosing those sources that fall into the  $\beta$ -quantile of the  $q_i$  values means selecting the  $K$  sources of smallest  $q_i$  value. By the above argument, these will either be not manipulated at all, or only in a way that does not have a negative effect on either the fairness or the accuracy of the training process. In practice, the regimes of *large* and *small*  $D$ -scores can overlap due to noise in the sampling process, and the perfect filtering property will only hold approximately. We later discuss a generalization bound that makes this reasoning rigorous.

Revisiting the above arguments one sees that the guarantees on the  $q_i$  follow also for any  $\beta > \frac{N-K}{N}$ , so in particular for  $\beta \geq \frac{1}{2}$ . To obtain the guarantee on the selected sources,  $\beta \leq \frac{K}{N}$  suffices. Therefore, even if the exact value of  $K$  is unknown in practice, setting  $\beta = \frac{1}{2} + \frac{1}{N}$  for even  $N$  and  $\beta = \frac{1}{2} + \frac{1}{2N}$  for odd  $N$  will always be working choices. These are also the values we use in our experiments.

In the heterogeneous situation, the  $D$ -scores between clean sources might not tend to zero for large  $n$  anymore. However, they will approach the true discrepancy, disparity and disbalance values between the sources' distributions. From this, one can obtain a guarantee that the selected sources are not more dissimilar from each other than the clean sources are, which is the best one can hope for in the heterogeneous setting.

### 3.2 Implementation

FLEA is straightforward to implement, with only the discrepancy and disparity estimates in the FILTER-SOURCES routine requiring some consideration. Naively, these would require optimizing combinatorial functions (the differences of fraction of errors or positive decisions) over all functions in the hypothesis class. This task is at least as hard as the problem of separating two point sets by a hyperplane, which is known to be NP-hard (Marcotte & Savard, 1992) and even difficult to approximate under any real-world conditions. Instead, we exploit the structure of the optimization problems to derive tractable approximations. For the discrepancy (2) such a method was originally proposed in the domain adaptation literature (Ben-David et al.,

---

#### Algorithm 1 FLEA

---

**Input:** datasets  $S_1, \dots, S_N$

**Input:** quantile parameter  $\beta$

**Input:** (fairness-aware) learning algorithm  $\mathcal{L}$

1:  $I \leftarrow \text{FILTERSOURCES}(S_1, \dots, S_N; \beta)$

2:  $S \leftarrow \bigcup_{i \in I} S_i$

3:  $f \leftarrow \mathcal{L}(S)$

**Output:** trained model  $f : \mathcal{X} \rightarrow \mathcal{Y}$

---



---

#### Subroutine FILTERSOURCES

---

**Input:**  $S_1, \dots, S_N; \beta$

1: **for**  $i = 1, \dots, N$  **do**

2:   **for**  $j = 1, \dots, N$  **do**

3:      $D_{i,j} \leftarrow \text{disc}(S_i, S_j) + \text{disp}(S_i, S_j) + \text{disb}(S_i, S_j)$

4:   **end for**

5:    $q_i \leftarrow \beta\text{-quantile}(D_{i,1}, \dots, D_{i,N})$

6: **end for**

7:  $I \leftarrow \{i : q_i \leq \beta\text{-quantile}(q_1, \dots, q_N)\}$

**Output:** index set  $I$

---

2010): finding the hypothesis with maximal accuracy difference between two datasets is equivalent to training a binary classifier on their union with the labels of one of the datasets flipped. From the solution one recovers the discrepancy value as  $1 - 2E$ , where  $E$  denotes the (training) error rate.

For the disparity (3), we propose an analogous route. Intuitively, the optimization step requires finding a hypothesis that is as unfair as possible on  $S_1$  (i.e. maximizes  $\Gamma_{S_1}$ ) while being as fair as possible on  $S_2$  (i.e. minimizes  $\Gamma_{S_2}$ ), or vice versa. From Equation (1) one sees that a hypothesis  $f$  is maximally unfair if it outputs  $f(x) = 1$  on  $S_1^{a=1}$  and  $f(x) = 0$  on  $S_1^{a=0}$ , or vice versa. This can be achieved by training a classifier to predict  $f(x) = a$  on  $S_1$ . To give both protected groups equal importance, as the definition requires, we use per-sample weights that are inversely proportional to the group sizes. To trade off the unfairness on  $S_1$  with the fairness on  $S_2$ , we simply add  $\Gamma_{S_2}$  as a regularizer to this optimization problem. Consequently, estimating  $\text{disp}(S_1, S_2)$  becomes as efficient as training a classifier with fairness regularization. As the above construction breaks the symmetry between the roles of  $S_1$  and  $S_2$ , we apply it a second time with the roles of  $S_1$  and  $S_2$  interchanged, and we keep the larger of the two values as estimate of  $\text{disp}$ .

### 3.3 Theoretical analysis

The informal justification of FLEA can be made precise in the form of a generalization bound. In this section we present our theoretical guarantees for FLEA. We begin by stating formally the assumptions we make on the data generating process, both for the heterogeneous and the homogeneous cases discussed above. We then state our main theoretical result, which certifies the performance of FLEA in the both the homogeneous and the more general heterogeneous case. Finally, we briefly outline the main proof steps. The full proofs can be found in Appendix D.

#### 3.3.1 Assumptions and formal adversarial model

First we present our formal set of assumptions, directly in the general setting of *heterogeneous* data sources. A crucial parameter here setup is  $\eta$ , which denotes the amount of variability between the clean sources' distributions. The case of  $\eta = 0$  recovers the *homogeneous* setup.

We assume the following data generation model, similar to the one of Qiao & Valiant (2018). By  $p(x, y, a)$  we denote the target distribution. It is unknown to the learning algorithm, though potentially known to the adversary. Initially, there are  $N$  datasets  $\tilde{S}_1, \dots, \tilde{S}_N$ , with the  $i$ -th set of samples being drawn i.i.d. from a distribution  $p_i(x, y, a)$ . These distributions might differ from the target distribution  $p$  by at most  $\eta$  in terms of total variation both with respect to the overall distributions as well as the conditional distributions with respect to  $a$ . Formally, we assume the following conditions for  $i = 1, \dots, N$ :

$$\text{TV}(p_i(x, y, a), p(x, y, a)) \leq \eta, \quad \text{and} \quad \max_{z \in \mathcal{A}} \{ \text{TV}(p_i(x, y|a=z), p(x, y|a=z)) \} \leq \eta, \quad (5)$$

where  $\text{TV}(p, q) = \sup_{B \in \mathcal{B}(\mathcal{X} \times \mathcal{Y} \times \mathcal{A})} |p(B) - q(B)|$  with  $\mathcal{B}(X)$  denoting the Borel  $\sigma$ -algebra on a topological space  $X$ .<sup>3</sup> This *total variation (TV) distance* is a standard measure of distance between distributions.

Once the clean datasets  $\tilde{S}_1, \dots, \tilde{S}_N$  are sampled, an *adversary* operates on them. This results in new datasets,  $S_1, \dots, S_N$ , which the learning algorithm receives as input. The adversary is an arbitrary (deterministic or randomized) function  $\mathcal{F} : \prod_{i=1}^N (\mathcal{X} \times \mathcal{Y} \times \mathcal{A})^{n_i} \rightarrow \prod_{i=1}^N (\mathcal{X} \times \mathcal{Y} \times \mathcal{A})^{n_i}$ , with the only restriction that for a fixed subset of indices,  $G \subset \{1, \dots, N\}$ , the data remains unchanged. That is,  $S_i = \tilde{S}_i$  for all  $i \in G$ , and  $S_i$  is arbitrary for  $i \notin G$ . For simplicity, we refer to a dataset  $S_i$  or a source  $i \in [N]$  as clean if  $i \in G$ .

#### 3.3.2 Theoretical guarantees on FLEA

We are now ready to state our theoretical guarantee on FLEA. For simplicity of notation, we present the case where all sources have the same number of samples. Results for general sample sizes can be obtain in an analogous way. We first state the guarantees for the homogeneous situations, which we obtain in fact as a corollary for  $\eta = 0$  of the general theorem later in this section.

<sup>3</sup>For the TV bounds we assume that the product  $\sigma$ -algebra of the corresponding Borel  $\sigma$ -algebras on  $\mathcal{X}$ ,  $\mathcal{Y}$  and  $\mathcal{A}$  is used.

**Theorem 1** (Homogeneous setting). *Assume that  $\mathcal{H}$  has a finite VC-dimension  $d \geq 1$ . Let  $p$  be an arbitrary target data distribution and without loss of generality let  $\tau = p(a = 0) \in (0, 0.5]$ . Let  $S_1, \dots, S_N$  be  $N$  datasets, each consisting of  $n_i$  samples, out of which  $K > \frac{N}{2}$  are sampled i.i.d. from the distribution  $p$ . For  $\frac{1}{2} < \beta \leq \frac{K}{N}$  and  $I = \text{FILTERSOURCES}(S_1, \dots, S_N; \beta)$  set  $S = \bigcup_{i \in I} S_i$ . Let  $\delta > 0$ . Then there exists a constant  $C = C(\delta, \tau, d, N, \eta)$ , such that for any  $n \geq C$ , the following inequalities hold with probability at least  $1 - \delta$  uniformly over all  $f \in \mathcal{H}$  and against any adversary:*

$$\Gamma_S(f) \leq \Gamma_p(f) + \tilde{\mathcal{O}}\left(\sqrt{\frac{1}{n}}\right), \quad \mathcal{R}_S(f) \leq \mathcal{R}_p(f) + \tilde{\mathcal{O}}\left(\sqrt{\frac{1}{n}}\right), \quad (6)$$

where  $\tilde{\mathcal{O}}$  indicates Landau’s big- $\mathcal{O}$  notation for function growth up to logarithmic factors (Cormen et al., 2009).

To analyze the statement, we observe that Equation (6) ensures that for large enough training sets the filtered training data  $S$  becomes an arbitrarily good representative of the true underlying data distribution with respect to the classification accuracy as well as the fairness. The approximation holds uniformly across all hypotheses in the class, therefore it is safe to use  $S$  for fairness-aware learning. Note that despite the intuitive conclusion, this is a highly non-trivial statement. For example, in the case of learning from a single datasource in which a constant fraction of the data can be manipulated, an analogous theorem is provably impossible (Kearns & Li, 1993; Konstantinov & Lampert, 2022).

For the general situation ( $\eta \geq 0$ ), we obtain the following guarantees:

**Theorem 1** (Heterogeneous setting). *Assume that  $\mathcal{H}$  has a finite VC-dimension  $d \geq 1$ . Let  $p$  be an arbitrary target data distribution and without loss of generality let  $\tau = p(a = 0) \in (0, 0.5]$ . Let  $S_1, \dots, S_N$  be  $N$  datasets, each consisting of  $n$  samples, out of which  $K > \frac{N}{2}$  are sampled i.i.d. from data distributions  $p_i$  that are  $\eta$ -close the distribution  $p$  in the sense of Section 3.3.1. Assume that  $9\eta < \tau$ . For  $\frac{1}{2} < \beta \leq \frac{K}{N}$  and  $I = \text{FILTERSOURCES}(S_1, \dots, S_N; \beta)$  set  $S = \bigcup_{i \in I} S_i$ . Let  $\delta > 0$ . Then there exists a constant  $C = C(\delta, \tau, d, N, \eta)$ , such that for any  $n \geq C$ , the following inequalities hold with probability at least  $1 - \delta$  uniformly over all  $f \in \mathcal{H}$  and against any adversary:*

$$\Gamma_S(f) \leq \Gamma_p(f) + \mathcal{O}(\eta) + \tilde{\mathcal{O}}\left(\sqrt{\frac{1}{n}}\right), \quad \mathcal{R}_S(f) \leq \mathcal{R}_p(f) + \mathcal{O}(\eta) + \tilde{\mathcal{O}}\left(\sqrt{\frac{1}{n}}\right). \quad (7)$$

In contrast to the homogeneous situation, an additional factor linear in  $\eta$  enters the right hand side of the bound. As discussed in Section 3.1, we believe that such a factor will be unavoidable in the heterogeneous case:  $\eta$  is a measure of the dissimilarity between the clean sources and the target distribution. Therefore, even if no corrupted data was present, the accuracy of a learned classifier for the unknown target distribution  $p$  will be limited by how close that is to the training distributions, i.e.  $\eta$  (Bartlett, 1992; Hanneke & Kpotufe, 2020).

**Proof sketch** The proof consists of three steps. First, we characterize a set of values into which the empirical risks and empirical deviation measures of the clean data sources fall with probability at least  $1 - \delta$ . Then we show that because the clean datasets cluster in such a way, any individual dataset that is accepted by the FILTERSOURCES algorithm provides good empirical estimates of the true risk and the true unfairness measure. Finally, we show that the same holds for the union of these sets,  $S$ , which implies the inequalities in the theorem. For the risk, the last step is a straightforward consequence of the second. For the fairness, which is not simply an expectation or average over per-sample contributions, a more careful derivation is needed that crucially uses the disbalance measure as well. For details of the steps, please see Appendix D.

## 4 Experiments

FLEA’s claim is that it allows learning classifiers that are fair even in the presence of perturbations in the training data. Due to its filtering approach it can be used in combination with any existing learning method. For our experiments, we run it in combination with three fairness-aware learning methods as well as



one fairness-unaware one against a variety of adversaries on five established fair classification datasets. We benchmark our method against the corresponding base learning algorithms without pre-filtering, as well as against three robust learning baselines from the literature.

#### 4.1 Experimental setup

We report experiments in two setups: for *homogeneous* and *heterogeneous* data sources.

**Datasets** For the homogeneous setup we use four standard benchmark datasets from the fair classification literature: COMPAS (Aingwin et al., 2016) (6171 examples), *adult* (48841), *germancredit*(1000) and *drugs* (1885) (Dua & Graff, 2017). To obtain multiple identically distributed sources, we randomly split each training set into  $N \in \{3, 5, 7, 9, 11\}$  equal-sized parts, out of which the adversary can manipulate  $\lfloor \frac{N-1}{2} \rfloor$ . For the heterogeneous case we use the 2018 US census data of the *folktables* dataset (Ding et al., 2021). We form 51 similarly but not identically distributed data sources by using up to 10000 examples from each of the US-states. Out of these 5, 10, 15, 20 or 25 can be manipulated. Details about the data preprocessing and feature extraction steps can be found in the supplemental material.

In all cases, we use *gender* as exemplary protected attribute, because it is present in all feature sets. We train linear classifiers by logistic regression without regularization, using 80% of the data for training and the remaining 20% for evaluation. All experiments are repeated ten times with different train-test splits and random seeds. We measure the mean and standard deviation of the accuracy and the fairness of the learned classifiers, where we compute fairness as  $1 - \Gamma_S$ , where  $\Gamma_S$  is the demographic parity violation on the test set.

**Fairness-Aware Learners** We use FLEA in combination with four fairness-aware learning methods that have found wide adoption in research and practice. In all cases, we use *logistic regression* as the underlying classification model.

- *Fairness regularization* (Kamishima et al., 2012) learns a fair classifier by minimizing a linear combination of the classification loss and the empirical unfairness measure  $\Gamma_S$ , where for numeric stability, in the latter the binary-valued classifier decisions  $f(x)$  are replaced by the real-valued confidences  $p(f(x) = 1|x)$ .
- *Data preprocessing* (Kamiran & Calders, 2012) modifies the training data to remove potential biases. Specifically, it creates a new dataset by *uniform resampling* (with repetition) from the original dataset, such that the fractions of positive and negative labels are the same for each protected group. On the resulting unbiased dataset it trains an ordinary fairness-unaware classifier.
- *Score postprocessing* (Hardt et al., 2016) first learns an ordinary (fairness-unaware) classifier on the available data. Afterwards, it determines which decision thresholds for each protected groups achieve (approximate) demographic parity on the training set, finally picking the fair thresholds with highest training accuracy.
- *Adversarial fairness* (Wadsworth et al., 2018) learns by minimizing a weighted difference between two terms. One is the loss of the actual classifier; the other is the loss of a classifier that tries to predict the protected attribute from the real-valued outputs of the main classifier.

For completeness, we also include plain logistic regression as a *fairness-unaware learner*. The supplemental material details the learners’ implementations and parameters.

**Adversaries** In a real-world setting, one does not know what kind of data quality issues will occur. Therefore, we test the baselines and FLEA for a range of adversaries that reflect potentially unintentional errors as well as intentional manipulations.

- *flip protected (FP)*, *flip label (FL)*, *flip both (FB)*: the adversary flips the value of protected attribute, of the label, or both, in all sources it can manipulate.
- *shuffle protected (SP)*: the adversary shuffles the protected attribute entry in each effected batch.
- *overwrite protected (OP)*, *overwrite label (OL)*: the adversary overwrites the protected attribute of each sample in the affected batch by its label, or vice versa.
- *resample protected (RP)*: the adversary samples new batches of data in the following ways: all original samples of protected group  $a = 0$  with labels  $y = 1$  are replaced by data samples from other sources which

also have  $a = 0$  but  $y = 0$ . Analogously, all samples of group  $a = 1$  with labels  $y = 0$  are replaced by data samples from other sources with  $a = 1$  and  $y = 1$ .

- *random (RND)*: the adversary randomly picks one of the strategies above for each source.
- *identity (ID)*: the adversary makes no changes to the data.

We include ID to certify that FLEA does not unnecessarily damage the learning process in the case when the training data is actually clean. The other adversaries either weaken the correlations between the protected attribute and the target data, thereby masking a potential existing bias in the data, or they strengthen the correlation between the protected attribute and the target label, thereby increasing the chance that the learned classifier will use the protected attribute as a basis for its decisions. In both cases, the dataset statistics at training time will differ from the situation at test time, and the efficacy of a potential mechanisms to ensure fairness at training time can be expected to suffer. For a more detailed discussion of the adversaries' effects, please see the supplemental material.

As our result in the following section will show, even these relatively simple adversaries are highly effective in disrupting the fairness-aware learning process. More complex attack strategies have also been proposed in the literature. E.g., similar to Solans et al. (2020), for real-valued input data an adversary could run a gradient-based optimization in order to construct data with maximally adverse effect on a given unfairness measure. This would be difficult in the setting of our experiments, where the training data is mostly categorical. It is also not necessary, though: for the linear classifiers that we consider, the effect of the data manipulations can be assessed explicitly, as we do above. There is no need for an attacker to differentiate through the model.

**Baselines** To the best of our knowledge, FLEA is the only existing method to tackle fair learning under arbitrary data manipulations. To nevertheless put our results into context, we compare it to three baselines: 1) a *robust ensemble* (similar to Smith & Martinez (2018)), which learns separate classifiers on each datasource and then combines their decisions by a majority vote. 2) A *distributionally robust optimization (DRO)* approach as proposed in Wang et al. (2020) to address noisy protected attributes. 3) The *filtering approach* of Konstantinov et al. (2020) which uses discrepancy to identify manipulated sources but does not specifically aim to preserve fairness. More details on these can be found in the supplemental material. Further candidates could be Roh et al. (2020); Konstantinov & Lampert (2019), but these are not applicable in our setting, as they require access to guaranteed clean validation data.

## 4.2 Results

The results of our experiments show a very consistent picture across different datasets, base learners and adversaries. For the sake of conciseness, we present only results for the regularization-based fairness-aware learner in the main manuscript. Results for other learners are qualitatively the same and can be found, together with more detailed results and ablation studies, in the supplemental material.

In Table 1 we report results for five learning methods: an ordinary learner that is fairness-aware but not protected against data manipulations (naive), the proposed FLEA, and the three baseline methods: the robust ensemble, DRO (adapted from Wang et al. (2020)), and discrepancy-based filtering Konstantinov et al. (2020). In addition, we report the value of a hypothetical oracle-based learner that knows which of the sources are actually clean and learns only on their data.

Each entry in the table is the *minimum* accuracy and fairness in the respective setting across all nine tested adversaries. We choose this worst-case measure because it allows a compact representation and reflects the fact that a real-world system should be robust against all possible data errors or manipulation simultaneously. Results broken down by individual adversaries are provided in the supplemental material.

First, a comparison of the *naive* results with the *oracle* confirms that the need for robust learning method is real: naive fairness-aware learning is not sufficient to ensure fair (or accurate) classifiers in the presence of unreliable data.

Table 1: Result of FLEA and baselines for robust fairness-aware multisource learning with homogeneous and heterogeneous data sources. Reported accuracy and fairness values are the minimal (worst-case) ones across all tested data manipulations in the respective settings. See main text for an explanation of the methods and details of the experimental setup.

(a) homogeneous: **adult**, **COMPAS**, **drugs** and **germancredit** datasets with 5 sources of which 2 are unreliable.

method	adult		COMPAS		drugs		germancredit	
	accuracy	fairness	accuracy	fairness	accuracy	fairness	accuracy	fairness
naive	66.2 $\pm$ 1.1	77.6 $\pm$ 1.2	63.5 $\pm$ 2.1	78.9 $\pm$ 2.3	60.1 $\pm$ 2.5	72.3 $\pm$ 3.1	58.1 $\pm$ 4.0	81.9 $\pm$ 5.0
robust ensemble	69.9 $\pm$ 0.4	90.9 $\pm$ 1.6	65.0 $\pm$ 1.1	88.4 $\pm$ 2.9	61.0 $\pm$ 2.1	75.7 $\pm$ 5.2	62.1 $\pm$ 2.9	71.2 $\pm$ 7.0
DRO (Wang et al., 2020)	52.8 $\pm$ 0.3	15.4 $\pm$ 1.3	54.5 $\pm$ 1.2	70.9 $\pm$ 5.7	55.1 $\pm$ 2.5	53.9 $\pm$ 23.3	65.0 $\pm$ 4.9	78.4 $\pm$ 13.1
(Konstantinov et al., 2020)	69.3 $\pm$ 0.4	77.6 $\pm$ 1.2	63.5 $\pm$ 2.1	78.9 $\pm$ 2.3	60.1 $\pm$ 2.5	72.3 $\pm$ 3.1	58.1 $\pm$ 4.0	81.9 $\pm$ 5.0
FLEA (proposed)	70.3 $\pm$ 0.4	98.1 $\pm$ 1.2	65.9 $\pm$ 1.1	95.3 $\pm$ 2.3	64.1 $\pm$ 1.3	91.7 $\pm$ 4.1	65.8 $\pm$ 3.2	92.6 $\pm$ 4.1
oracle	70.3 $\pm$ 0.4	98.2 $\pm$ 1.0	66.2 $\pm$ 1.1	96.2 $\pm$ 1.3	64.4 $\pm$ 1.5	93.6 $\pm$ 3.3	67.3 $\pm$ 3.0	94.4 $\pm$ 4.0

(b) heterogeneous: **folktables** dataset with  $N = 51$  sources of which  $N - K \in \{5, 10, 15, 20, 25\}$  are unreliable.

method	$N - K = 5$		$N - K = 10$		$N - K = 15$		$N - K = 20$		$N - K = 25$	
	accuracy	fairness	accuracy	fairness	accuracy	fairness	accuracy	fairness	accuracy	fairness
naive	74.4 $\pm$ 0.2	93.4 $\pm$ 0.8	73.8 $\pm$ 0.2	87.0 $\pm$ 0.8	72.9 $\pm$ 0.5	80.1 $\pm$ 0.9	71.2 $\pm$ 0.8	73.3 $\pm$ 0.6	58.2 $\pm$ 6.2	73.9 $\pm$ 1.0
robust ensemble	74.9 $\pm$ 0.2	97.3 $\pm$ 0.4	74.3 $\pm$ 0.2	93.8 $\pm$ 0.4	73.5 $\pm$ 0.3	89.0 $\pm$ 0.4	71.9 $\pm$ 0.3	81.7 $\pm$ 0.7	65.9 $\pm$ 1.0	60.4 $\pm$ 2.2
DRO (Wang et al., 2020)	65.2 $\pm$ 0.8	96.3 $\pm$ 0.4	68.6 $\pm$ 1.2	95.4 $\pm$ 0.9	67.4 $\pm$ 1.6	85.4 $\pm$ 2.3	66.1 $\pm$ 1.2	77.8 $\pm$ 12.7	58.2 $\pm$ 5.6	6.9 $\pm$ 8.4
(Konstantinov et al., 2020)	74.4 $\pm$ 0.2	93.4 $\pm$ 0.8	73.8 $\pm$ 0.2	87.0 $\pm$ 0.8	72.9 $\pm$ 0.5	80.1 $\pm$ 0.9	71.2 $\pm$ 0.8	73.3 $\pm$ 0.6	58.2 $\pm$ 6.2	73.9 $\pm$ 1.0
FLEA (proposed)	75.4 $\pm$ 0.2	99.5 $\pm$ 0.2	75.4 $\pm$ 0.2	99.4 $\pm$ 0.3	75.3 $\pm$ 0.3	99.3 $\pm$ 0.3	75.3 $\pm$ 0.2	98.3 $\pm$ 1.1	66.4 $\pm$ 11.5	86.9 $\pm$ 14.1
oracle	75.2 $\pm$ 0.2	99.5 $\pm$ 0.3	75.2 $\pm$ 0.2	99.6 $\pm$ 0.2	75.3 $\pm$ 0.2	99.7 $\pm$ 0.2	75.3 $\pm$ 0.2	99.7 $\pm$ 0.2	75.1 $\pm$ 0.3	99.6 $\pm$ 0.4

An ideal robust method should achieve results approximately as good as the *oracle* result, as this would indicate that the adversary was indeed not able to negatively affect the learning process beyond the unavoidable loss of some training data. The results show that FLEA comes close to this behavior, but none of the other methods does. In the homogeneous setting (Table 1a), for the largest dataset, **adult**, FLEA reliably suppresses the effects of all tested adversaries. It learns classifiers with accuracy and fairness almost exactly those of a fair classifier trained only on the clean data sources. For the other datasets, **COMPAS**, **drugs** and **germancredit**, FLEA increases the accuracy and fairness to levels only slightly below the oracle. In all cases, FLEA’s results are as good as or better than the baselines; the robust ensemble is also able to improve fairness to some extent, but it does not reach the oracle results. The DRO-based approach shows highly volatile behavior. For some adversaries it improves fairness or accuracy, but for some adversaries it fail severely. Consequently, the min-aggregated values in the table are often even lower than the naive method.<sup>4</sup> The approach from Konstantinov et al. (2020) has almost no effect. Only for the largest dataset, **adult**, it yields a slight accuracy improvement. This can be explained by the fact that the method only removes sources that it can confidently identify as manipulated. The theory-derived thresholds for this are quite strict, so the method is ineffective unless a lot of data is available. The observed characteristics of the different methods hold also for other number of data sources and for the other base learners, see the supplemental material.

In the heterogeneous setting (Table 1b) the results show similar trends: for  $N - K \in \{5, 10, 15, 20\}$ , FLEA manages reliably to filter out the malignant sources, such that the accuracy and fairness of the learned classifiers matches the one of the oracle method almost perfectly. The robust ensemble has a positive effect, but less so than FLEA. For this data, DRO somewhat improves fairness, but this comes at a loss of accuracy. The method from Konstantinov et al. (2020) has no noticeable effect. For  $N - K = 25$ , FLEA still performs best, although not close to the quality of the hypothetical best oracle. Presumably, this is because the combined effect of distribution differences between the sources and the uncertainty due to finite sampling when estimating disc, disp and disb are too large to perfectly allow a decision which 26 sources to keep and which 25 to exclude.

<sup>4</sup>Note that these results should not negatively reflect on Wang et al. (2020), as our data manipulation model is different and substantially harder than what the DRO method was designed for.

## 5 Conclusion

We studied the problem of fairness-aware classification in the setting when data from multiple sources is available, but some of them might be noisy, contain unintentional errors, or even have been maliciously manipulated. Ordinary fairness-aware learning methods are not robust against such problem and often fail to produce fair classifiers. We proposed a filtering-based algorithm, FLEA, that is able to identify and suppress those data sources that would negatively affect the training process, thereby restoring the property that fairness-aware learning methods actually produce fair classifiers. We showed the effectiveness of FLEA experimentally, and we also presented a theorem that provides formal guarantees of FLEA’s efficacy.

Despite our promising results, we consider FLEA just a first step on the path toward making fairness-aware learning more robust. One potential future step is to include other notions of fairness besides *demographic parity*. So far, FLEA can already be used as it is with classifiers that enforce other fairness criteria. However, our theoretical guarantees do not hold for these, as the *disparity* measure that enters our filtering step is not tailored to them. We do not see fundamental problems in deriving filtering steps for other fairness notions that are also defined in terms of properties of the joint distribution of inputs, outputs, and protected attributes, such as *equality of opportunity* or *equalized odds*. However, the theoretical analysis and the practical implementation could get more involved.

On the algorithmic side, FLEA as we formulated it requires computing all pairwise similarities between the sources. This could render it inefficient when the number of sources is very large (e.g. thousands). We expect that it will be possible to overcome this, for example by randomization of the sources, but we leave this step to future work.

## References

- Agarwal, A., Beygelzimer, A., Dudik, M., Langford, J., and Wallach, H. A reductions approach to fair classification. In *International Conference on Machine Learning (ICML)*, 2018.
- Aingwin, J., Larson, J., Mattu, S., and Kirchner, L. Machine bias: There’s software used across the country to predict future criminals and its biased against blacks., 2016. URL <https://github.com/propublica/compas-analysis>.
- Barnett, V. and Lewis, T. *Outliers in statistical data*. Wiley, 1984.
- Barocas, S., Hardt, M., and Narayanan, A. *Fairness and Machine Learning*. fairmlbook.org, 2019.
- Bartlett, P. L. Learning with a slowly changing distribution. In *Workshop on Computational Learning Theory (COLT)*, pp. 243–252, 1992.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *Machine Learning (ML)*, 2010.
- Beutel, A., Chen, J., Zhao, Z., and Chi, E. H. Data decisions and theoretical implications when adversarially learning fair representations. In *Conference on Fairness, Accountability and Transparency (FAccT)*, 2017.
- Calders, T., Kamiran, F., and Pechenizkiy, M. Building classifiers with independency constraints. In *International Conference on Data Mining Workshops (IDCMW)*, 2009.
- Calmon, F., Wei, D., Vinzamuri, B., Natesan Ramamurthy, K., and Varshney, K. R. Optimized pre-processing for discrimination prevention. *Advances in neural information processing systems*, 30, 2017.
- Celis, L. E., Keswani, V., and Vishnoi, N. Data preprocessing to mitigate bias: A maximum entropy based approach. In *International Conference on Machine Learning*, pp. 1349–1359. PMLR, 2020.
- Celis, L. E., Huang, L., Keswani, V., and Vishnoi, N. K. Fair classification with noisy protected attributes: A framework with provable guarantees. In *International Conference on Machine Learning (ICML)*, 2021a.
- Celis, L. E., Mehrotra, A., and Vishnoi, N. K. Fair classification with adversarial perturbations. *Conference on Neural Information Processing Systems (NeurIPS)*, 2021b.

- Charikar, M., Steinhardt, J., and Valiant, G. Learning from untrusted data. In *Symposium on Theory of Computing (STOC)*, 2017.
- Chen, S., Li, J., and Moitra, A. Efficiently learning structured distributions from untrusted batches. In *Symposium on Theory of Computing (STOC)*, 2019.
- Cho, J., Hwang, G., and Suh, C. A fair classifier using kernel density estimation. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Chuang, C.-Y. and Mroueh, Y. Fair mixup: Fairness via interpolation. In *International Conference on Learning Representations (ICLR)*, 2021.
- Chzhen, E., Denis, C., Hebiri, M., Oneto, L., and Pontil, M. Fair regression with Wasserstein barycenters. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Corder, G. W. and Foreman, D. I. *Nonparametric statistics: A step-by-step approach*. John Wiley & Sons, 2014.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. *Introduction to algorithms*. The MIT Press, 2009.
- Crammer, K., Kearns, M., and Wortman, J. Learning from multiple sources. *Journal of Machine Learning Research (JMLR)*, 2008.
- Ding, F., Hardt, M., Miller, J., and Schmidt, L. Retiring adult: New datasets for fair machine learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J. S., and Pontil, M. Empirical risk minimization under fairness constraints. In *Conference on Neural Information Processing Systems (NeurIPS)*, volume 31, 2018.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In *Innovations in Theoretical Computer Science Conference (ITCS)*, 2012.
- Erfani, S., Baktashmotlagh, M., Moshtaghi, M., Nguyen, V., Leckie, C., Bailey, J., and Ramamohanarao, K. From shared subspaces to shared landmarks: A robust multi-source classification approach. In *Conference on Artificial Intelligence (AAAI)*, 2017.
- Feng, J., Cai, Q.-Z., and Zhou, Z.-H. Learning to confuse: Generating training time adversarial data with auto-encoder. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- Fowl, L., Goldblum, M., Chiang, P.-y., Geiping, J., Czaja, W., and Goldstein, T. Adversarial examples make strong poisons. *arXiv preprint arXiv:2106.10807*, 2021.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2014.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *Journal of Machine Learning Research (JMLR)*, 13(25):723–773, 2012.
- Hanneke, S. and Kpotufe, S. A no-free-lunch theorem for multitask learning. *arXiv preprint arXiv:2006.15785*, 2020.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2016.

- Jain, A. and Orlitsky, A. Optimal robust learning of discrete distributions from batches. In *International Conference on Machine Learning (ICML)*, 2020a.
- Jain, A. and Orlitsky, A. A general method for robust learning from batches. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020b.
- Kallus, N., Mao, X., and Zhou, A. Assessing algorithmic fairness with unobserved protected class using data combination. In *Conference on Fairness, Accountability and Transparency (FAccT)*, 2020.
- Kamiran, F. and Calders, T. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems (KAIS)*, 33(1):1–33, 2012.
- Kamishima, T., Akaho, S., Asoh, H., and Sakuma, J. Fairness-aware classifier with prejudice remover regularizer. In *European Conference on Machine Learning and Data Mining (ECML PKDD)*, 2012.
- Kearns, M. and Li, M. Learning in the presence of malicious errors. *SIAM Journal on Computing*, 1993.
- Kifer, D., Ben-David, S., and Gehrke, J. Detecting change in data streams. In *International Conference on Very Large Data Bases (VLDB)*, 2004.
- Konstantinov, N. and Lampert, C. H. Robust learning from untrusted sources. In *International Conference on Machine Learning (ICML)*, 2019.
- Konstantinov, N. and Lampert, C. H. Fairness-aware learning from corrupted data. *Journal of Machine Learning Research (JMLR)*, 2022.
- Konstantinov, N., Frantar, E., Alistarh, D., and Lampert, C. On the sample complexity of adversarial multi-source PAC learning. In *International Conference on Machine Learning (ICML)*, 2020.
- Kullback, S. and Leibler, R. A. On information and sufficiency. *The annals of mathematical statistics*, 22(1): 79–86, 1951.
- Lahoti, P., Beutel, A., Chen, J., Lee, K., Prost, F., Thain, N., Wang, X., and Chi, E. Fairness without demographics through adversarially reweighted learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Lamy, A., Zhong, Z., Menon, A. K., and Verma, N. Noise-tolerant fair classification. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- Li, T., Hu, S., Beirami, A., and Smith, V. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning (ICML)*, 2021.
- Mandal, D., Deng, S., Jana, S., Wing, J., and Hsu, D. J. Ensuring fairness beyond the training data. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Marcotte, P. and Savard, G. Novel approaches to the discrimination problem. *Zeitschrift für Operations Research*, 1992.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021a.
- Mehrabi, N., Naveed, M., Morstatter, F., and Galstyan, A. Exacerbating algorithmic bias through fairness attacks. In *Conference on Artificial Intelligence (AAAI)*, 2021b.
- Mehrotra, A. and Celis, L. E. Mitigating bias in set selection with noisy protected attributes. In *Conference on Fairness, Accountability and Transparency (FAccT)*, 2021.
- Mohri, M. and Medina, A. M. New analysis and algorithm for learning with drifting distributions. In *Algorithmic Learning Theory (ALT)*, 2012.
- Natarajan, N., Dhillon, I. S., Ravikumar, P., and Tewari, A. Learning with noisy labels. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2013.

- Qiao, M. and Valiant, G. Learning discrete distributions from untrusted batches. In *Innovations in Theoretical Computer Science Conference (ITCS)*, 2018.
- Rezaei, A., Fathony, R., Memarrast, O., and Ziebart, B. Fairness for robust log loss classification. In *Conference on Artificial Intelligence (AAAI)*, 2020.
- Roh, Y., Lee, K., Whang, S., and Suh, C. FR-train: A mutual information-based approach to fair and robust training. In *International Conference on Machine Learning (ICML)*, 2020.
- Roh, Y., Lee, K., Whang, S., and Suh, C. Sample selection for fair and robust training. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- Shalev-Shwartz, S. and Ben-David, S. *Understanding machine learning: from theory to algorithms*. Cambridge University Press, 2014.
- Smith, M. R. and Martinez, T. R. The robustness of majority voting compared to filtering misclassified instances in supervised classification tasks. *Artificial Intelligence Review*, 49(1):105–130, 2018.
- Solans, D., Biggio, B., and Castillo, C. Poisoning attacks on algorithmic fairness. *European Conference on Machine Learning and Data Mining (ECML PKDD)*, 2020.
- Student. The probable error of a mean. *Biometrika*, pp. 1–25, 1908.
- Tan, Z., Yeom, S., Fredrikson, M., and Talwalkar, A. Learning fair representations for kernel models. In *Conference on Uncertainty in Artificial Intelligence (AISTATS)*, 2020.
- Ting, K. M. and Low, B. T. Model combination in the multiple-data-batches scenario. In *European Conference on Machine Learning (ECML)*, 1997.
- Tsybakov, A. B. *Introduction to Nonparametric Estimation*. Springer series in statistics. Springer, 2009.
- Vapnik, V. *The nature of statistical learning theory*. Springer Science & Business Media, 2013.
- Villani, C. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Wadsworth, C., Vera, F., and Piech, C. Achieving fairness through adversarial learning: an application to recidivism prediction. In *Conference on Fairness, Accountability and Transparency (FAccT)*, 2018.
- Wang, H., Ustun, B., and Calmon, F. Repairing without retraining: Avoiding disparate impact with counterfactual distributions. In *International Conference on Machine Learning (ICML)*, 2019.
- Wang, S., Guo, W., Narasimhan, H., Cotter, A., Gupta, M., and Jordan, M. Robust optimization for fairness with noisy protected groups. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Woodworth, B., Gunasekar, S., Ohannessian, M. I., and Srebro, N. Learning non-discriminatory predictors. In *Workshop on Computational Learning Theory (COLT)*, 2017.
- Zafar, M. B., Valera, I., Gomez Rodriguez, M., and Gummadi, K. P. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *International World Wide Web Conference (WWW)*, 2017a.
- Zafar, M. B., Valera, I., Rogriguez, M. G., and Gummadi, K. P. Fairness constraints: Mechanisms for fair classification. In *Conference on Uncertainty in Artificial Intelligence (AISTATS)*, 2017b.
- Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. Learning fair representations. In *International Conference on Machine Learning (ICML)*, 2013.
- Zhang, B. H., Lemoine, B., and Mitchell, M. Mitigating unwanted biases with adversarial learning. In *Conference on AI, Ethics, and Society (AIES)*, 2018.

# Appendix

## Table of Contents

- **A: Experimental Setup**
- **B: Detailed Experimental Results**
- **C: Discussion of the Role of `disb`, `disc` and `disp` and Ablation Study**
- **D: Complete Formulation and Proof of Theorem 1**

## A Experimental setup

### A.1 Dataset preparation

The datasets we use are publicly available and frequently used to evaluate fair classification methods.

The COMPAS dataset was introduced by ProPublica. It contains data from the US criminal justice system and was obtained by a public records request. The dataset contains personal information. To mitigate negative side effects, we delete the *name*, *first*, *last* and *dob* (date of birth) entries from the dataset before processing it further. We then exclude entries that do not fit the problem setting of predicting two year recidivism, following the steps of the original analysis.<sup>5</sup> Specifically, this means keeping only cases from Broward county, Florida, for which data has been entered within 30 days of the arrest. Traffic offenses and cases with insufficient information are also excluded. This steps leave 6171 examples out of the original 7214 cases. The categorical features and numerical features that we extract from the data are provided in Table 3a.

`adult`, `germancredit`, and `drugs` are available in the UCI data repository as well as multiple other online sources.<sup>6</sup> We use them in unmodified form, except for binning some of the feature values; see Tables 2 and 3.

Table 2: Dataset information

(a) <code>adult</code>		
dataset size	48842	
categorical features	<i>workclass</i> <i>education</i> <i>hours-per-week</i> <i>age</i> <i>native-country</i> <i>race</i>	federal-gov, local-gov, never-worked, private, self-emp-inc, self-emp-not-inc, state-gov, without-pay, unknown 1st-4th, 5th-6th, 7th-8th, 9th, 10th, 11th, 12th, Assoc-acdm, Assoc-voc, Bachelors, Doctorate, HS-grad, Masters, Preschool, Prof-school, Some-college $\leq 19$ , 20–29, 30–39, $\geq 40$ $\leq 24$ , 25–34, 35–44, 45–54, 55–64, $\geq 65$ United States, other Amer-Indian-Eskimo, Asian-Pac-Islander, Black, White, other
numerical features	—	
protected attribute	<i>gender</i>	values: female (33.2%), male (66.8%)
target variable	<i>income</i>	$\leq 50K$ (76.1%), $> 50K$ (33.9%)

<sup>5</sup><https://github.com/propublica/compas-analysis>

<sup>6</sup>`adult`: <https://archive.ics.uci.edu/ml/datasets/adult>,  
`germancredit`: [https://github.com/praisan/hello-world/blob/master/german\\_credit\\_data.csv](https://github.com/praisan/hello-world/blob/master/german_credit_data.csv),  
`drugs`: [https://raw.githubusercontent.com/deepak525/Drug-Consumption/master/drug\\_consumption.csv](https://raw.githubusercontent.com/deepak525/Drug-Consumption/master/drug_consumption.csv)



Table 3: Dataset information (continued)

## (a) COMPAS

dataset size	6171 (7214 before filtering)	
categorical features	<i>c-charge-degree</i> <i>age-cat</i> <i>race</i>	values: F (felony), M (misconduct) values: <25, 25–45, >45 values: African-American, Caucasian, Hispanic, Other
numerical features	<i>priors-count</i>	
protected attribute	<i>sex</i>	Female (19.0%), Male (81.0%)
target variable	<i>two-year-recid</i>	0 (54.9%), 1 (45.1%)

## (b) drugs

dataset size	1885	
categorical features	—	
numerical features	<i>Age, Gender, Education, Country, Ethnicity, Nscore, Escore, Oscore, Ascore, Cscore, Impulsive, SS</i>	(precomputed numeric values in dataset)
protected attribute	<i>Gender</i>	female (31.0%), male (69.0%)
target variable	<i>Coke</i>	never used (55.1%), used (44.9%)

## (c) germancredit

dataset size	1000	
categorical features	<i>Age</i> <i>Saving accounts</i> <i>Checking account</i>	values: $\leq 24$ , 25–34, 35–44, 45–54, 55–64, $\geq 65$ little, moderate, quite rich, rich little, moderate, rich
numerical features	<i>Duration, Credit amount</i>	
protected attribute	<i>Sex</i>	female (31.0%), male (69.0%)
target variable	<i>Risk</i>	bad (30%), good (70%)

## (d) folktables

dataset size	255078	
categorical features	<i>AGE</i> (age; binned) <i>COW</i> (class of worker) <i>SCHL</i> (education) <i>MAR</i> (marital status)  <i>OCCP</i> (occupation code) <i>POBP</i> (place of birth) <i>RELHP</i> (relationship in household) <i>WKHP</i> (weekly working hours; binned) <i>RAC1P</i> (race code)	values: $\leq 14$ , 15–24, 25–34, 35–44, 45–54, 55–64, $\geq 65$ values: 1, ..., 9 values: 1, ..., 24 values: married, widowed, divorced, separated, never married values: 0, 1, ..., 9 values: USA, other values: 0, 1, ..., 17 values: $\leq 19$ , 20–29, 30–39, $\geq 40$ values: 1, ..., 9
numerical features	—	
protected attribute	<i>SEX</i>	female (52.1%), male (47.9%)
target variable	<i>income</i>	$\leq 50K$ (64.8%), $> 50K$ (35.2%)

For details of the numeric codes, see [https://www2.census.gov/programs-surveys/acs/tech\\_docs/pums/data\\_dict/PUMS\\_Data\\_Dictionary\\_2018.pdf](https://www2.census.gov/programs-surveys/acs/tech_docs/pums/data_dict/PUMS_Data_Dictionary_2018.pdf)

## A.2 Training objectives

All training objectives are derived from logistic regression classifiers. For data  $S = \{(x_1, y_1), \dots, (x_n, y_n)\} \subset \mathbb{R}^d \times \{\pm 1\}$  we learn a prediction function  $g(x) = w^\top x + b$  by solving

$$\min_{w \in \mathbb{R}^d, b \in \mathbb{R}} \mathcal{L}_S(w, b) + \lambda \|w\|^2 \quad (8)$$

with

$$\mathcal{L}_S(w, b) = \frac{1}{|S|} \sum_{(x, y) \in S} y \log(1 + e^{-g(x)}) + (1 - y) \log(1 + e^{g(x)}) \quad (9)$$

We use the `LogisticRegression` routine of the `sklearn` package for this, which runs a LBFGS optimizer for up to 500 iterations. By default, we do not use a regularizer, i.e.  $\lambda = 0$ . From  $g(x)$  we obtain classification decisions as  $f(x) = \text{sign } g(x)$  and probability estimates as  $\sigma(x; w, b) = p(y = 1|x) = \frac{1}{1 + e^{-g(x)}}$ , where we clip the output of  $g$  to the interval  $[-20, 20]$  to avoid numeric issues.

To train with fairness regularization, we solve the optimization problem

$$\min_{w \in \mathbb{R}^d, b \in \mathbb{R}} \mathcal{L}_S(w, b) + \eta \Gamma_S(w, b) \quad (10)$$

with

$$\Gamma_S(w, b) = \left| \frac{1}{|S^{a=0}|} \sum_{x \in S^{a=0}} \sigma(x; w, b) - \frac{1}{|S^{a=1}|} \sum_{x \in S^{a=1}} \sigma(x; w, b) \right|_\epsilon \quad (11)$$

where for reasons of numeric stability, we use  $|t|_\epsilon = \sqrt{\frac{t^2}{t^2 + \epsilon}}$  with  $\epsilon = 10^{-8}$ . To do so, we use the `scipy.minimize` routine with "BFGS" optimizer for up to 500 iterations. The necessary gradients are computed automatically using `jax`.<sup>7</sup> To initialize  $(w, b)$ , we use the result of training a (fairness-unaware) logistic regression with  $\lambda = 1$ , where the regularization is meant to ensure that the parameters do not take too extreme values. When estimating the disparity, we use the same objective, but with different datasets,  $S_1, S_2$  for the two terms in (10), with the protected attributes as target labels for  $S_1$ , and the inverse of the protected attributes as target labels for  $S_2$ .

To train with adversarial regularization, we parameterize an adversary  $g' : \mathbb{R} \rightarrow \mathbb{R}$  as  $g'(x') = w'x' + b'$  and solve the optimization problem

$$\min_{w \in \mathbb{R}^d, b \in \mathbb{R}} \max_{w' \in \mathbb{R}, b' \in \mathbb{R}} \mathcal{L}_S(w, b) - \eta \mathcal{L}'_S(w', b') \quad (12)$$

with

$$\mathcal{L}'_S(w, b, w', b') = \frac{1}{|S|} \sum_{(x, a) \in S} a \log(1 + e^{-g'(g(x))}) + (1 - a) \log(1 + e^{g'(g(x))}) \quad (13)$$

To do so, we use the `optax` package with gradient updates by the Adam rule for up to 1000 steps. The learning rates for classifier and adversary are 0.001. The gradients are again computed using `jax`. We initialize  $(w, b)$  the same way as for (10).  $(w', b')$  we simply initialize with zeros.

To perform score postprocessing, we evaluate the linear prediction function on the training set and determine the thresholds that result in a fraction of  $r \in \{0, 0.01, \dots, 0.99, 1\}$  positive decision separately for each protected group. For each  $r$  we then compute the overall accuracy of the classifier that results from using these group-specific thresholds and select the value for  $r$  that leads to the highest accuracy. We then modify the classifier to use the corresponding thresholds for each group by adjusting the classifier weights of the protected attributes.

<sup>7</sup><https://github.com/google/jax> (version 0.2.12)

### A.3 Baselines

In this section, we provide more details about the baselines.

**Robust ensemble** For this baseline, we train  $N$  classifiers, one per data source, using the respective base learner. For prediction, we compute the median value of the predicted probabilities and threshold it at 0.5 to obtain a binary label. Since in our experiments the number of sources is always odd, this is also equivalent to classifying using the majority vote rule.

**Filtering method from Konstantinov et al. (2020)** The method proposed in Konstantinov et al. (2020) uses a filtering step to suppress unreliable sources, like we do, but that differs from FLEA’s in two main aspects: it uses only the discrepancy score for its decisions, and its decision criterion is threshold-based, not quantile-based.

For its implementation, one first computes the pairwise discrepancy scores,  $\text{disc}(S_i, S_j)$ , between all sources. Then, one determines a threshold,  $t = \sqrt{\frac{8d \log(2en/d) + 8 \log(8N/\delta)}{n}}$ , where  $d$  is the VC dimension of the hypothesis class (for us: the dimensionality of the feature vectors plus 1).  $\delta$  is a freely choosable confidence parameter. In the limited data regime of our experiments, its value has little influence on the threshold, so we leave it at a default of  $\delta = 0.1$ . Finally, for each source,  $S_i$ , we check for how many other sources,  $S_j$ , their pairwise discrepancy to  $S_i$  is less than  $t$  (i.e.  $\sum_{j \neq i} \mathbb{I}[\text{disc}(S_i, S_j) < t]$ ). If the number of such sources is at least  $K - 1$ , the source  $S_i$  is made part of the overall training set, otherwise is it discarded.

One can check that in the setting of our experiments, only for the `adult` dataset one obtains values for  $t$  substantially below 1. Therefore, only for this dataset, the filtering step can have a non-trivial effect.

**DRO method from Wang et al. (2020)** The DRO method was proposed originally for the *equal opportunity* or *equalized odds* fairness measures. We adapt it to *demographic parity* by imposing constraints on the fraction of positive decisions instead of the true and false positive rates.

Our implementation follows the publicly available github repository,<sup>8</sup> which implements an approximate version of the method described in the publication. The main step is learning a classifier with fairness constraints. This is implemented by deriving a Lagrangian objective and performing simultaneous gradient descent on the classifier parameters and gradient ascent on the Lagrange multipliers. This construction has one hyperparameter,  $\xi$ , the permitted slack up to which the constraints have to be fulfilled. We set this adaptively, starting with a small value  $\xi = 0.01$ , but then doubling  $\xi$  until the optimization results in a non-degenerate solution (i.e. not a constant classifier).

Additionally, the constraint term of the objective is optimized in a distributionally robust (DRO) way. For this, sample weights are introduced, and the Lagrangian term is maximized also with respect to these weights, subject to  $L^1$ -ball constraints around uniform weights, and  $L^1$ -simplex constraints to ensure that the weights encode a discrete probability distribution. Following the original code, we use a projected gradient algorithm for the ball constraint, while the simplex constraint is approximated by implicit renormalization. The DRO also has one hyperparameter,  $s$ , the radius of the  $L^1$ -ball. Following the derivation in the original work, we set this to twice the maximal total variation distance between the data distribution of the protected attribute in the original data and in the manipulated data, which in our case is  $s = 2(1 - \alpha)$ .

Additional hyperparameters are the learning rates of for the classifier parameters, for the Lagrangian multipliers, and for the sample weights. After some initial sanity checks we keep these at the values that worked best in the original publication, which is 0.01 in all three cases.

### A.4 Computing resources

All experiments were run on CPU-only compute servers. For each train/test split of each dataset and each adversary, one experimental run across all learning methods takes between 3 and 18 minutes on two CPU cores, depending on the number of sources, the size of the data sets, and the CPU architecture. The time

<sup>8</sup><https://github.com/wenshuoguo/robust-fairness-code>

needed for each row in the ablation study is similar, except for the `folktables` data, which each took 8-12 hours. The combined time for all reported experiments (5 datasets, 5 values for  $N$  or  $K$ , 10 adversaries, 10 train-test splits, 5 base learners) is approximately 2000 core hours.

For the baselines we are able to reuse many already computed parts. If implemented individually, we'd estimate that the training time for the robust ensemble would be slightly lower than for FLEA (but it is slower at prediction time), while the training time for Konstantinov et al. (2020) and the DRO method from Wang et al. (2020) would be comparable to FLEA's.

## A.5 Hyperparameters

We avoid hyperparameter tuning as far as possible. We do not use  $L^2$ -regularization (hyperparameter  $\lambda$ ) except to create initializers, where we found the value used to hardly matter. For the fairness-regularizer and fairness-adversary we use fixed values of  $\eta = \frac{1}{2}$ . We found these to result in generally fair classifiers for unperturbed data without causing classifiers to degenerate (i.e. become constant). Hence we, did not tune these values on a case-by-case basis. When estimating the disparity, we use  $\eta = 1$  to be consistent with the theory.

As learning rate for the adversarial fairness training,  $\text{lr}_{\text{adv}} = 0.001$  was found by trial and error to ensure convergence at a reasonable speed. Once we identified a reliably working setting, we did not try to tune it further.

## A.6 Adversaries

In this section, we describe the adversaries and their motivation in more detail.

- *flip protected (FP)*: the adversary flips the value of protected attribute.

This is a straightforward attack on fairness. FP inverts the correlation between the protected attribute and the rest of the data. After the sources have been combined, the correlation is therefore weakened, which makes the training data look "less unfair". On the one hand, this can cause fairness-enforcing mechanisms as used, e.g., in postprocessing fairness, to erroneously believe that little or no compensation for dataset unfairness is required. Consequently, the resulting classifier is actually unfair when applied to future unmanipulated data. On the other hand, it is possible that the training process actually learns to ignore the protected attribute during training, because it is uncorrelated with the target labels. This could make the classifier more fair, e.g. when used with fairness-unaware training.

Our detailed experimental results (Fig. 2 –26) show that both of these effect do, in fact, occur. FP typically increases unfairness when regularization-based or postprocessing-based base learners are used, but it has the opposite effect for the fairness-unaware base learner.

- *flip label (FL)*: the adversary flips the value of the label.

This is a straightforward attack on accuracy. Following an analog reasoning as above, FL reduces the correlation between the target label and all other data, which makes it harder for the learner to identify a strong classifier.

Indeed, the experiments shows that the FL adversary often succeeds in reducing the accuracy, while the fairness is relatively unaffected. The adverse effect is small for the large datasets (`adult`, `COMPAS`), and larger for the small ones (`drugs`, `germancredit`), presumably because having more data increases the robustness of the learners against mislabeled data.

- *flip both (FB)*: the adversary flips the value of the protected attribute and the label.

This attack influences fairness and accuracy at the same time. It preserves the correlation between the protected attribute and the labels, but reduces the correlation between these two and all the other features. Consequently, the learned classifier might rely heavily on the protected attribute to predict the label, which would make it maximally unfair, but potentially also less accurate.

Our experiments show that this is, indeed, often the observed effect, though the exact amount depends strongly on the dataset and the base learner.

- *shuffle protected (SP)*: the adversary shuffles the protected attribute entries of each batch it modifies, i.e. each example gets assigned the protected attribute of another example that has been chosen at random (without replacement).

This adversary is similar to FP in that it reduces the overall correlation between the protected attribute and the other data. Its effect is weaker, since it does not explicitly introduce anti-correlation in the manipulated sources. However, its manipulations are less likely to be detected by automatic or manual inspection, since it does not change the marginal statistics of the data, i.e. even after the manipulation, the statistical distribution of each feature dimension, including the protected attribute, is the same as for clean sources.

In experimental results, SP indeed performs similarly to FP for the fairness-aware base learners, and its effect is somewhat weaker for the fairness-unaware base learner.

- *overwrite protected (OP)*: the adversary overwrites the protected attribute of each sample in the affected batch by its label.

This manipulation creates a strong artificial correlation between the protected attribute and the target label. In fact, the maximally unfair classifier that predicts the label directly from the protected attribute will have perfect accuracy on the manipulated data, and still a much higher accuracy than what would be correct on the overall training data. Consequently, the learned classifier might make strong use of the protected attribute, which leads to unfair and potentially incorrect decisions on clean data.

Our experiments show that OP indeed often leads to large increases in unfairness. However, there are also cases where the unfairness is actually reduced, but then typically this is accompanied by loss of accuracy.

- *overwrite label (OL)*: the adversary overwrites the label of each sample in the affected batch by its protected attribute.

Like the OP adversary, this manipulation leads to a perfect correlation between the target labels and the protected attributes. However, it achieves this without changing the marginal distribution of the protected attribute, instead influencing the statistics of the labels. Depending on the specific situation, it might be easier or harder to detect from automatic or manual inspection. OL is also more likely to negatively affect the accuracy, since the classifier will try to predict incorrect labels.

The experiments show that OL indeed almost always reduces the accuracy, while at the same time often increasing unfairness.

- *resample protected (RP)*: the adversary samples new batches of data in the following ways: all original samples of protected group  $a = 0$  with labels  $y = 1$  are replaced by data samples from other sources which also have  $a = 0$ , but  $y = 0$ . Analogously, all samples of group  $a = 1$  with labels  $y = 0$  are replaced by data samples from other sources with  $a = 1$  and  $y = 1$ .

Like OL and OP, RP results in a perfect correlation between protected attributes and labels, thereby facilitating unfairness and reducing accuracy. It does so in a more subtle and harder-to-detect way, however, as it achieves the effect using original data samples.

Indeed, in our experimental results RP influences fairness and accuracy in similar ways as the other two methods.

- *random (RND)*: the adversary randomly picks one of the strategies above (except ID) for each source.

This adversary reflects the observation that different sources might be manipulated in different ways. One reason for this could be that in a real-world system, multiple adversaries exist who manipulate individual data sources without coordinating their actions. Alternatively, there might be just one adversary who manipulates all sources, but chooses to manipulate them in different ways, e.g. to avoid easy detection.

The experimental results show that this strategy does, indeed, work to some extent, with RND often having an effect where some of the other methods do not, but the effect is weaker.

- *identity (ID)*: the adversary makes no changes to the data.

The ID adversary serves as a useful check that FLEA does not damage the learning process in the case that all data is actually clean. It also reflects the fact that even though the adversary has the power to manipulate the data it does not have to. Ideally, the learning method will notice this and achieve even better results in presence of the ID adversary than for the *oracle*.

In the experimental results, this effect is only rarely visible for any method, though.

Note that even though we introduced the adversaries above as intentional manipulations, many of them could also occur accidentally when data from different sources is collected, e.g. as problems during data entering or numeric encoding.

## B Detailed experimental results

In addition to the experiments with a regularization-based base learner that were reported in the main manuscript, we also run experiments with postprocessing-based fairness, adversarial fairness, and fairness-unaware learning. Additionally, besides the situation with  $N = 5$  sources out of which  $K = 3$  are clean, we also ran experiments with  $N = 3$  ( $K = 2$ ),  $N = 7$  ( $K = 4$ ),  $N = 9$  ( $K = 5$ ), and  $N = 11$  ( $K = 6$ ) in order to identify the *breaking point* of FLEA with respect to the size of the data sources. The results are depicted in Fig. 2–21 for the homogeneous setting and in Fig. 22–26 for the heterogeneous setting. Also included are results for two of the baselines, robust ensemble and Konstantinov et al. (2020). The DRO method from Wang et al. (2020) that cannot be combined with arbitrary base learners is reported together with results for the regularization-based base learners, as these are methodologically the closest.

The format of the figures is as follows: for each datasets and method, we report the accuracy and fairness results for different adversaries. Each panel contains 10 bars. The left-most one in each diagram shows the result of the hypothetical *oracle* setting, where the learning algorithm trains only on the clean data sources, i.e. the ones which the adversary cannot modify. The remaining bars correspond to the outcome when different adversaries have perturbed the data. An ideal robust method should achieve results approximately as good as the oracle result, as this would indicate that the adversary was indeed not able to negatively affect the learning process.

From the results, one can see that FLEA works almost perfectly in the homogeneous setting as long as the individual datasets do not become too small. The latter is the case for two smaller datasets, **drugs** and **germancredit**, in the setting with many sources. In that case, for some adversaries FLEA does not always match the *oracle* results, though it still performs better than the baselines. In the heterogeneous case, FLEA works reliably in all settings, except with  $N - K = 25$ , as we had already discussed in the main manuscript.

The results also show that different base learners achieve different accuracy/fairness trade-offs, but FLEA is effective with each of them.






Figure 2: `adult` dataset, regularization-based fairness(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

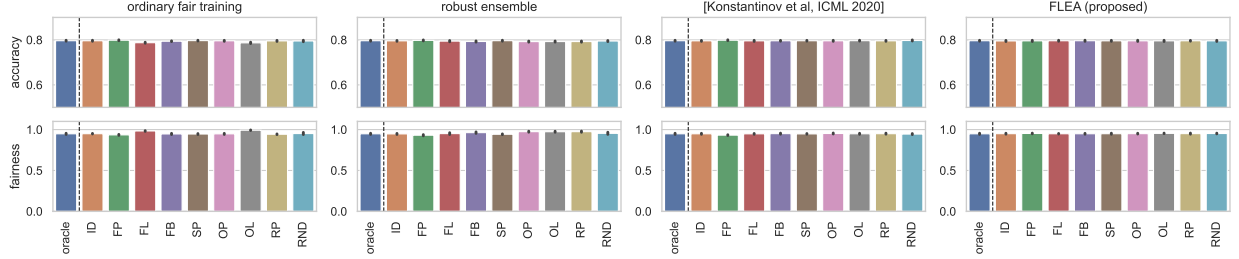

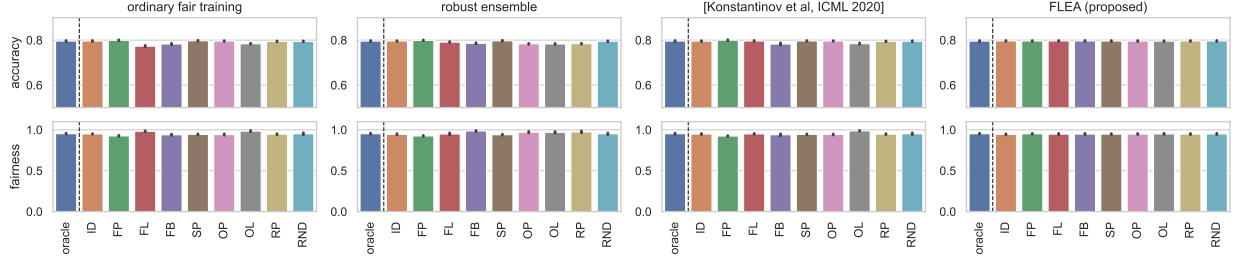

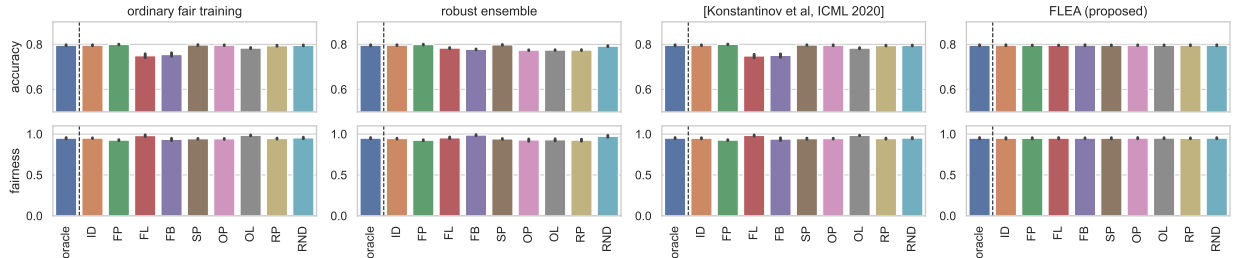
Figure 3: `adult` dataset, preprocessing-based fairness(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 








Figure 4: `adult` dataset, postprocessing-based fairness(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 5: adult dataset, adversarial fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 6: adult dataset, fairness-unaware






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 7: COMPAS dataset, regularization-based fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 8: COMPAS dataset, preprocessing-based fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 9: COMPAS dataset, postprocessing-based fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 10: COMPAS dataset, adversarial fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 11: COMPAS dataset, fairness-unaware






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 



Figure 12: drugs dataset, regularization-based fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 13: drugs dataset, preprocessing-based fairness



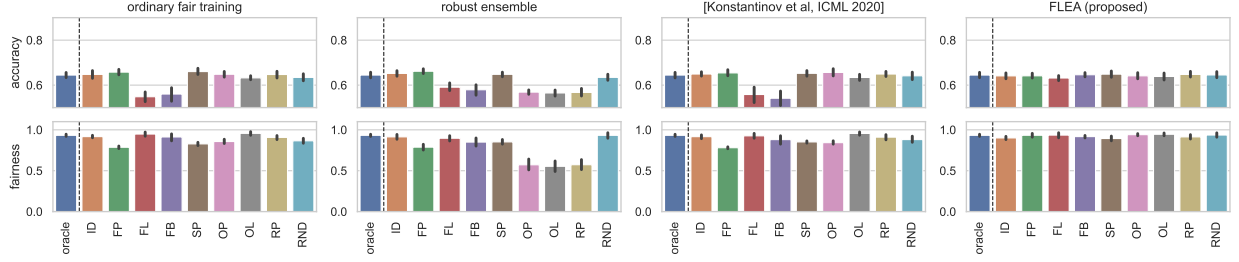

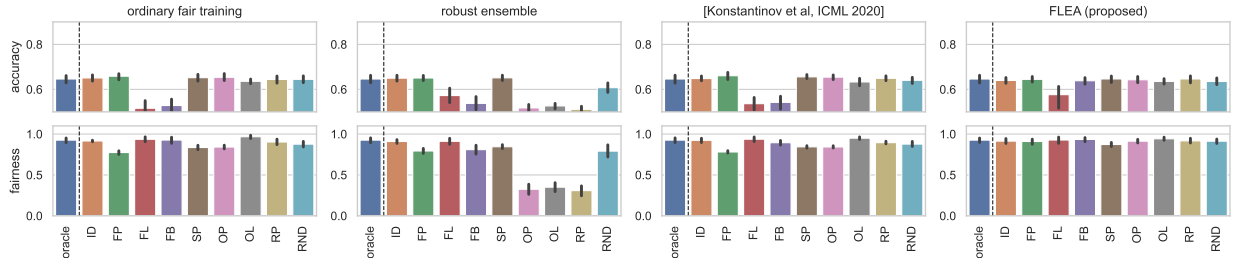
(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 14: drugs dataset, postprocessing-based fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 15: drugs dataset, adversarial fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 16: drugs dataset, fairness-unaware






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 






Figure 17: `germancredit` dataset, regularization-based fairness(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 18: germancredit dataset, preprocessing-based fairness






(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 






Figure 19: `germancredit` dataset, postprocessing-based fairness(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 





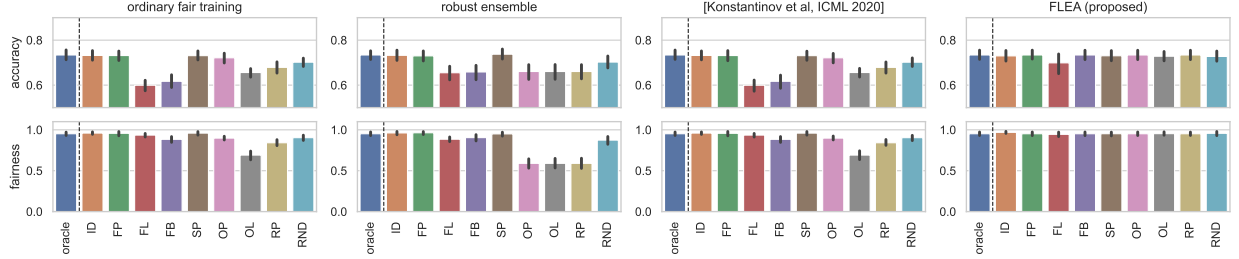
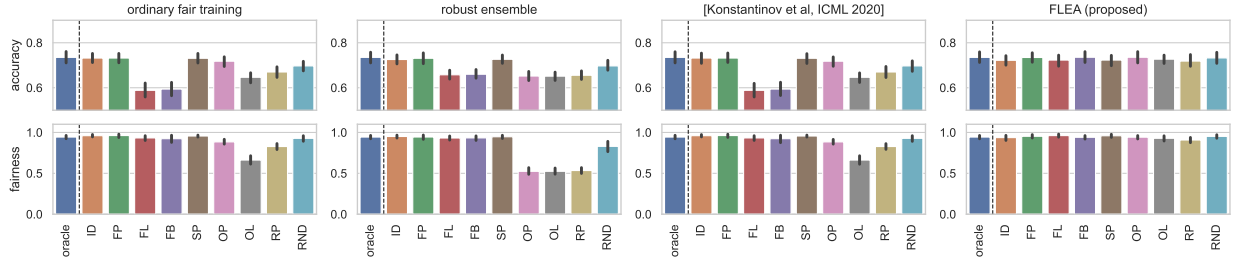
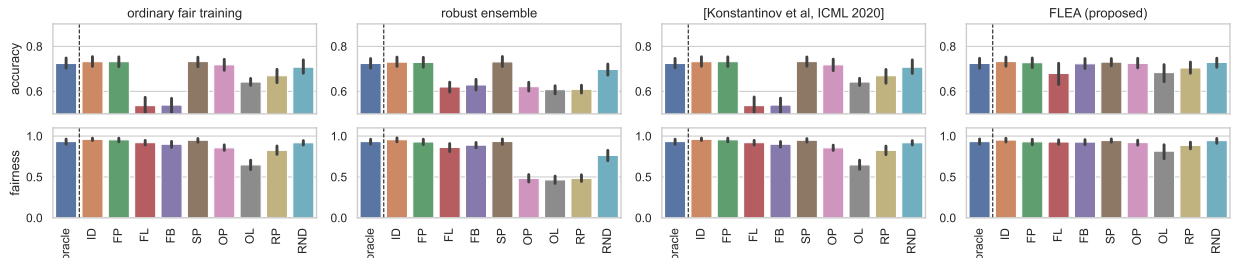
Figure 20: `germancredit` dataset, adversarial fairness(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 


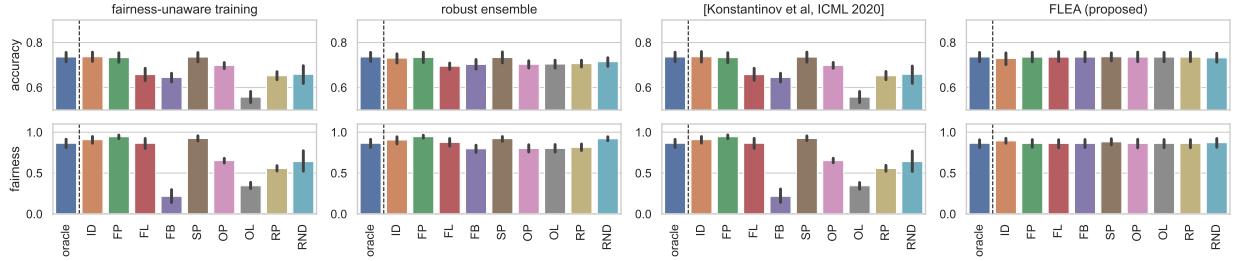
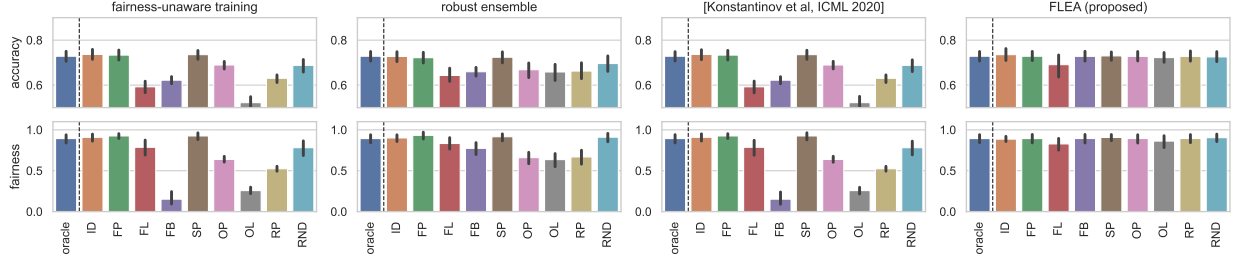
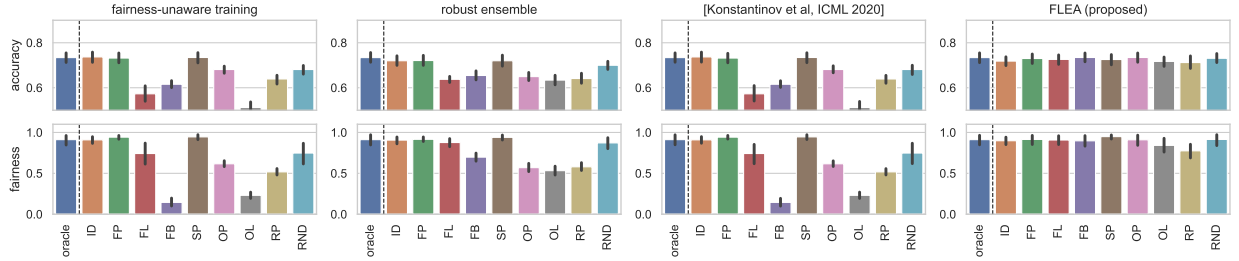
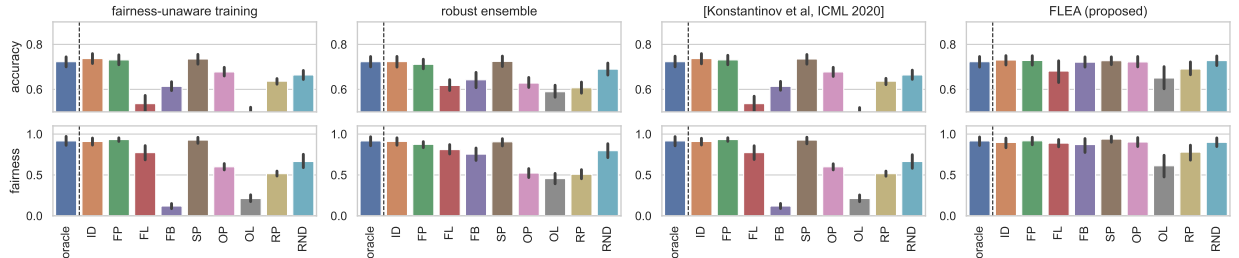
Figure 21: `germancredit` dataset, fairness-unaware(a)  $N = 3, N - K = 1$ (b)  $N = 5, N - K = 2$ (c)  $N = 7, N - K = 3$ (d)  $N = 9, N - K = 4$ (e)  $N = 11, N - K = 5$ 

Figure 22: folktables dataset, regularization-based fairness






(a)  $N = 51, N - K = 5$ (b)  $N = 51, N - K = 10$ (c)  $N = 51, N - K = 15$ (d)  $N = 51, N - K = 20$ (e)  $N = 51, N - K = 25$ 

Figure 23: folktables dataset, preprocessing-based fairness






(a)  $N = 51, N - K = 5$ (b)  $N = 51, N - K = 10$ (c)  $N = 51, N - K = 15$ (d)  $N = 51, N - K = 20$ (e)  $N = 51, N - K = 25$ 

Figure 24: folktables dataset, postprocessing-based fairness






(a)  $N = 51, N - K = 5$ (b)  $N = 51, N - K = 10$ (c)  $N = 51, N - K = 15$ (d)  $N = 51, N - K = 20$ (e)  $N = 51, N - K = 25$ 

Figure 25: folktables dataset, adversarial fairness











(a)  $N = 51, N - K = 5$ (b)  $N = 51, N - K = 10$ (c)  $N = 51, N - K = 15$ (d)  $N = 51, N - K = 20$ (e)  $N = 51, N - K = 25$ 

Figure 26: folktables dataset, fairness-unaware

(a)  $N = 51, N - K = 5$ (b)  $N = 51, N - K = 10$ (c)  $N = 51, N - K = 15$ (d)  $N = 51, N - K = 20$ (e)  $N = 51, N - K = 25$ 

## C Discussion of the role of `disc`, `disp` and `disb` and ablation study

FLEA relies on the combination of three dissimilarity measures: `disc`, `disp` and `disb`. In these section we discuss the importance of each of them and report on an ablation study to verify their practical significance.

Of the three measures, `disc` is indispensable to ensure classifier accuracy, as it is the only measure that depends on the label values. At the same time, `disc` is blind to changes in the protected attributes whenever those are not part of the feature set.

Even in the case when the protected attribute is among the features, the `disc` measure may not detect changes in the data that may harm fairness. For example, if one of the protected groups is much more rare than the other, changing even a small number of data points (e.g. the points from that group) can cause a large change in the conditional distributions of the data given the value of the protected attribute. At the same time, the discrepancy will remain largely unaffected, since only a few points have been changed in total. Therefore, filtering only based on `disc` is insufficient - sources with a different conditional distribution may get in through the filtering step, potentially causing unfair classifiers (on clean data) to appear fair on the (corrupted) dataset.

FLEA avoids such issues by additionally adopting the `disp` measure, which can reliably detect changes in the conditional distributions of the data given the value of the protected attribute, thereby ensuring reliable fairness estimates based on the sources that are returned by the filtering procedure.

Finally, `disb` is required for technical reasons: it ensures that the disparity of the union of some sources is close to their average individual disparity. This is important because a fairness-aware learner works on top of `FILTERSOURCES` by merging the data from the sources returned by the filtering algorithm. This aspect becomes apparent only in the proof of Theorem 1, see Section D.

### C.1 Ablation study

To understand the respective contributions of `disb`, `disc`, and `disp` on real-world data, we performed an ablation study that runs variants of FLEA in which any subsets of the three measures are used to compute the  $D$ -scores. The variant with all measures active is identical to FLEA. The variant with all measures inactive randomly chooses subsets to train on.

The results are presented in Tables 4a–4e. One can see that for all datasets, not using `disc` (column 3) for the filtering step has the most drastic impact. This makes sense, because several of the adversaries make large changes to the labels and features, and `disc` is well suited to identify these.

Not using `disp` (column 2) has usually less of an effect, but in some situations it does lead to a noticeable drop in fairness, see e.g. `COMPAS`, `drugs` with  $N \in \{3, 5, 7\}$ , and `folktables` with  $N - K \in \{5, 10, 15, 20\}$ . This is also consistent with our expectations, as the measure is specifically able to detect even subtle manipulation that would negatively affect fairness. However, in some cases when the amount of data per source is small (e.g. `germancredit`, and `drugs` with  $N \in \{9, 11\}$ ), or the amount of manipulated data is very close to half (`folktables` with  $N - K = 25$ ), not using `disp` is actually beneficial for the system. We attribute this to the fact that as a difference of ratios, `disp` is harder to estimate from small sample sets than the other two measures. A noisy estimate, however, can lead to clean sources to be suppressed, and manipulated ones to be selected. This explanation is also supported by the fact that for small datasets the variability of results is bigger when `disp` is included than when it is not.

The effect of not using `disb` (column 4) is small on real data. It rarely exceeds the standard deviation of the estimates. This, again, is expected, as the other two measures are typically able to ensure accuracy and fairness, whereas the role of `disb` is mainly to handle corner cases that are unlikely to occur in real data.

Dropping two measures from the filtering step only makes sense, if the remaining measure is `disc` (column 6). Even then, a decrease in accuracy and/or fairness is quite common, or at least an increase in variability.



Table 4: Performance of FLEA with different combination of disb, disc, and disp activated or deactivated (crossed-out). Reported results are in the same format at Table 1: minimal accuracy (A) and fairness (F) against any of the tested adversaries

(a) regularization-based fairness									
		disb disc disp	disb disc <del>disp</del>	disb <del>disc</del> disp	<del>disb</del> disc disp	disb <del>disc</del> <del>disp</del>	<del>disb</del> disc <del>disp</del>	<del>disb</del> <del>disc</del> disp	<del>disb</del> <del>disc</del> <del>disp</del>
adult $N = 3, N - K = 1$	A:	70.4 $\pm$ 0.3	70.4 $\pm$ 0.3	56.0 $\pm$ 10.3	70.4 $\pm$ 0.3	53.2 $\pm$ 9.7	70.4 $\pm$ 0.3	58.3 $\pm$ 10.8	63.3 $\pm$ 11.1
	F:	98.4 $\pm$ 0.9	98.4 $\pm$ 0.9	97.3 $\pm$ 1.2	98.4 $\pm$ 0.9	84.8 $\pm$ 7.0	90.0 $\pm$ 6.1	97.5 $\pm$ 1.4	89.2 $\pm$ 13.5
adult $N = 5, N - K = 2$	A:	70.3 $\pm$ 0.4	70.3 $\pm$ 0.4	60.4 $\pm$ 13.8	70.3 $\pm$ 0.4	53.3 $\pm$ 16.6	70.3 $\pm$ 0.4	60.1 $\pm$ 13.4	57.5 $\pm$ 16.5
	F:	98.2 $\pm$ 1.1	98.2 $\pm$ 1.1	86.0 $\pm$ 23.4	98.2 $\pm$ 1.1	64.0 $\pm$ 22.4	89.2 $\pm$ 5.5	84.7 $\pm$ 24.9	72.1 $\pm$ 23.5
adult $N = 7, N - K = 3$	A:	70.3 $\pm$ 0.3	70.3 $\pm$ 0.3	47.8 $\pm$ 10.3	70.3 $\pm$ 0.3	58.9 $\pm$ 11.8	70.3 $\pm$ 0.3	47.0 $\pm$ 10.4	56.1 $\pm$ 9.5
	F:	98.1 $\pm$ 0.9	98.0 $\pm$ 0.9	79.6 $\pm$ 35.0	98.1 $\pm$ 0.9	77.5 $\pm$ 23.0	85.7 $\pm$ 4.5	89.3 $\pm$ 25.6	76.5 $\pm$ 8.5
adult $N = 9, N - K = 4$	A:	70.2 $\pm$ 0.3	70.2 $\pm$ 0.3	56.5 $\pm$ 14.6	70.2 $\pm$ 0.3	56.7 $\pm$ 15.8	70.2 $\pm$ 0.3	53.3 $\pm$ 16.2	53.3 $\pm$ 16.8
	F:	97.6 $\pm$ 1.2	97.6 $\pm$ 1.3	87.9 $\pm$ 14.1	97.6 $\pm$ 1.2	78.3 $\pm$ 8.5	85.0 $\pm$ 4.7	78.5 $\pm$ 29.0	70.3 $\pm$ 24.6
adult $N = 11, N - K = 5$	A:	70.2 $\pm$ 0.3	70.2 $\pm$ 0.3	51.8 $\pm$ 12.3	70.2 $\pm$ 0.3	60.0 $\pm$ 11.5	70.2 $\pm$ 0.3	54.9 $\pm$ 13.3	52.0 $\pm$ 15.1
	F:	98.0 $\pm$ 1.0	97.9 $\pm$ 1.1	84.4 $\pm$ 24.9	98.0 $\pm$ 1.0	79.9 $\pm$ 14.3	87.6 $\pm$ 4.9	84.7 $\pm$ 25.0	66.0 $\pm$ 22.3
compas $N = 3, N - K = 1$	A:	66.0 $\pm$ 1.1	65.9 $\pm$ 1.1	59.9 $\pm$ 8.5	66.0 $\pm$ 1.1	54.8 $\pm$ 8.4	65.9 $\pm$ 1.1	58.2 $\pm$ 9.1	61.0 $\pm$ 8.3
	F:	96.2 $\pm$ 1.9	94.7 $\pm$ 3.1	94.9 $\pm$ 2.1	96.2 $\pm$ 1.9	86.8 $\pm$ 7.4	91.1 $\pm$ 4.0	94.9 $\pm$ 1.8	91.1 $\pm$ 7.6
compas $N = 5, N - K = 2$	A:	66.1 $\pm$ 1.2	66.0 $\pm$ 1.2	55.4 $\pm$ 13.6	66.1 $\pm$ 1.2	43.8 $\pm$ 12.9	66.0 $\pm$ 1.1	55.5 $\pm$ 13.6	56.4 $\pm$ 13.2
	F:	96.2 $\pm$ 1.3	93.1 $\pm$ 1.8	94.0 $\pm$ 2.5	96.2 $\pm$ 1.3	86.4 $\pm$ 7.6	90.2 $\pm$ 1.9	93.8 $\pm$ 2.8	82.7 $\pm$ 7.7
compas $N = 7, N - K = 3$	A:	65.8 $\pm$ 1.3	65.8 $\pm$ 1.2	52.4 $\pm$ 8.8	65.7 $\pm$ 1.1	59.2 $\pm$ 8.8	65.6 $\pm$ 1.0	52.0 $\pm$ 7.3	56.0 $\pm$ 8.7
	F:	94.3 $\pm$ 2.8	94.1 $\pm$ 1.8	89.1 $\pm$ 14.9	92.8 $\pm$ 3.6	90.1 $\pm$ 4.2	87.8 $\pm$ 2.7	87.7 $\pm$ 14.4	83.4 $\pm$ 4.5
compas $N = 9, N - K = 4$	A:	65.9 $\pm$ 1.1	66.0 $\pm$ 1.1	60.1 $\pm$ 8.8	65.9 $\pm$ 1.2	56.2 $\pm$ 11.1	65.7 $\pm$ 0.9	58.7 $\pm$ 10.6	52.8 $\pm$ 13.8
	F:	94.7 $\pm$ 2.8	94.4 $\pm$ 2.6	93.0 $\pm$ 2.2	93.5 $\pm$ 3.7	87.3 $\pm$ 4.6	88.4 $\pm$ 3.6	92.7 $\pm$ 3.3	75.6 $\pm$ 20.9
compas $N = 11, N - K = 5$	A:	65.9 $\pm$ 1.1	65.9 $\pm$ 1.3	51.4 $\pm$ 9.0	65.6 $\pm$ 1.0	52.2 $\pm$ 11.5	65.8 $\pm$ 1.1	53.6 $\pm$ 8.7	51.2 $\pm$ 13.2
	F:	94.4 $\pm$ 2.6	93.1 $\pm$ 2.6	90.3 $\pm$ 3.1	91.5 $\pm$ 3.2	78.1 $\pm$ 24.0	88.8 $\pm$ 2.0	89.6 $\pm$ 2.5	76.4 $\pm$ 9.0
drugs $N = 3, N - K = 1$	A:	64.7 $\pm$ 1.3	64.3 $\pm$ 1.5	55.1 $\pm$ 6.8	64.7 $\pm$ 1.3	51.6 $\pm$ 4.9	64.4 $\pm$ 1.5	57.8 $\pm$ 7.4	60.8 $\pm$ 6.4
	F:	94.2 $\pm$ 4.2	88.0 $\pm$ 6.1	93.0 $\pm$ 4.3	94.2 $\pm$ 4.2	74.9 $\pm$ 8.0	89.7 $\pm$ 4.7	93.6 $\pm$ 3.7	86.5 $\pm$ 12.4
drugs $N = 5, N - K = 2$	A:	64.2 $\pm$ 1.4	64.3 $\pm$ 1.5	59.3 $\pm$ 8.5	64.2 $\pm$ 1.4	56.4 $\pm$ 9.9	64.3 $\pm$ 1.5	57.6 $\pm$ 9.1	53.5 $\pm$ 11.6
	F:	92.8 $\pm$ 4.1	88.6 $\pm$ 5.7	78.8 $\pm$ 19.9	92.8 $\pm$ 4.1	72.9 $\pm$ 12.4	86.3 $\pm$ 6.5	80.7 $\pm$ 19.4	72.1 $\pm$ 18.7
drugs $N = 7, N - K = 3$	A:	63.3 $\pm$ 2.2	63.8 $\pm$ 2.0	54.9 $\pm$ 9.7	63.7 $\pm$ 2.0	50.2 $\pm$ 9.4	63.9 $\pm$ 2.0	52.6 $\pm$ 10.3	53.5 $\pm$ 8.6
	F:	91.0 $\pm$ 5.6	82.9 $\pm$ 7.5	78.2 $\pm$ 25.5	88.8 $\pm$ 6.4	64.1 $\pm$ 18.5	84.4 $\pm$ 5.9	69.6 $\pm$ 32.8	72.8 $\pm$ 7.0
drugs $N = 9, N - K = 4$	A:	62.4 $\pm$ 1.6	64.1 $\pm$ 2.1	53.1 $\pm$ 8.8	62.0 $\pm$ 1.6	50.5 $\pm$ 9.9	63.8 $\pm$ 2.3	49.4 $\pm$ 8.4	51.3 $\pm$ 10.3
	F:	86.1 $\pm$ 3.9	77.2 $\pm$ 8.6	82.8 $\pm$ 5.0	85.8 $\pm$ 5.8	63.5 $\pm$ 22.2	81.7 $\pm$ 6.2	82.1 $\pm$ 5.3	69.7 $\pm$ 24.1
drugs $N = 11, N - K = 5$	A:	60.3 $\pm$ 7.2	62.3 $\pm$ 4.2	54.3 $\pm$ 7.4	57.3 $\pm$ 8.8	47.5 $\pm$ 11.3	64.1 $\pm$ 2.3	51.7 $\pm$ 7.4	51.5 $\pm$ 10.5
	F:	85.4 $\pm$ 5.8	76.8 $\pm$ 7.7	77.2 $\pm$ 10.5	85.9 $\pm$ 4.7	61.9 $\pm$ 17.3	79.0 $\pm$ 7.1	75.4 $\pm$ 12.0	63.4 $\pm$ 23.9
germancredit $N = 3, N - K = 1$	A:	66.1 $\pm$ 3.1	66.2 $\pm$ 2.5	55.7 $\pm$ 8.8	66.1 $\pm$ 3.1	56.5 $\pm$ 7.1	66.2 $\pm$ 2.5	55.0 $\pm$ 8.5	60.4 $\pm$ 9.8
	F:	92.4 $\pm$ 3.2	92.9 $\pm$ 5.3	88.0 $\pm$ 4.8	92.4 $\pm$ 3.2	87.2 $\pm$ 6.8	92.3 $\pm$ 5.2	88.4 $\pm$ 3.8	87.9 $\pm$ 9.7
germancredit $N = 5, N - K = 2$	A:	66.3 $\pm$ 3.2	65.8 $\pm$ 3.0	58.6 $\pm$ 10.7	66.3 $\pm$ 3.1	60.3 $\pm$ 7.5	66.3 $\pm$ 2.6	52.2 $\pm$ 12.0	56.6 $\pm$ 11.5
	F:	92.8 $\pm$ 3.9	92.4 $\pm$ 3.7	89.0 $\pm$ 5.9	92.2 $\pm$ 4.2	88.3 $\pm$ 4.4	91.5 $\pm$ 4.2	85.6 $\pm$ 6.2	79.6 $\pm$ 16.0
germancredit $N = 7, N - K = 3$	A:	61.8 $\pm$ 6.6	65.4 $\pm$ 3.2	52.8 $\pm$ 11.2	64.6 $\pm$ 3.2	49.5 $\pm$ 8.5	64.9 $\pm$ 2.2	52.5 $\pm$ 9.0	54.4 $\pm$ 7.9
	F:	91.2 $\pm$ 5.3	91.5 $\pm$ 5.3	79.4 $\pm$ 13.9	89.9 $\pm$ 5.4	79.6 $\pm$ 16.6	89.3 $\pm$ 6.0	83.6 $\pm$ 6.7	82.5 $\pm$ 5.4
germancredit $N = 9, N - K = 4$	A:	64.2 $\pm$ 3.1	66.2 $\pm$ 3.0	57.0 $\pm$ 7.9	62.6 $\pm$ 3.2	48.8 $\pm$ 9.8	64.7 $\pm$ 2.2	50.6 $\pm$ 8.1	53.0 $\pm$ 10.2
	F:	90.2 $\pm$ 4.0	91.6 $\pm$ 3.6	82.1 $\pm$ 10.2	88.5 $\pm$ 4.7	79.4 $\pm$ 17.6	91.0 $\pm$ 3.6	72.9 $\pm$ 22.8	75.1 $\pm$ 25.5
germancredit $N = 11, N - K = 5$	A:	60.9 $\pm$ 4.9	65.2 $\pm$ 3.1	57.3 $\pm$ 8.2	60.6 $\pm$ 3.7	53.9 $\pm$ 9.8	64.7 $\pm$ 2.6	54.8 $\pm$ 5.9	51.4 $\pm$ 10.2
	F:	88.8 $\pm$ 3.4	92.0 $\pm$ 3.8	83.4 $\pm$ 5.9	88.2 $\pm$ 4.5	81.2 $\pm$ 7.7	91.7 $\pm$ 3.1	77.3 $\pm$ 11.6	80.0 $\pm$ 13.8
folktables $N = 51, N - K = 5$	A:	75.5 $\pm$ 0.2	75.5 $\pm$ 0.2	74.7 $\pm$ 0.2	75.4 $\pm$ 0.2	74.4 $\pm$ 0.5	75.4 $\pm$ 0.2	74.7 $\pm$ 0.4	74.3 $\pm$ 0.5
	F:	99.5 $\pm$ 0.2	99.0 $\pm$ 0.5	99.5 $\pm$ 0.2	99.6 $\pm$ 0.3	93.0 $\pm$ 2.6	99.1 $\pm$ 0.6	99.4 $\pm$ 0.2	92.8 $\pm$ 4.2
folktables $N = 51, N - K = 10$	A:	75.4 $\pm$ 0.2	75.4 $\pm$ 0.2	73.9 $\pm$ 0.3	75.4 $\pm$ 0.2	73.7 $\pm$ 0.3	75.3 $\pm$ 0.2	73.8 $\pm$ 0.4	73.1 $\pm$ 1.0
	F:	99.4 $\pm$ 0.3	98.1 $\pm$ 0.7	99.2 $\pm$ 0.1	99.4 $\pm$ 0.3	88.3 $\pm$ 3.0	98.2 $\pm$ 0.5	99.1 $\pm$ 0.3	86.8 $\pm$ 3.8
folktables $N = 51, N - K = 15$	A:	75.3 $\pm$ 0.3	75.4 $\pm$ 0.3	72.9 $\pm$ 1.1	75.3 $\pm$ 0.2	72.3 $\pm$ 0.9	75.3 $\pm$ 0.3	72.8 $\pm$ 0.8	68.9 $\pm$ 8.2
	F:	99.3 $\pm$ 0.4	97.0 $\pm$ 0.6	99.2 $\pm$ 0.2	99.4 $\pm$ 0.3	82.4 $\pm$ 2.4	97.1 $\pm$ 1.0	99.1 $\pm$ 0.3	78.9 $\pm$ 5.4
folktables $N = 51, N - K = 20$	A:	75.3 $\pm$ 0.2	75.3 $\pm$ 0.2	70.3 $\pm$ 3.8	75.3 $\pm$ 0.2	70.5 $\pm$ 2.7	75.3 $\pm$ 0.2	66.9 $\pm$ 9.1	64.8 $\pm$ 10.9
	F:	98.3 $\pm$ 1.1	95.6 $\pm$ 1.0	98.8 $\pm$ 0.4	98.4 $\pm$ 1.0	79.1 $\pm$ 3.1	93.9 $\pm$ 1.5	98.3 $\pm$ 0.9	75.4 $\pm$ 3.5
folktables $N = 51, N - K = 25$	A:	66.4 $\pm$ 11.5	75.0 $\pm$ 0.3	59.1 $\pm$ 14.6	65.4 $\pm$ 13.9	41.7 $\pm$ 14.2	75.1 $\pm$ 0.3	52.4 $\pm$ 14.4	56.6 $\pm$ 13.3
	F:	86.7 $\pm$ 14.0	89.9 $\pm$ 1.6	88.6 $\pm$ 5.5	89.0 $\pm$ 7.9	79.4 $\pm$ 6.6	89.3 $\pm$ 1.8	87.9 $\pm$ 8.1	74.6 $\pm$ 3.6

(b) postprocessing-based fairness

			disb disc disp	disb disc <del>disp</del>	disb <del>disc</del> disp	<del>disb</del> disc disp	disb <del>disc</del> <del>disp</del>	<del>disb</del> disc <del>disp</del>	<del>disb</del> <del>disc</del> disp	<del>disb</del> <del>disc</del> <del>disp</del>
$N = 3, N - K = 1$	adult	A:	78.9 $\pm$ 0.3	78.9 $\pm$ 0.3	53.8 $\pm$ 17.3	78.9 $\pm$ 0.3	53.7 $\pm$ 13.1	78.9 $\pm$ 0.3	57.2 $\pm$ 18.4	69.6 $\pm$ 14.3
		F:	99.3 $\pm$ 0.5	99.3 $\pm$ 0.5	98.9 $\pm$ 0.6	99.3 $\pm$ 0.5	89.5 $\pm$ 5.2	94.5 $\pm$ 3.4	99.0 $\pm$ 0.6	95.4 $\pm$ 5.8
$N = 5, N - K = 2$	adult	A:	78.9 $\pm$ 0.4	78.9 $\pm$ 0.4	66.9 $\pm$ 22.9	78.9 $\pm$ 0.4	55.6 $\pm$ 27.9	78.9 $\pm$ 0.3	67.0 $\pm$ 22.8	60.7 $\pm$ 26.7
		F:	99.2 $\pm$ 0.6	99.2 $\pm$ 0.6	98.9 $\pm$ 0.6	99.2 $\pm$ 0.6	88.6 $\pm$ 1.4	93.0 $\pm$ 3.1	98.9 $\pm$ 0.5	90.6 $\pm$ 4.6
$N = 7, N - K = 3$	adult	A:	78.8 $\pm$ 0.3	78.8 $\pm$ 0.3	50.1 $\pm$ 17.7	78.8 $\pm$ 0.3	60.6 $\pm$ 22.6	78.8 $\pm$ 0.3	52.2 $\pm$ 17.4	61.5 $\pm$ 14.2
		F:	99.2 $\pm$ 0.6	98.9 $\pm$ 1.0	98.3 $\pm$ 1.0	99.2 $\pm$ 0.6	89.8 $\pm$ 3.3	92.0 $\pm$ 2.0	98.6 $\pm$ 0.7	87.8 $\pm$ 2.1
$N = 9, N - K = 4$	adult	A:	78.9 $\pm$ 0.3	78.9 $\pm$ 0.3	61.1 $\pm$ 25.7	78.9 $\pm$ 0.3	61.2 $\pm$ 25.5	78.9 $\pm$ 0.3	55.7 $\pm$ 27.5	55.1 $\pm$ 27.9
		F:	99.1 $\pm$ 0.5	98.9 $\pm$ 0.8	98.6 $\pm$ 0.5	99.1 $\pm$ 0.5	88.7 $\pm$ 2.5	91.9 $\pm$ 2.5	98.8 $\pm$ 0.4	88.2 $\pm$ 3.3
$N = 11, N - K = 5$	adult	A:	78.9 $\pm$ 0.3	78.9 $\pm$ 0.3	51.5 $\pm$ 20.4	78.9 $\pm$ 0.3	67.6 $\pm$ 18.4	78.9 $\pm$ 0.3	56.8 $\pm$ 23.6	50.9 $\pm$ 25.0
		F:	99.0 $\pm$ 0.7	98.7 $\pm$ 1.1	98.3 $\pm$ 1.1	99.0 $\pm$ 0.8	89.2 $\pm$ 2.9	92.2 $\pm$ 2.8	98.7 $\pm$ 0.6	88.3 $\pm$ 3.1
$N = 3, N - K = 1$	compas	A:	65.2 $\pm$ 1.1	65.2 $\pm$ 1.1	58.4 $\pm$ 8.8	65.2 $\pm$ 1.1	53.6 $\pm$ 8.5	65.2 $\pm$ 1.1	57.5 $\pm$ 8.7	60.4 $\pm$ 8.4
		F:	97.1 $\pm$ 2.1	94.9 $\pm$ 4.5	96.1 $\pm$ 2.6	97.1 $\pm$ 2.1	90.2 $\pm$ 4.9	91.1 $\pm$ 6.1	96.6 $\pm$ 1.9	91.3 $\pm$ 9.2
$N = 5, N - K = 2$	compas	A:	65.5 $\pm$ 1.1	65.2 $\pm$ 1.2	54.7 $\pm$ 12.2	65.5 $\pm$ 1.1	43.7 $\pm$ 12.2	65.2 $\pm$ 1.2	54.3 $\pm$ 12.0	54.5 $\pm$ 13.1
		F:	96.0 $\pm$ 2.1	93.6 $\pm$ 2.8	95.2 $\pm$ 1.7	96.0 $\pm$ 2.1	89.0 $\pm$ 3.0	90.1 $\pm$ 2.7	94.8 $\pm$ 1.9	82.9 $\pm$ 9.9
$N = 7, N - K = 3$	compas	A:	65.5 $\pm$ 1.1	65.4 $\pm$ 1.0	52.1 $\pm$ 9.1	65.2 $\pm$ 0.9	58.2 $\pm$ 8.6	65.0 $\pm$ 0.8	52.1 $\pm$ 7.6	53.7 $\pm$ 8.8
		F:	93.5 $\pm$ 1.6	93.6 $\pm$ 1.6	94.0 $\pm$ 1.5	92.3 $\pm$ 4.4	91.0 $\pm$ 4.7	88.0 $\pm$ 3.9	94.3 $\pm$ 1.7	80.7 $\pm$ 4.8
$N = 9, N - K = 4$	compas	A:	65.5 $\pm$ 1.0	65.6 $\pm$ 0.9	58.1 $\pm$ 8.1	65.3 $\pm$ 1.2	55.0 $\pm$ 10.5	65.3 $\pm$ 1.0	57.2 $\pm$ 10.5	51.6 $\pm$ 12.9
		F:	95.2 $\pm$ 1.9	94.8 $\pm$ 3.6	93.8 $\pm$ 2.0	93.2 $\pm$ 3.7	88.8 $\pm$ 2.9	88.2 $\pm$ 3.9	93.2 $\pm$ 3.2	78.5 $\pm$ 7.4
$N = 11, N - K = 5$	compas	A:	64.9 $\pm$ 1.0	65.2 $\pm$ 1.1	50.7 $\pm$ 8.6	64.9 $\pm$ 0.8	51.4 $\pm$ 9.9	65.1 $\pm$ 1.1	53.2 $\pm$ 7.8	50.6 $\pm$ 12.8
		F:	96.0 $\pm$ 2.1	94.0 $\pm$ 2.4	93.6 $\pm$ 3.8	93.5 $\pm$ 3.0	89.0 $\pm$ 4.5	89.0 $\pm$ 2.9	90.5 $\pm$ 3.2	78.2 $\pm$ 7.2
$N = 3, N - K = 1$	drugs	A:	64.3 $\pm$ 1.4	64.1 $\pm$ 1.5	55.7 $\pm$ 6.5	64.3 $\pm$ 1.4	50.8 $\pm$ 4.7	64.2 $\pm$ 1.5	58.1 $\pm$ 6.7	60.5 $\pm$ 5.8
		F:	96.1 $\pm$ 2.8	91.6 $\pm$ 4.7	94.0 $\pm$ 3.4	96.1 $\pm$ 2.8	75.1 $\pm$ 8.5	93.1 $\pm$ 5.3	94.9 $\pm$ 2.7	88.3 $\pm$ 12.7
$N = 5, N - K = 2$	drugs	A:	63.4 $\pm$ 2.0	63.3 $\pm$ 2.0	59.4 $\pm$ 7.1	63.4 $\pm$ 2.0	56.5 $\pm$ 9.1	63.3 $\pm$ 2.0	56.8 $\pm$ 8.8	53.6 $\pm$ 11.0
		F:	95.3 $\pm$ 2.6	91.1 $\pm$ 6.5	93.1 $\pm$ 3.1	95.3 $\pm$ 2.6	76.4 $\pm$ 9.4	89.2 $\pm$ 7.3	92.0 $\pm$ 2.4	78.8 $\pm$ 13.9
$N = 7, N - K = 3$	drugs	A:	63.0 $\pm$ 2.5	63.5 $\pm$ 1.9	54.4 $\pm$ 8.9	63.4 $\pm$ 1.9	50.3 $\pm$ 9.9	63.3 $\pm$ 1.9	52.7 $\pm$ 9.4	53.7 $\pm$ 8.1
		F:	95.0 $\pm$ 3.1	84.5 $\pm$ 8.0	88.0 $\pm$ 8.6	92.2 $\pm$ 5.7	69.5 $\pm$ 4.8	86.8 $\pm$ 5.7	88.1 $\pm$ 6.6	74.0 $\pm$ 7.8
$N = 9, N - K = 4$	drugs	A:	61.1 $\pm$ 2.3	63.0 $\pm$ 2.1	53.0 $\pm$ 8.2	61.2 $\pm$ 2.2	51.2 $\pm$ 10.4	63.3 $\pm$ 2.2	50.4 $\pm$ 8.2	50.6 $\pm$ 10.1
		F:	90.4 $\pm$ 4.3	80.5 $\pm$ 8.7	82.2 $\pm$ 6.7	88.6 $\pm$ 4.6	70.2 $\pm$ 12.1	84.4 $\pm$ 6.3	83.5 $\pm$ 5.4	72.8 $\pm$ 12.3
$N = 11, N - K = 5$	drugs	A:	58.4 $\pm$ 6.6	62.6 $\pm$ 3.2	54.1 $\pm$ 6.4	55.7 $\pm$ 7.2	47.3 $\pm$ 10.0	63.0 $\pm$ 1.8	51.9 $\pm$ 6.8	51.9 $\pm$ 10.8
		F:	89.0 $\pm$ 5.3	80.2 $\pm$ 7.2	81.8 $\pm$ 3.8	87.3 $\pm$ 3.9	71.8 $\pm$ 7.9	81.8 $\pm$ 7.8	82.1 $\pm$ 6.6	72.4 $\pm$ 13.3
$N = 3, N - K = 1$	germancredit	A:	72.0 $\pm$ 3.6	71.7 $\pm$ 2.7	53.7 $\pm$ 14.4	72.0 $\pm$ 3.6	53.2 $\pm$ 14.1	71.7 $\pm$ 2.7	55.5 $\pm$ 15.5	63.4 $\pm$ 13.1
		F:	92.4 $\pm$ 3.1	91.8 $\pm$ 5.1	87.1 $\pm$ 7.4	92.4 $\pm$ 3.1	84.8 $\pm$ 7.3	91.5 $\pm$ 5.1	87.6 $\pm$ 7.1	87.3 $\pm$ 8.7
$N = 5, N - K = 2$	germancredit	A:	71.9 $\pm$ 3.9	71.9 $\pm$ 2.6	59.4 $\pm$ 13.2	72.0 $\pm$ 3.4	64.7 $\pm$ 10.9	72.3 $\pm$ 3.1	51.9 $\pm$ 16.1	60.6 $\pm$ 13.2
		F:	93.5 $\pm$ 3.6	92.1 $\pm$ 5.1	87.9 $\pm$ 7.1	92.1 $\pm$ 3.7	86.6 $\pm$ 8.5	91.7 $\pm$ 3.8	85.6 $\pm$ 8.0	82.8 $\pm$ 9.1
$N = 7, N - K = 3$	germancredit	A:	68.8 $\pm$ 6.7	71.5 $\pm$ 2.9	55.1 $\pm$ 15.2	70.7 $\pm$ 3.6	54.4 $\pm$ 10.8	70.9 $\pm$ 3.4	53.5 $\pm$ 14.1	57.0 $\pm$ 12.3
		F:	91.4 $\pm$ 4.5	93.0 $\pm$ 4.9	80.8 $\pm$ 8.2	91.3 $\pm$ 4.6	80.6 $\pm$ 9.1	91.6 $\pm$ 4.9	82.9 $\pm$ 8.5	79.8 $\pm$ 8.6
$N = 9, N - K = 4$	germancredit	A:	69.5 $\pm$ 4.7	71.9 $\pm$ 3.9	57.5 $\pm$ 11.6	67.7 $\pm$ 3.7	48.4 $\pm$ 14.9	70.6 $\pm$ 2.3	47.9 $\pm$ 14.4	53.3 $\pm$ 16.0
		F:	91.7 $\pm$ 4.4	92.9 $\pm$ 2.4	78.8 $\pm$ 8.2	90.1 $\pm$ 3.2	80.4 $\pm$ 7.6	91.0 $\pm$ 4.7	80.1 $\pm$ 7.7	79.6 $\pm$ 10.4
$N = 11, N - K = 5$	germancredit	A:	65.0 $\pm$ 9.6	71.3 $\pm$ 3.8	61.4 $\pm$ 7.6	66.5 $\pm$ 6.6	55.5 $\pm$ 13.6	70.8 $\pm$ 3.4	55.2 $\pm$ 10.9	53.0 $\pm$ 15.3
		F:	87.4 $\pm$ 4.3	91.9 $\pm$ 3.1	82.0 $\pm$ 5.0	86.1 $\pm$ 5.0	81.6 $\pm$ 6.8	89.8 $\pm$ 4.0	79.3 $\pm$ 5.6	80.6 $\pm$ 7.3
$N = 51, N - K = 5$	folktables	A:	77.3 $\pm$ 0.2	77.3 $\pm$ 0.2	77.2 $\pm$ 0.2	77.3 $\pm$ 0.2	77.1 $\pm$ 0.2	77.3 $\pm$ 0.2	77.2 $\pm$ 0.2	77.1 $\pm$ 0.2
		F:	99.6 $\pm$ 0.3	99.3 $\pm$ 0.3	99.4 $\pm$ 0.2	99.6 $\pm$ 0.3	93.7 $\pm$ 2.6	99.3 $\pm$ 0.4	99.5 $\pm$ 0.2	94.0 $\pm$ 3.7
$N = 51, N - K = 10$	folktables	A:	77.3 $\pm$ 0.2	77.4 $\pm$ 0.1	77.0 $\pm$ 0.3	77.3 $\pm$ 0.2	76.9 $\pm$ 0.2	77.3 $\pm$ 0.2	77.0 $\pm$ 0.3	76.6 $\pm$ 0.7
		F:	99.5 $\pm$ 0.3	98.6 $\pm$ 0.6	99.5 $\pm$ 0.2	99.6 $\pm$ 0.2	89.1 $\pm$ 2.9	98.8 $\pm$ 0.4	99.4 $\pm$ 0.3	88.5 $\pm$ 3.2
$N = 51, N - K = 15$	folktables	A:	77.3 $\pm$ 0.2	77.4 $\pm$ 0.2	76.0 $\pm$ 1.1	77.3 $\pm$ 0.2	76.1 $\pm$ 0.6	77.3 $\pm$ 0.2	76.2 $\pm$ 0.8	72.2 $\pm$ 7.4
		F:	99.4 $\pm$ 0.4	97.9 $\pm$ 0.7	99.3 $\pm$ 0.4	99.5 $\pm$ 0.4	83.3 $\pm$ 2.8	97.9 $\pm$ 1.1	99.4 $\pm$ 0.4	81.0 $\pm$ 5.3
$N = 51, N - K = 20$	folktables	A:	77.3 $\pm$ 0.2	77.3 $\pm$ 0.2	73.4 $\pm$ 4.3	77.3 $\pm$ 0.2	73.2 $\pm$ 4.6	77.3 $\pm$ 0.2	70.8 $\pm$ 8.3	67.6 $\pm$ 13.5
		F:	98.7 $\pm$ 0.7	96.6 $\pm$ 0.7	99.2 $\pm$ 0.4	98.9 $\pm$ 0.6	79.3 $\pm$ 4.0	95.1 $\pm$ 1.4	99.2 $\pm$ 0.3	76.5 $\pm$ 4.1
$N = 51, N - K = 25$	folktables	A:	69.4 $\pm$ 12.8	77.2 $\pm$ 0.1	59.9 $\pm$ 16.3	67.8 $\pm$ 15.5	40.4 $\pm$ 16.8	77.2 $\pm$ 0.1	52.9 $\pm$ 15.9	59.9 $\pm$ 15.2
		F:	94.0 $\pm$ 2.2	91.1 $\pm$ 1.4	84.4 $\pm$ 11.9	93.9 $\pm$ 2.5	71.6 $\pm$ 3.7	90.6 $\pm$ 1.7	85.0 $\pm$ 12.1	73.4 $\pm$ 3.2

(c) adversarial fairness

			disb disc disp	disb disc <del>disc</del> disp	disb <del>disc</del> disp	<del>disb</del> disc disp	disb <del>disc</del> disp	<del>disb</del> disc <del>disc</del> disp	<del>disb</del> disc <del>disc</del> disp	<del>disb</del> disc <del>disc</del> disp
$N = 3, N - K = 1$	adult	A:	80.7 $\pm$ 0.3	80.7 $\pm$ 0.3	56.6 $\pm$ 16.3	80.7 $\pm$ 0.3	54.1 $\pm$ 14.0	80.2 $\pm$ 0.4	60.4 $\pm$ 17.1	70.2 $\pm$ 16.1
		F:	86.8 $\pm$ 0.8	86.8 $\pm$ 0.8	79.2 $\pm$ 7.4	86.8 $\pm$ 0.8	80.8 $\pm$ 4.0	86.8 $\pm$ 0.8	79.6 $\pm$ 7.7	85.2 $\pm$ 3.1
$N = 5, N - K = 2$	adult	A:	80.6 $\pm$ 0.4	80.6 $\pm$ 0.4	66.8 $\pm$ 23.3	80.6 $\pm$ 0.4	55.1 $\pm$ 28.2	80.1 $\pm$ 0.3	66.7 $\pm$ 23.1	61.2 $\pm$ 27.0
		F:	87.0 $\pm$ 1.0	87.0 $\pm$ 1.0	77.0 $\pm$ 20.1	87.0 $\pm$ 1.0	66.9 $\pm$ 24.6	87.0 $\pm$ 1.0	77.0 $\pm$ 19.8	73.6 $\pm$ 23.5
$N = 7, N - K = 3$	adult	A:	80.6 $\pm$ 0.3	80.6 $\pm$ 0.3	45.5 $\pm$ 15.4	80.6 $\pm$ 0.3	64.4 $\pm$ 19.7	80.0 $\pm$ 0.3	47.1 $\pm$ 16.5	60.2 $\pm$ 15.4
		F:	86.8 $\pm$ 1.0	86.8 $\pm$ 1.0	65.9 $\pm$ 28.8	86.8 $\pm$ 1.0	76.2 $\pm$ 23.7	86.8 $\pm$ 1.0	72.1 $\pm$ 21.5	83.1 $\pm$ 4.7
$N = 9, N - K = 4$	adult	A:	80.7 $\pm$ 0.3	80.6 $\pm$ 0.3	61.1 $\pm$ 25.7	80.7 $\pm$ 0.3	61.2 $\pm$ 25.6	80.0 $\pm$ 0.4	55.4 $\pm$ 27.9	55.2 $\pm$ 28.2
		F:	86.7 $\pm$ 1.0	86.7 $\pm$ 1.0	76.8 $\pm$ 15.3	86.7 $\pm$ 1.0	76.8 $\pm$ 14.9	86.7 $\pm$ 1.0	68.4 $\pm$ 26.6	70.2 $\pm$ 27.5
$N = 11, N - K = 5$	adult	A:	80.6 $\pm$ 0.4	80.5 $\pm$ 0.3	52.1 $\pm$ 20.9	80.6 $\pm$ 0.4	67.8 $\pm$ 18.6	80.0 $\pm$ 0.3	57.1 $\pm$ 23.1	53.0 $\pm$ 24.2
		F:	87.3 $\pm$ 1.2	87.3 $\pm$ 1.2	72.7 $\pm$ 18.8	87.3 $\pm$ 1.2	80.9 $\pm$ 15.0	87.3 $\pm$ 1.2	73.5 $\pm$ 19.2	70.1 $\pm$ 22.9
$N = 3, N - K = 1$	compas	A:	66.0 $\pm$ 1.0	65.9 $\pm$ 1.0	59.2 $\pm$ 9.2	66.0 $\pm$ 1.0	54.6 $\pm$ 7.9	65.8 $\pm$ 1.0	57.8 $\pm$ 9.4	61.4 $\pm$ 7.8
		F:	85.6 $\pm$ 3.7	83.7 $\pm$ 4.5	78.1 $\pm$ 11.8	85.6 $\pm$ 3.7	72.3 $\pm$ 10.7	83.5 $\pm$ 4.3	76.4 $\pm$ 11.0	79.6 $\pm$ 9.3
$N = 5, N - K = 2$	compas	A:	65.9 $\pm$ 1.0	65.8 $\pm$ 1.1	54.5 $\pm$ 12.3	65.9 $\pm$ 1.0	43.5 $\pm$ 12.5	65.7 $\pm$ 1.1	54.4 $\pm$ 12.3	54.5 $\pm$ 12.7
		F:	85.1 $\pm$ 3.5	83.1 $\pm$ 3.2	64.2 $\pm$ 32.4	85.1 $\pm$ 3.5	36.0 $\pm$ 30.2	82.3 $\pm$ 4.0	64.2 $\pm$ 32.4	62.6 $\pm$ 31.5
$N = 7, N - K = 3$	compas	A:	65.9 $\pm$ 1.1	65.8 $\pm$ 1.1	51.5 $\pm$ 9.1	65.9 $\pm$ 1.1	58.2 $\pm$ 8.2	65.7 $\pm$ 1.3	51.7 $\pm$ 7.8	54.5 $\pm$ 9.1
		F:	80.9 $\pm$ 6.1	83.0 $\pm$ 5.2	64.8 $\pm$ 23.6	81.7 $\pm$ 5.1	80.1 $\pm$ 6.3	83.0 $\pm$ 4.4	63.2 $\pm$ 22.7	73.1 $\pm$ 8.5
$N = 9, N - K = 4$	compas	A:	65.9 $\pm$ 1.2	66.0 $\pm$ 1.1	58.6 $\pm$ 8.6	66.0 $\pm$ 1.2	54.4 $\pm$ 10.5	65.8 $\pm$ 1.2	57.1 $\pm$ 10.7	51.8 $\pm$ 13.2
		F:	84.1 $\pm$ 3.6	84.0 $\pm$ 3.3	76.9 $\pm$ 16.8	84.3 $\pm$ 3.2	66.8 $\pm$ 25.0	83.8 $\pm$ 3.3	71.7 $\pm$ 20.9	60.4 $\pm$ 28.2
$N = 11, N - K = 5$	compas	A:	65.6 $\pm$ 1.1	65.7 $\pm$ 1.0	51.2 $\pm$ 8.8	65.3 $\pm$ 0.9	50.9 $\pm$ 11.1	65.5 $\pm$ 0.9	53.8 $\pm$ 7.1	51.0 $\pm$ 12.3
		F:	84.4 $\pm$ 3.9	82.4 $\pm$ 4.6	67.7 $\pm$ 22.8	84.0 $\pm$ 3.6	62.3 $\pm$ 29.2	82.1 $\pm$ 4.4	75.1 $\pm$ 10.7	58.1 $\pm$ 30.5
$N = 3, N - K = 1$	drugs	A:	65.7 $\pm$ 1.9	65.5 $\pm$ 2.0	55.9 $\pm$ 7.1	65.7 $\pm$ 1.9	51.2 $\pm$ 4.9	65.6 $\pm$ 2.0	58.5 $\pm$ 7.6	61.4 $\pm$ 6.8
		F:	87.7 $\pm$ 3.0	83.3 $\pm$ 3.7	74.5 $\pm$ 11.6	87.7 $\pm$ 3.0	72.4 $\pm$ 7.3	84.9 $\pm$ 3.5	76.4 $\pm$ 12.3	83.3 $\pm$ 7.9
$N = 5, N - K = 2$	drugs	A:	65.1 $\pm$ 2.0	65.1 $\pm$ 2.0	59.5 $\pm$ 8.9	65.1 $\pm$ 2.0	56.5 $\pm$ 10.0	65.0 $\pm$ 1.9	57.7 $\pm$ 8.8	54.1 $\pm$ 11.6
		F:	86.7 $\pm$ 2.5	83.9 $\pm$ 3.8	70.6 $\pm$ 31.2	86.7 $\pm$ 2.5	66.1 $\pm$ 28.4	83.2 $\pm$ 4.1	72.0 $\pm$ 28.9	61.0 $\pm$ 32.7
$N = 7, N - K = 3$	drugs	A:	64.6 $\pm$ 2.2	65.2 $\pm$ 1.6	55.3 $\pm$ 9.6	65.0 $\pm$ 1.8	50.4 $\pm$ 9.8	65.0 $\pm$ 1.8	53.0 $\pm$ 10.9	54.3 $\pm$ 8.5
		F:	87.1 $\pm$ 3.9	80.8 $\pm$ 5.0	66.0 $\pm$ 24.6	84.6 $\pm$ 4.1	63.1 $\pm$ 21.8	81.7 $\pm$ 4.4	59.6 $\pm$ 31.0	70.1 $\pm$ 10.0
$N = 9, N - K = 4$	drugs	A:	62.2 $\pm$ 2.6	64.3 $\pm$ 1.5	52.4 $\pm$ 8.0	62.3 $\pm$ 2.4	50.7 $\pm$ 10.4	64.2 $\pm$ 1.8	49.1 $\pm$ 8.9	51.9 $\pm$ 10.9
		F:	83.1 $\pm$ 3.7	79.4 $\pm$ 6.0	58.7 $\pm$ 25.5	82.1 $\pm$ 4.1	53.5 $\pm$ 30.1	80.8 $\pm$ 4.7	53.7 $\pm$ 25.9	59.6 $\pm$ 30.7
$N = 11, N - K = 5$	drugs	A:	59.6 $\pm$ 7.6	62.7 $\pm$ 3.4	53.7 $\pm$ 6.9	56.2 $\pm$ 9.0	47.4 $\pm$ 10.4	64.7 $\pm$ 1.8	51.5 $\pm$ 7.4	52.1 $\pm$ 11.2
		F:	83.1 $\pm$ 4.5	78.3 $\pm$ 4.2	69.1 $\pm$ 18.5	82.2 $\pm$ 2.8	46.9 $\pm$ 30.2	79.3 $\pm$ 3.4	60.0 $\pm$ 23.0	56.9 $\pm$ 34.5
$N = 3, N - K = 1$	germancredit	A:	73.2 $\pm$ 3.6	72.8 $\pm$ 3.3	55.5 $\pm$ 12.9	73.2 $\pm$ 3.6	55.7 $\pm$ 12.4	72.7 $\pm$ 3.2	56.4 $\pm$ 13.3	64.1 $\pm$ 14.2
		F:	93.9 $\pm$ 2.8	93.4 $\pm$ 3.4	59.7 $\pm$ 19.8	93.9 $\pm$ 2.8	67.5 $\pm$ 20.2	93.5 $\pm$ 3.4	59.1 $\pm$ 19.5	82.1 $\pm$ 20.7
$N = 5, N - K = 2$	germancredit	A:	72.4 $\pm$ 3.4	72.4 $\pm$ 3.4	61.0 $\pm$ 13.4	72.6 $\pm$ 3.1	66.0 $\pm$ 10.8	72.2 $\pm$ 2.9	53.8 $\pm$ 16.7	60.0 $\pm$ 15.7
		F:	92.8 $\pm$ 4.1	92.0 $\pm$ 4.9	69.7 $\pm$ 31.1	92.5 $\pm$ 4.4	76.6 $\pm$ 21.1	91.7 $\pm$ 3.9	55.4 $\pm$ 33.6	65.4 $\pm$ 33.0
$N = 7, N - K = 3$	germancredit	A:	69.5 $\pm$ 6.7	72.8 $\pm$ 3.4	56.0 $\pm$ 14.6	72.1 $\pm$ 3.4	52.9 $\pm$ 11.0	72.2 $\pm$ 3.7	54.5 $\pm$ 13.3	57.6 $\pm$ 11.8
		F:	92.4 $\pm$ 4.0	93.3 $\pm$ 4.4	61.5 $\pm$ 26.3	93.5 $\pm$ 4.7	55.6 $\pm$ 20.0	91.9 $\pm$ 4.2	67.2 $\pm$ 14.9	65.7 $\pm$ 21.7
$N = 9, N - K = 4$	germancredit	A:	70.5 $\pm$ 3.5	72.2 $\pm$ 3.9	56.0 $\pm$ 11.4	67.6 $\pm$ 3.2	47.5 $\pm$ 14.4	71.1 $\pm$ 3.1	47.6 $\pm$ 12.8	53.5 $\pm$ 15.9
		F:	90.8 $\pm$ 5.1	93.4 $\pm$ 2.5	63.8 $\pm$ 24.7	88.2 $\pm$ 5.8	50.6 $\pm$ 31.7	90.4 $\pm$ 5.1	48.8 $\pm$ 30.8	56.7 $\pm$ 32.6
$N = 11, N - K = 5$	germancredit	A:	66.8 $\pm$ 7.8	71.3 $\pm$ 3.8	61.2 $\pm$ 12.2	66.2 $\pm$ 6.0	55.6 $\pm$ 12.9	71.3 $\pm$ 3.3	58.1 $\pm$ 11.2	53.1 $\pm$ 15.5
		F:	82.0 $\pm$ 11.4	89.6 $\pm$ 4.4	70.2 $\pm$ 20.4	83.3 $\pm$ 7.5	60.9 $\pm$ 21.1	91.0 $\pm$ 4.8	63.3 $\pm$ 22.8	56.2 $\pm$ 31.0
$N = 51, N - K = 5$	folktables	A:	78.6 $\pm$ 0.2	78.6 $\pm$ 0.2	77.9 $\pm$ 0.4	78.6 $\pm$ 0.2	77.7 $\pm$ 0.4	78.7 $\pm$ 0.2	77.8 $\pm$ 0.5	77.8 $\pm$ 0.6
		F:	86.4 $\pm$ 0.6	86.2 $\pm$ 0.5	86.5 $\pm$ 0.5	86.3 $\pm$ 0.6	86.3 $\pm$ 0.5	86.0 $\pm$ 0.5	86.5 $\pm$ 0.5	86.9 $\pm$ 0.5
$N = 51, N - K = 10$	folktables	A:	78.7 $\pm$ 0.2	78.6 $\pm$ 0.2	76.8 $\pm$ 0.4	78.7 $\pm$ 0.2	77.0 $\pm$ 0.4	78.6 $\pm$ 0.2	76.8 $\pm$ 0.5	76.8 $\pm$ 0.9
		F:	86.3 $\pm$ 0.6	86.1 $\pm$ 0.6	86.4 $\pm$ 0.6	86.3 $\pm$ 0.6	86.5 $\pm$ 0.4	86.0 $\pm$ 0.6	86.4 $\pm$ 0.5	87.6 $\pm$ 0.5
$N = 51, N - K = 15$	folktables	A:	78.6 $\pm$ 0.2	78.5 $\pm$ 0.2	75.9 $\pm$ 0.7	78.6 $\pm$ 0.2	76.1 $\pm$ 0.5	78.5 $\pm$ 0.3	76.0 $\pm$ 0.5	72.3 $\pm$ 7.7
		F:	86.2 $\pm$ 0.5	86.0 $\pm$ 0.6	86.1 $\pm$ 1.0	86.2 $\pm$ 0.6	86.3 $\pm$ 0.6	86.0 $\pm$ 0.5	86.3 $\pm$ 0.6	85.5 $\pm$ 4.7
$N = 51, N - K = 20$	folktables	A:	78.5 $\pm$ 0.3	78.4 $\pm$ 0.2	73.0 $\pm$ 4.0	78.5 $\pm$ 0.3	73.8 $\pm$ 3.5	78.4 $\pm$ 0.2	70.1 $\pm$ 9.0	67.6 $\pm$ 13.3
		F:	86.2 $\pm$ 0.5	86.1 $\pm$ 0.6	83.6 $\pm$ 3.3	86.2 $\pm$ 0.6	85.7 $\pm$ 3.6	86.1 $\pm$ 0.6	81.6 $\pm$ 6.7	82.5 $\pm$ 8.5
$N = 51, N - K = 25$	folktables	A:	67.0 $\pm$ 13.3	78.3 $\pm$ 0.2	60.0 $\pm$ 16.6	66.1 $\pm$ 15.3	40.1 $\pm$ 16.9	78.3 $\pm$ 0.2	54.0 $\pm$ 16.2	60.0 $\pm$ 15.5
		F:	71.2 $\pm$ 21.4	86.7 $\pm$ 0.5	75.9 $\pm$ 11.9	71.3 $\pm$ 18.0	66.6 $\pm$ 16.1	86.7 $\pm$ 0.5	74.7 $\pm$ 7.3	74.8 $\pm$ 10.8

(d) preprocessing-based fairness

		disb disc disp	disb disc <del>disp</del>	disb <del>disc</del> disp	<del>disb</del> disc disp	disb <del>disc</del> <del>disp</del>	<del>disb</del> disc <del>disp</del>	<del>disb</del> <del>disc</del> disp	<del>disb</del> <del>disc</del> <del>disp</del>
$N = 3, N - K = 1$	adult A:	79.5 $\pm$ 0.3	79.5 $\pm$ 0.3	56.9 $\pm$ 15.0	79.5 $\pm$ 0.3	53.1 $\pm$ 13.3	79.5 $\pm$ 0.3	60.4 $\pm$ 16.3	69.3 $\pm$ 15.5
	F:	94.2 $\pm$ 1.1	94.4 $\pm$ 0.7	94.3 $\pm$ 0.9	94.2 $\pm$ 1.3	93.6 $\pm$ 1.1	92.9 $\pm$ 1.0	94.1 $\pm$ 1.2	93.4 $\pm$ 1.8
$N = 5, N - K = 2$	adult A:	79.5 $\pm$ 0.3	79.5 $\pm$ 0.3	67.1 $\pm$ 23.1	79.4 $\pm$ 0.3	55.5 $\pm$ 28.1	79.5 $\pm$ 0.3	67.0 $\pm$ 22.9	61.1 $\pm$ 26.9
	F:	94.0 $\pm$ 0.9	94.2 $\pm$ 1.2	94.3 $\pm$ 1.2	94.3 $\pm$ 1.1	93.3 $\pm$ 0.8	92.4 $\pm$ 1.1	94.0 $\pm$ 1.2	92.4 $\pm$ 1.6
$N = 7, N - K = 3$	adult A:	79.5 $\pm$ 0.3	79.5 $\pm$ 0.3	43.0 $\pm$ 11.2	79.5 $\pm$ 0.3	65.2 $\pm$ 18.8	79.5 $\pm$ 0.3	49.8 $\pm$ 13.4	59.9 $\pm$ 15.5
	F:	94.4 $\pm$ 1.2	94.1 $\pm$ 1.4	94.2 $\pm$ 1.3	94.1 $\pm$ 1.3	93.3 $\pm$ 1.3	92.1 $\pm$ 1.3	94.2 $\pm$ 1.4	92.5 $\pm$ 1.7
$N = 9, N - K = 4$	adult A:	79.5 $\pm$ 0.3	79.5 $\pm$ 0.3	61.2 $\pm$ 25.3	79.5 $\pm$ 0.3	61.2 $\pm$ 25.4	79.5 $\pm$ 0.3	55.6 $\pm$ 27.1	55.4 $\pm$ 27.6
	F:	94.5 $\pm$ 1.4	94.3 $\pm$ 1.3	94.5 $\pm$ 1.4	94.1 $\pm$ 1.2	93.6 $\pm$ 0.8	92.4 $\pm$ 1.2	94.2 $\pm$ 1.4	92.6 $\pm$ 1.8
$N = 11, N - K = 5$	adult A:	79.4 $\pm$ 0.3	79.5 $\pm$ 0.4	53.4 $\pm$ 20.4	79.4 $\pm$ 0.3	66.6 $\pm$ 19.3	79.4 $\pm$ 0.3	57.4 $\pm$ 22.6	52.5 $\pm$ 24.3
	F:	94.3 $\pm$ 1.0	94.2 $\pm$ 0.9	94.6 $\pm$ 0.8	94.3 $\pm$ 1.1	93.6 $\pm$ 1.2	92.3 $\pm$ 1.6	94.5 $\pm$ 0.8	91.9 $\pm$ 1.8
$N = 3, N - K = 1$	compas A:	65.1 $\pm$ 1.0	65.0 $\pm$ 1.2	58.8 $\pm$ 9.1	65.1 $\pm$ 1.0	53.5 $\pm$ 8.9	65.1 $\pm$ 1.2	58.4 $\pm$ 8.1	58.7 $\pm$ 10.0
	F:	94.1 $\pm$ 2.9	92.5 $\pm$ 4.4	94.1 $\pm$ 1.9	93.7 $\pm$ 2.5	91.1 $\pm$ 3.6	89.3 $\pm$ 4.6	93.1 $\pm$ 3.5	91.8 $\pm$ 3.5
$N = 5, N - K = 2$	compas A:	65.1 $\pm$ 1.1	65.0 $\pm$ 0.8	53.5 $\pm$ 11.2	65.0 $\pm$ 1.0	43.7 $\pm$ 10.7	64.9 $\pm$ 1.2	52.6 $\pm$ 11.3	53.3 $\pm$ 12.0
	F:	92.1 $\pm$ 2.5	91.2 $\pm$ 3.4	91.8 $\pm$ 2.3	91.4 $\pm$ 2.8	88.2 $\pm$ 3.0	88.2 $\pm$ 2.5	92.6 $\pm$ 2.4	87.5 $\pm$ 3.0
$N = 7, N - K = 3$	compas A:	65.1 $\pm$ 1.3	64.9 $\pm$ 0.7	50.8 $\pm$ 8.8	65.0 $\pm$ 1.0	57.8 $\pm$ 7.7	65.0 $\pm$ 1.0	53.5 $\pm$ 7.2	52.9 $\pm$ 8.9
	F:	91.8 $\pm$ 2.4	91.0 $\pm$ 3.9	90.8 $\pm$ 2.5	89.6 $\pm$ 4.6	92.3 $\pm$ 2.5	86.9 $\pm$ 3.6	91.9 $\pm$ 1.8	87.8 $\pm$ 3.4
$N = 9, N - K = 4$	compas A:	65.2 $\pm$ 1.2	65.0 $\pm$ 0.6	57.2 $\pm$ 7.7	64.9 $\pm$ 0.8	53.9 $\pm$ 8.6	65.1 $\pm$ 0.9	56.5 $\pm$ 8.9	50.8 $\pm$ 11.3
	F:	91.7 $\pm$ 4.3	92.6 $\pm$ 2.8	92.1 $\pm$ 2.8	91.6 $\pm$ 4.3	93.4 $\pm$ 1.9	87.9 $\pm$ 1.9	91.2 $\pm$ 2.9	86.7 $\pm$ 3.8
$N = 11, N - K = 5$	compas A:	64.9 $\pm$ 0.8	65.2 $\pm$ 1.0	49.6 $\pm$ 9.0	64.8 $\pm$ 0.9	50.4 $\pm$ 9.7	65.0 $\pm$ 0.8	53.2 $\pm$ 7.1	49.4 $\pm$ 11.7
	F:	92.6 $\pm$ 3.6	91.6 $\pm$ 3.7	90.1 $\pm$ 2.8	89.9 $\pm$ 2.9	92.3 $\pm$ 4.3	87.2 $\pm$ 4.5	90.0 $\pm$ 2.3	86.5 $\pm$ 4.6
$N = 3, N - K = 1$	drugs A:	64.0 $\pm$ 1.8	64.2 $\pm$ 1.3	55.1 $\pm$ 6.7	63.5 $\pm$ 1.4	50.8 $\pm$ 5.3	63.7 $\pm$ 2.5	57.1 $\pm$ 6.4	58.5 $\pm$ 8.3
	F:	89.2 $\pm$ 4.3	87.0 $\pm$ 4.2	87.4 $\pm$ 3.0	88.5 $\pm$ 4.0	79.0 $\pm$ 4.7	85.8 $\pm$ 4.7	87.5 $\pm$ 4.2	84.7 $\pm$ 3.5
$N = 5, N - K = 2$	drugs A:	62.9 $\pm$ 1.8	63.4 $\pm$ 2.1	57.2 $\pm$ 7.3	62.8 $\pm$ 1.9	55.8 $\pm$ 8.7	63.0 $\pm$ 1.8	57.1 $\pm$ 7.0	52.9 $\pm$ 10.7
	F:	90.3 $\pm$ 3.0	85.5 $\pm$ 4.9	87.3 $\pm$ 2.8	87.7 $\pm$ 3.3	79.3 $\pm$ 6.7	83.1 $\pm$ 4.9	87.0 $\pm$ 1.8	80.9 $\pm$ 6.5
$N = 7, N - K = 3$	drugs A:	62.8 $\pm$ 1.7	63.1 $\pm$ 2.1	52.6 $\pm$ 9.9	62.6 $\pm$ 1.2	49.0 $\pm$ 8.0	62.7 $\pm$ 2.0	54.3 $\pm$ 9.0	53.2 $\pm$ 7.7
	F:	89.5 $\pm$ 2.6	82.3 $\pm$ 5.8	87.8 $\pm$ 3.1	84.5 $\pm$ 5.8	76.7 $\pm$ 3.5	84.7 $\pm$ 4.5	86.3 $\pm$ 2.6	78.1 $\pm$ 5.2
$N = 9, N - K = 4$	drugs A:	59.6 $\pm$ 3.7	63.1 $\pm$ 2.1	52.2 $\pm$ 6.9	60.5 $\pm$ 3.1	51.2 $\pm$ 9.8	62.7 $\pm$ 2.4	49.8 $\pm$ 7.5	51.9 $\pm$ 9.7
	F:	86.1 $\pm$ 3.9	79.3 $\pm$ 6.5	86.3 $\pm$ 4.1	84.3 $\pm$ 4.2	76.0 $\pm$ 7.4	80.9 $\pm$ 4.6	84.5 $\pm$ 2.5	77.4 $\pm$ 7.9
$N = 11, N - K = 5$	drugs A:	56.9 $\pm$ 6.8	59.8 $\pm$ 4.9	55.5 $\pm$ 6.6	55.0 $\pm$ 7.8	45.8 $\pm$ 8.4	62.6 $\pm$ 2.2	51.5 $\pm$ 5.5	49.5 $\pm$ 10.0
	F:	84.9 $\pm$ 6.3	78.1 $\pm$ 6.7	83.5 $\pm$ 4.7	85.4 $\pm$ 4.0	76.3 $\pm$ 5.6	79.8 $\pm$ 5.4	83.3 $\pm$ 5.2	76.7 $\pm$ 8.6
$N = 3, N - K = 1$	germancredit A:	70.5 $\pm$ 3.1	70.8 $\pm$ 3.2	54.1 $\pm$ 12.2	70.5 $\pm$ 3.4	55.7 $\pm$ 12.0	71.0 $\pm$ 3.3	53.0 $\pm$ 12.9	60.6 $\pm$ 14.7
	F:	90.2 $\pm$ 4.3	88.7 $\pm$ 4.5	86.2 $\pm$ 6.3	91.3 $\pm$ 3.9	89.5 $\pm$ 4.1	90.7 $\pm$ 5.1	85.5 $\pm$ 4.4	88.3 $\pm$ 4.3
$N = 5, N - K = 2$	germancredit A:	70.0 $\pm$ 1.8	70.6 $\pm$ 3.4	60.4 $\pm$ 9.9	70.0 $\pm$ 2.9	64.3 $\pm$ 8.8	70.0 $\pm$ 3.7	50.4 $\pm$ 14.6	56.4 $\pm$ 13.6
	F:	88.3 $\pm$ 5.6	91.5 $\pm$ 4.0	88.8 $\pm$ 2.8	87.8 $\pm$ 4.1	85.2 $\pm$ 4.3	89.2 $\pm$ 5.9	87.4 $\pm$ 4.6	83.0 $\pm$ 6.1
$N = 7, N - K = 3$	germancredit A:	66.2 $\pm$ 7.0	70.2 $\pm$ 3.5	56.6 $\pm$ 12.5	70.0 $\pm$ 3.9	52.8 $\pm$ 9.7	70.2 $\pm$ 3.2	54.1 $\pm$ 12.4	55.5 $\pm$ 12.2
	F:	89.0 $\pm$ 5.3	89.3 $\pm$ 4.9	83.8 $\pm$ 4.3	86.7 $\pm$ 6.6	84.4 $\pm$ 6.6	87.2 $\pm$ 5.5	84.4 $\pm$ 5.5	83.5 $\pm$ 5.3
$N = 9, N - K = 4$	germancredit A:	67.9 $\pm$ 3.3	70.3 $\pm$ 2.5	54.8 $\pm$ 10.1	65.8 $\pm$ 3.7	49.4 $\pm$ 11.5	69.5 $\pm$ 2.9	46.8 $\pm$ 11.3	51.9 $\pm$ 13.8
	F:	90.0 $\pm$ 4.0	89.6 $\pm$ 3.6	85.6 $\pm$ 5.1	87.8 $\pm$ 4.5	84.3 $\pm$ 3.7	87.4 $\pm$ 2.7	84.5 $\pm$ 4.6	85.4 $\pm$ 5.0
$N = 11, N - K = 5$	germancredit A:	64.9 $\pm$ 7.3	69.0 $\pm$ 3.6	57.0 $\pm$ 10.5	63.9 $\pm$ 8.1	53.1 $\pm$ 12.1	69.8 $\pm$ 3.2	56.1 $\pm$ 9.8	51.9 $\pm$ 13.4
	F:	87.2 $\pm$ 4.7	86.3 $\pm$ 4.1	85.5 $\pm$ 6.0	85.4 $\pm$ 6.8	86.8 $\pm$ 4.5	87.7 $\pm$ 5.5	85.5 $\pm$ 5.3	84.3 $\pm$ 6.5
$N = 51, N - K = 5$	folktables A:	78.1 $\pm$ 0.2	78.1 $\pm$ 0.2	77.6 $\pm$ 0.3	78.1 $\pm$ 0.2	77.5 $\pm$ 0.3	78.1 $\pm$ 0.2	77.6 $\pm$ 0.4	77.5 $\pm$ 0.4
	F:	92.2 $\pm$ 0.4	92.1 $\pm$ 0.4	92.4 $\pm$ 0.5	92.3 $\pm$ 0.6	91.8 $\pm$ 0.4	92.1 $\pm$ 0.4	92.4 $\pm$ 0.4	91.8 $\pm$ 0.8
$N = 51, N - K = 10$	folktables A:	78.1 $\pm$ 0.2	78.1 $\pm$ 0.2	77.0 $\pm$ 0.3	78.1 $\pm$ 0.2	77.1 $\pm$ 0.2	78.1 $\pm$ 0.2	77.0 $\pm$ 0.3	76.7 $\pm$ 0.9
	F:	92.3 $\pm$ 0.4	92.1 $\pm$ 0.4	92.4 $\pm$ 0.5	92.2 $\pm$ 0.6	91.9 $\pm$ 0.5	92.1 $\pm$ 0.6	92.3 $\pm$ 0.4	91.0 $\pm$ 1.0
$N = 51, N - K = 15$	folktables A:	78.1 $\pm$ 0.2	78.1 $\pm$ 0.2	76.2 $\pm$ 0.8	78.1 $\pm$ 0.2	76.2 $\pm$ 0.6	78.1 $\pm$ 0.2	76.2 $\pm$ 0.7	72.4 $\pm$ 7.6
	F:	92.3 $\pm$ 0.4	91.9 $\pm$ 0.3	92.3 $\pm$ 0.5	92.2 $\pm$ 0.5	91.4 $\pm$ 0.4	91.9 $\pm$ 0.6	92.4 $\pm$ 0.5	90.0 $\pm$ 1.4
$N = 51, N - K = 20$	folktables A:	78.1 $\pm$ 0.2	78.1 $\pm$ 0.2	73.1 $\pm$ 4.0	78.0 $\pm$ 0.2	73.8 $\pm$ 3.4	78.1 $\pm$ 0.2	70.0 $\pm$ 9.0	67.5 $\pm$ 13.5
	F:	92.1 $\pm$ 0.3	91.6 $\pm$ 0.3	92.0 $\pm$ 0.4	92.1 $\pm$ 0.5	91.3 $\pm$ 0.6	91.1 $\pm$ 0.5	91.8 $\pm$ 0.6	88.7 $\pm$ 1.3
$N = 51, N - K = 25$	folktables A:	69.3 $\pm$ 12.3	78.0 $\pm$ 0.2	59.8 $\pm$ 16.2	67.9 $\pm$ 15.2	40.5 $\pm$ 16.7	78.0 $\pm$ 0.2	54.0 $\pm$ 15.3	59.6 $\pm$ 15.2
	F:	90.4 $\pm$ 1.0	90.3 $\pm$ 0.6	90.8 $\pm$ 1.2	90.5 $\pm$ 1.1	89.9 $\pm$ 1.0	90.2 $\pm$ 0.7	89.9 $\pm$ 2.0	88.2 $\pm$ 1.3

(e) fairness-unaware

			disb disc disp	disb disc <del>disp</del>	disb <del>disc</del> disp	<del>disb</del> disc disp	disb <del>disc</del> <del>disp</del>	<del>disb</del> disc <del>disp</del>	<del>disb</del> <del>disc</del> disp	<del>disb</del> <del>disc</del> <del>disp</del>
$N = 3, N - K = 1$	adult	A:	80.9 $\pm$ 0.4	80.9 $\pm$ 0.4	56.8 $\pm$ 16.6	80.9 $\pm$ 0.4	54.1 $\pm$ 14.2	80.2 $\pm$ 0.5	60.6 $\pm$ 17.4	70.6 $\pm$ 15.9
		F:	80.5 $\pm$ 0.8	80.5 $\pm$ 0.8	50.2 $\pm$ 20.3	80.5 $\pm$ 0.8	45.8 $\pm$ 17.5	80.5 $\pm$ 0.8	54.4 $\pm$ 21.7	67.7 $\pm$ 19.5
$N = 5, N - K = 2$	adult	A:	81.0 $\pm$ 0.4	81.0 $\pm$ 0.4	63.0 $\pm$ 22.0	81.0 $\pm$ 0.4	51.5 $\pm$ 26.2	80.2 $\pm$ 0.5	62.4 $\pm$ 21.2	58.6 $\pm$ 25.8
		F:	80.7 $\pm$ 1.1	80.7 $\pm$ 1.1	61.4 $\pm$ 30.3	80.7 $\pm$ 1.1	45.7 $\pm$ 36.5	80.7 $\pm$ 1.1	60.7 $\pm$ 30.0	54.7 $\pm$ 35.0
$N = 7, N - K = 3$	adult	A:	81.0 $\pm$ 0.4	81.0 $\pm$ 0.3	44.0 $\pm$ 13.7	81.0 $\pm$ 0.4	61.9 $\pm$ 18.9	79.9 $\pm$ 0.6	46.5 $\pm$ 16.6	59.2 $\pm$ 14.5
		F:	80.8 $\pm$ 1.0	80.8 $\pm$ 1.0	30.2 $\pm$ 15.1	80.8 $\pm$ 1.0	58.3 $\pm$ 26.9	80.8 $\pm$ 1.0	38.2 $\pm$ 18.0	55.9 $\pm$ 22.9
$N = 9, N - K = 4$	adult	A:	81.0 $\pm$ 0.3	81.0 $\pm$ 0.3	56.2 $\pm$ 23.9	81.0 $\pm$ 0.3	57.0 $\pm$ 24.4	79.9 $\pm$ 0.6	51.4 $\pm$ 25.6	52.1 $\pm$ 26.5
		F:	80.5 $\pm$ 1.1	80.5 $\pm$ 1.1	49.0 $\pm$ 27.2	80.5 $\pm$ 1.1	51.5 $\pm$ 28.7	80.5 $\pm$ 1.1	43.4 $\pm$ 30.4	46.4 $\pm$ 33.0
$N = 11, N - K = 5$	adult	A:	80.9 $\pm$ 0.3	80.8 $\pm$ 0.3	50.9 $\pm$ 20.1	80.9 $\pm$ 0.3	64.3 $\pm$ 17.8	80.1 $\pm$ 0.5	53.2 $\pm$ 20.5	51.9 $\pm$ 23.5
		F:	81.1 $\pm$ 1.0	81.0 $\pm$ 1.0	43.0 $\pm$ 27.7	81.1 $\pm$ 1.0	61.6 $\pm$ 26.2	81.0 $\pm$ 1.0	44.6 $\pm$ 28.6	43.4 $\pm$ 32.9
$N = 3, N - K = 1$	compas	A:	66.4 $\pm$ 1.1	66.4 $\pm$ 1.1	58.8 $\pm$ 10.3	66.4 $\pm$ 1.1	53.5 $\pm$ 8.5	66.2 $\pm$ 1.1	57.3 $\pm$ 10.3	61.3 $\pm$ 8.5
		F:	72.2 $\pm$ 4.7	72.2 $\pm$ 4.7	44.6 $\pm$ 33.4	72.2 $\pm$ 4.7	25.8 $\pm$ 34.2	72.2 $\pm$ 4.7	38.2 $\pm$ 32.3	50.8 $\pm$ 30.3
$N = 5, N - K = 2$	compas	A:	66.5 $\pm$ 1.1	66.4 $\pm$ 1.2	49.7 $\pm$ 9.7	66.5 $\pm$ 1.1	40.8 $\pm$ 9.1	65.9 $\pm$ 1.3	49.7 $\pm$ 9.7	51.8 $\pm$ 11.7
		F:	71.4 $\pm$ 4.8	71.4 $\pm$ 4.8	28.2 $\pm$ 20.5	71.4 $\pm$ 4.8	8.4 $\pm$ 12.8	70.9 $\pm$ 5.4	28.2 $\pm$ 20.5	31.5 $\pm$ 24.6
$N = 7, N - K = 3$	compas	A:	66.4 $\pm$ 1.2	66.4 $\pm$ 1.2	49.8 $\pm$ 8.0	66.3 $\pm$ 1.0	55.4 $\pm$ 7.1	66.0 $\pm$ 1.2	49.2 $\pm$ 6.6	52.2 $\pm$ 7.3
		F:	70.2 $\pm$ 4.2	70.2 $\pm$ 4.5	16.8 $\pm$ 23.2	70.4 $\pm$ 4.3	36.9 $\pm$ 25.9	70.5 $\pm$ 4.3	11.4 $\pm$ 18.5	27.7 $\pm$ 23.9
$N = 9, N - K = 4$	compas	A:	66.3 $\pm$ 1.0	66.2 $\pm$ 1.0	53.5 $\pm$ 7.1	66.2 $\pm$ 1.0	49.3 $\pm$ 6.8	65.8 $\pm$ 1.1	53.9 $\pm$ 9.2	49.1 $\pm$ 11.1
		F:	69.7 $\pm$ 3.7	69.8 $\pm$ 3.7	31.2 $\pm$ 20.9	69.8 $\pm$ 3.7	16.0 $\pm$ 11.0	69.8 $\pm$ 3.7	36.2 $\pm$ 25.4	25.6 $\pm$ 25.5
$N = 11, N - K = 5$	compas	A:	66.2 $\pm$ 1.0	66.2 $\pm$ 1.0	48.4 $\pm$ 6.3	66.0 $\pm$ 0.8	48.3 $\pm$ 9.3	66.1 $\pm$ 1.0	50.7 $\pm$ 5.3	49.2 $\pm$ 11.4
		F:	69.7 $\pm$ 5.0	70.5 $\pm$ 4.3	15.2 $\pm$ 18.1	69.6 $\pm$ 4.6	20.1 $\pm$ 22.0	70.4 $\pm$ 4.5	19.0 $\pm$ 18.5	23.8 $\pm$ 24.9
$N = 3, N - K = 1$	drugs	A:	66.7 $\pm$ 2.7	66.2 $\pm$ 2.4	50.9 $\pm$ 10.7	66.7 $\pm$ 2.7	48.7 $\pm$ 7.6	66.5 $\pm$ 2.7	55.5 $\pm$ 11.3	60.0 $\pm$ 10.6
		F:	65.2 $\pm$ 4.2	65.2 $\pm$ 4.2	24.6 $\pm$ 26.3	65.2 $\pm$ 4.2	10.8 $\pm$ 18.3	65.2 $\pm$ 4.2	35.4 $\pm$ 29.1	47.5 $\pm$ 26.7
$N = 5, N - K = 2$	drugs	A:	66.6 $\pm$ 2.2	66.5 $\pm$ 2.1	54.0 $\pm$ 8.2	66.6 $\pm$ 2.2	52.6 $\pm$ 9.5	66.5 $\pm$ 2.1	52.8 $\pm$ 8.9	52.0 $\pm$ 11.2
		F:	66.8 $\pm$ 5.6	65.6 $\pm$ 5.8	36.5 $\pm$ 20.5	66.8 $\pm$ 5.6	30.7 $\pm$ 23.3	65.7 $\pm$ 5.8	35.4 $\pm$ 19.6	27.0 $\pm$ 24.7
$N = 7, N - K = 3$	drugs	A:	65.5 $\pm$ 2.4	65.7 $\pm$ 2.1	49.7 $\pm$ 10.8	65.9 $\pm$ 2.5	49.8 $\pm$ 8.5	65.9 $\pm$ 2.4	49.4 $\pm$ 10.4	51.2 $\pm$ 8.2
		F:	63.3 $\pm$ 5.8	63.9 $\pm$ 4.6	27.0 $\pm$ 27.8	63.9 $\pm$ 4.1	14.2 $\pm$ 13.6	64.5 $\pm$ 5.3	28.3 $\pm$ 29.6	21.0 $\pm$ 18.9
$N = 9, N - K = 4$	drugs	A:	63.2 $\pm$ 2.9	65.1 $\pm$ 1.9	46.8 $\pm$ 6.8	63.2 $\pm$ 2.7	47.4 $\pm$ 8.9	65.0 $\pm$ 2.0	45.5 $\pm$ 8.2	49.9 $\pm$ 9.9
		F:	67.5 $\pm$ 5.3	67.9 $\pm$ 5.3	20.2 $\pm$ 23.7	65.9 $\pm$ 5.4	17.4 $\pm$ 15.6	67.4 $\pm$ 4.6	15.2 $\pm$ 20.3	22.0 $\pm$ 21.3
$N = 11, N - K = 5$	drugs	A:	60.3 $\pm$ 9.5	62.6 $\pm$ 4.1	47.5 $\pm$ 5.9	56.5 $\pm$ 10.6	42.5 $\pm$ 7.8	65.2 $\pm$ 2.4	45.6 $\pm$ 5.1	49.4 $\pm$ 10.0
		F:	62.0 $\pm$ 6.1	50.1 $\pm$ 15.8	19.5 $\pm$ 14.4	61.9 $\pm$ 6.2	5.7 $\pm$ 7.4	62.6 $\pm$ 5.9	13.3 $\pm$ 14.8	20.0 $\pm$ 20.4
$N = 3, N - K = 1$	germancredit	A:	72.5 $\pm$ 4.3	73.0 $\pm$ 3.4	51.6 $\pm$ 13.2	72.5 $\pm$ 4.3	54.0 $\pm$ 12.7	73.0 $\pm$ 3.3	52.0 $\pm$ 13.2	63.4 $\pm$ 14.1
		F:	89.0 $\pm$ 5.3	88.9 $\pm$ 5.4	31.3 $\pm$ 28.2	89.0 $\pm$ 5.3	41.0 $\pm$ 33.0	88.9 $\pm$ 5.4	32.1 $\pm$ 28.5	67.9 $\pm$ 34.6
$N = 5, N - K = 2$	germancredit	A:	72.0 $\pm$ 3.9	72.4 $\pm$ 3.5	57.9 $\pm$ 13.0	72.3 $\pm$ 3.4	63.2 $\pm$ 10.9	72.9 $\pm$ 3.4	50.4 $\pm$ 15.1	56.2 $\pm$ 14.9
		F:	85.1 $\pm$ 7.0	84.5 $\pm$ 6.5	45.1 $\pm$ 26.4	84.8 $\pm$ 6.9	52.7 $\pm$ 22.8	85.4 $\pm$ 7.4	26.0 $\pm$ 20.6	35.8 $\pm$ 30.3
$N = 7, N - K = 3$	germancredit	A:	68.7 $\pm$ 8.0	72.6 $\pm$ 3.5	52.9 $\pm$ 14.4	71.4 $\pm$ 4.1	48.7 $\pm$ 11.2	72.2 $\pm$ 3.7	54.4 $\pm$ 11.4	54.4 $\pm$ 11.9
		F:	79.4 $\pm$ 12.3	87.2 $\pm$ 6.9	35.5 $\pm$ 27.9	84.0 $\pm$ 11.7	25.3 $\pm$ 19.2	85.7 $\pm$ 7.9	33.0 $\pm$ 20.8	28.8 $\pm$ 21.4
$N = 9, N - K = 4$	germancredit	A:	69.9 $\pm$ 3.2	71.8 $\pm$ 3.6	54.1 $\pm$ 10.0	66.2 $\pm$ 5.0	46.1 $\pm$ 13.4	71.2 $\pm$ 3.0	46.5 $\pm$ 11.6	52.5 $\pm$ 14.8
		F:	80.6 $\pm$ 14.3	89.6 $\pm$ 8.7	32.3 $\pm$ 19.7	70.3 $\pm$ 18.8	25.4 $\pm$ 28.5	84.5 $\pm$ 7.9	17.3 $\pm$ 14.7	29.5 $\pm$ 24.8
$N = 11, N - K = 5$	germancredit	A:	64.8 $\pm$ 7.9	71.2 $\pm$ 3.8	56.8 $\pm$ 11.3	64.0 $\pm$ 7.0	53.5 $\pm$ 13.1	71.5 $\pm$ 3.4	54.0 $\pm$ 9.6	50.5 $\pm$ 15.1
		F:	63.1 $\pm$ 20.2	83.3 $\pm$ 5.9	42.9 $\pm$ 21.0	62.6 $\pm$ 20.0	32.6 $\pm$ 24.1	83.6 $\pm$ 6.1	29.7 $\pm$ 18.2	28.2 $\pm$ 25.4
$N = 51, N - K = 5$	folktables	A:	78.8 $\pm$ 0.3	78.8 $\pm$ 0.2	78.1 $\pm$ 0.7	78.8 $\pm$ 0.2	77.8 $\pm$ 0.7	78.8 $\pm$ 0.2	78.0 $\pm$ 0.9	77.8 $\pm$ 1.3
		F:	78.9 $\pm$ 0.6	78.8 $\pm$ 0.5	79.0 $\pm$ 0.5	79.0 $\pm$ 0.6	78.8 $\pm$ 0.6	78.5 $\pm$ 0.6	79.0 $\pm$ 0.6	73.5 $\pm$ 3.6
$N = 51, N - K = 10$	folktables	A:	78.8 $\pm$ 0.2	78.8 $\pm$ 0.3	75.0 $\pm$ 2.0	78.9 $\pm$ 0.2	75.8 $\pm$ 1.5	78.8 $\pm$ 0.2	75.0 $\pm$ 2.3	75.2 $\pm$ 2.2
		F:	78.8 $\pm$ 0.6	78.7 $\pm$ 0.6	77.5 $\pm$ 3.5	78.8 $\pm$ 0.7	78.6 $\pm$ 1.2	78.3 $\pm$ 0.7	77.2 $\pm$ 3.6	66.6 $\pm$ 5.2
$N = 51, N - K = 15$	folktables	A:	78.8 $\pm$ 0.2	78.8 $\pm$ 0.2	70.4 $\pm$ 4.6	78.8 $\pm$ 0.2	71.2 $\pm$ 2.6	78.8 $\pm$ 0.2	70.4 $\pm$ 3.5	67.4 $\pm$ 7.5
		F:	78.6 $\pm$ 0.6	78.5 $\pm$ 0.7	68.6 $\pm$ 13.1	78.5 $\pm$ 0.7	70.5 $\pm$ 6.9	78.4 $\pm$ 0.7	68.3 $\pm$ 10.0	46.4 $\pm$ 19.9
$N = 51, N - K = 20$	folktables	A:	78.8 $\pm$ 0.2	78.7 $\pm$ 0.2	64.0 $\pm$ 5.8	78.8 $\pm$ 0.2	65.7 $\pm$ 5.4	78.7 $\pm$ 0.2	63.2 $\pm$ 7.2	60.4 $\pm$ 10.9
		F:	78.6 $\pm$ 0.6	78.6 $\pm$ 0.6	50.7 $\pm$ 16.7	78.6 $\pm$ 0.6	53.7 $\pm$ 13.7	78.6 $\pm$ 0.6	49.2 $\pm$ 19.3	33.3 $\pm$ 15.8
$N = 51, N - K = 25$	folktables	A:	58.0 $\pm$ 13.2	78.3 $\pm$ 0.2	53.2 $\pm$ 12.5	57.9 $\pm$ 15.2	42.4 $\pm$ 13.2	78.3 $\pm$ 0.3	51.1 $\pm$ 8.0	53.1 $\pm$ 11.8
		F:	39.6 $\pm$ 25.6	78.0 $\pm$ 1.7	27.5 $\pm$ 15.7	40.7 $\pm$ 27.6	18.9 $\pm$ 12.1	78.3 $\pm$ 1.0	25.7 $\pm$ 11.3	16.1 $\pm$ 12.3

## D Complete formulation and proof of Theorem 1

In this section we present the full proof of Theorem 1. We begin by reminding the reader of our notation and formal assumption from in D.1. Next in Section D.2 we state a few standard concentration results that are used in our main proof. In Section D.3 we define the population counterparts of the empirical discrepancy, disparity and disbalance measures, as understanding how well these measures are estimated from finite data is crucial for the proof of our results.

Finally, we present the full proof of Theorem 1 in Section D and the proof of the concentration lemmas from Section D.2 in Section D.5.

### D.1 Assumptions and formal adversary model

For convenience of the reader, we repeat the formal notation and assumptions stated in 3.3.1. Initially, there are  $N$  datasets  $\tilde{S}_1, \dots, \tilde{S}_N$ , with the  $i$ -th set of samples being drawn i.i.d. from a distribution  $p_i(x, y, a)$ . We assume that all these distributions are clean, in the sense that they are close to the true target distribution  $p$ . Formally, we assume that each of the following conditions hold:

$$\text{TV}(p_i(x, y, a), p(x, y, a)) \leq \eta, \quad \text{and} \quad \max_{z \in \mathcal{A}} \{ \text{TV}(p_i(x, y|a=z), p(x, y|a=z)) \} \leq \eta, \quad (14)$$

where  $\mathcal{B}(X)$  denotes the Borel  $\sigma$ -algebra on a topological space  $X$ .

Once the clean datasets  $\tilde{S}_1, \dots, \tilde{S}_N$  are sampled, an *adversary* operates on them. This results in new datasets,  $S_1, \dots, S_N$ , which the learning algorithm receives as input. The adversary is an arbitrary (deterministic or randomized) function  $\mathcal{F} : \prod_{i=1}^N (\mathcal{X} \times \mathcal{Y} \times \mathcal{A})^n \rightarrow \prod_{i=1}^N (\mathcal{X} \times \mathcal{Y} \times \mathcal{A})^n$ , with the only restriction that for a fixed subset of indices,  $G \subset \{1, \dots, N\}$ , the data remains unchanged. That is,  $S_i = \tilde{S}_i$  for all  $i \in G$ , and  $S_i$  is arbitrary for  $i \notin G$ .

Note that the learner only observes the datasets  $S_i$  and outputs a hypothesis based on them. Therefore, in the proof we will only work with the datasets  $S_i$  and not with  $\tilde{S}_i$ , using that  $S_i = \tilde{S}_i$  whenever  $i \in G$ , so that  $S_i$  is i.i.d. from  $p_i$ . For simplicity, we refer to a dataset  $S_i$  or a source  $i \in [N]$  as clean if  $i \in G$ .

We assume without loss of generality that  $\tau = \mathbb{P}_{(X,Y,A) \sim p}(A=0) \in (0, \frac{1}{2}]$ . For technical reasons, we also assume that  $9\eta < \tau = \mathbb{P}_{(X,Y,A) \sim p}(A=0)$ .

### D.2 Concentration tools and notation

We first present the two lemmas which demonstrate uniform convergence of the empirical risk and the empirical fairness deviation measure respectively, for any hypothesis set  $\mathcal{H}$  with finite VC dimension. The first is just the classic VC generalization bound, as given in Chapter 28.1 of Shalev-Shwartz & Ben-David (2014). The proof of the second lemma closely follows the proofs of similar results from Woodworth et al. (2017); Agarwal et al. (2018); Konstantinov & Lampert (2022) and is presented in Section D.5 for completeness.

**Lemma 1** (Uniform Convergence for Binary Loss). *Let  $d$  be the VC-dimension of  $\mathcal{H}$ . Then for any dataset  $S$  of size  $n$  sampled i.i.d. from a distribution  $p$ ,  $\forall \delta \in (0, 1)$ ,*

$$\mathbb{P} \left( \sup_{h \in \mathcal{H}} |\mathcal{R}_S(h) - \mathcal{R}_p(h)| > 2 \sqrt{\frac{8d \log \left( \frac{en}{d} \right) + 2 \log \left( \frac{4}{\delta} \right)}{n}} \right) \leq \delta.$$

**Lemma 2** (Uniform Convergence for demographic parity). *Let  $p$  be a distribution on  $\mathcal{X} \times \mathcal{A} \times \mathcal{Y}$ . Let  $d = \text{VC}(\mathcal{H}) \geq 1$  and let  $\tau = \min_{a \in \{0,1\}} \mathbb{P}_{(X,Y,A) \sim p}(A=a)$  for some constant  $\tau \in (0, 0.5]$ . Then for any dataset  $S$  of size  $n \geq \max \left\{ \frac{8 \log \left( \frac{8}{\delta} \right)}{\tau}, \frac{d}{2} \right\}$  sampled i.i.d. from  $p$ ,  $\forall \delta \in (0, 1/2)$ :*

$$\mathbb{P}_S \left( \sup_{h \in \mathcal{H}} |\Gamma_S(h) - \Gamma_p(h)| \geq 16 \sqrt{2 \frac{d \log \left( \frac{2en}{d} \right) + \log \left( \frac{24}{\delta} \right)}{n\tau}} \right) \leq \delta \quad (15)$$

For the dataset  $S_i$ , denote by

$$c_i := \sum_{(x,y,a) \in S_i} \mathbb{1}\{a = 0\} = |S_i^{a=0}|. \quad (16)$$

Denote  $\tau_i = \mathbb{P}_{(X,Y,A) \sim p_i}(A = 0)$ . Then for a clean data source we have that  $c_i \sim \text{Bin}(n, \tau_i)$ . Therefore, by the Hoeffding bound, for any  $\delta > 0$ :

$$\mathbb{P} \left( |c_i - n\tau_i| \geq n \sqrt{\frac{\log \left( \frac{2}{\delta} \right)}{2n}} \right) \leq 2 \exp \left( - \frac{2 \left( \sqrt{\frac{n}{2} \log \left( \frac{2}{\delta} \right)} \right)^2}{n} \right) = \delta. \quad (17)$$

We will assume without loss of generality that  $\tau = \mathbb{P}_{(X,Y,A) \sim p_i}(A = 0) = \min_{a \in \{0,1\}} \mathbb{P}_{(X,Y,A) \sim p}(A = a)$ . Note that since  $TV(p_i, p) \leq \eta$ , for any clean dataset  $S_i$ ,

$$\tau_i = \mathbb{P}_{(X,Y,A) \sim p_i}(A = 0) \geq \mathbb{P}_{(X,Y,A) \sim p}(A = 0) - \eta = \tau - \eta.$$

In addition,

$$1 - \tau_i = \mathbb{P}_{(X,Y,A) \sim p_i}(A = 1) \geq \mathbb{P}_{(X,Y,A) \sim p}(A = 1) - \eta \geq \mathbb{P}_{(X,Y,A) \sim p}(A = 0) - \eta = \tau - \eta.$$

Recall also that  $\tau - \eta \geq \tau - 9\eta > 0$  by assumption. Denote by:

$$\Delta(\delta) = \max \left\{ 2 \sqrt{\frac{8d \log \left( \frac{en}{d} \right) + 2 \log \left( \frac{4}{\delta} \right)}{n}}, 16 \sqrt{2 \frac{d \log \left( \frac{2en}{d} \right) + \log \left( \frac{24}{\delta} \right)}{n(\tau - \eta)}}, \sqrt{\frac{\log \left( \frac{2}{\delta} \right)}{2n}} \right\} \quad (18)$$

$$= 16 \sqrt{2 \frac{d \log \left( \frac{2en}{d} \right) + \log \left( \frac{24}{\delta} \right)}{n(\tau - \eta)}}. \quad (19)$$

The lemmas above, as well as the observation that  $\min\{\tau_i, 1 - \tau_i\} \geq \tau - \eta$  for any clean source  $i$ , readily imply that:

$$\mathbb{P} \left( \sup_{h \in \mathcal{H}} |\mathcal{R}_{S_i}(h) - \mathcal{R}_{p_i}(h)| \geq \Delta(\delta) \right) \leq \delta, \quad (20)$$

$$\mathbb{P}_S \left( \sup_{h \in \mathcal{H}} |\Gamma_{S_i}(h) - \Gamma_{p_i}(h)| \geq \Delta(\delta) \right) \leq \delta \quad (21)$$

and

$$\mathbb{P}_S (|c_i - n\tau_i| \geq n\Delta(\delta)) \leq \delta, \quad (22)$$

for any clean  $i$ .

### D.3 Discrepancy, disparity and disbalance between distributions

In our proof we will consider the population counterparts of the between-dataset distances that we defined in the main body of the text. In particular, the discrepancy distance between two distributions  $p$  and  $q$  is

$$\text{disc}(p, q) = \sup_{h \in \mathcal{H}} |\mathcal{R}_p(h) - \mathcal{R}_q(h)|. \quad (23)$$

Similarly, the disparity is

$$\text{disp}(p, q) = \sup_{h \in \mathcal{H}} |\Gamma_p(h) - \Gamma_q(h)|, \quad (24)$$

where as an unfairness measure  $\Gamma_p$  we will use:

$$\Gamma_p(h) = |\mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1 | A = 0) - \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1 | A = 1)|.$$

Finally, the disbalance is simply

$$\text{disb}(p, q) = |\mathbb{P}_p(A = 0) - \mathbb{P}_q(A = 0)|. \quad (25)$$

Next we study these distances, between a distribution  $p_i$  of a clean source and the true target distribution  $p$ . Recall our assumptions about the bounded TV distances from Section D.1. Clearly, we have that  $\text{disb}(p_i, p) \leq TV(p_i, p) \leq \eta$ . Note also that:

$$\text{disc}(p_i, p) = \sup_{h \in \mathcal{H}} |\mathcal{R}_{p_i}(h) - \mathcal{R}_p(h)| = \sup_{h \in \mathcal{H}} |\mathbb{P}_{(X,Y,A) \sim p_i}(h(X) \neq Y) - \mathbb{P}_{(X,Y,A) \sim p}(h(X) \neq Y)| \leq \eta,$$

because any (measurable) classifier  $h : \mathcal{X} \rightarrow \mathcal{Y}$  can be associated with a (Borel) set  $S_h = \{(x, y, a) \in (\mathcal{X} \times \mathcal{Y} \times \mathcal{A}) : h(x) \neq y\}$ . Finally, we bound the disparity in terms of  $\eta$ . Note that:

$$\begin{aligned} \text{disp}(p_i, p) &= \sup_{h \in \mathcal{H}} \left| \mathbb{P}_{(X,Y,A) \sim p_i}(h(X) = 1|A = 0) - \mathbb{P}_{(X,Y,A) \sim p_i}(h(X) = 1|A = 1) \right| \\ &\quad - \left| \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1|A = 0) - \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1|A = 1) \right| \\ &\leq \sup_{h \in \mathcal{H}} \left| \mathbb{P}_{(X,Y,A) \sim p_i}(h(X) = 1|A = 0) - \mathbb{P}_{(X,Y,A) \sim p_i}(h(X) = 1|A = 1) \right| \\ &\quad - \left| \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1|A = 0) - \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1|A = 1) \right| \\ &\leq \sup_{h \in \mathcal{H}} \left( \left| \mathbb{P}_{(X,Y,A) \sim p_i}(h(X) = 1|A = 0) - \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1|A = 0) \right| \right. \\ &\quad \left. + \left| \mathbb{P}_{(X,Y,A) \sim p_i}(h(X) = 1|A = 1) - \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1|A = 1) \right| \right) \\ &\leq \sup_{h \in \mathcal{H}} \left| \mathbb{P}_{(X,Y,A) \sim p_i}(h(X) = 1|A = 0) - \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1|A = 0) \right| \\ &\quad + \sup_{h \in \mathcal{H}} \left| \mathbb{P}_{(X,Y,A) \sim p_i}(h(X) = 1|A = 1) - \mathbb{P}_{(X,Y,A) \sim p}(h(X) = 1|A = 1) \right| \\ &\leq 2\eta. \end{aligned}$$

#### D.4 Proof

**Theorem 1.** Assume that  $\mathcal{H}$  has a finite VC-dimension  $d \geq 1$ . Let  $p$  be an arbitrary target data distribution and without loss of generality let  $\tau = p(a = 0) \in (0, 0.5]$ . Let  $S_1, \dots, S_N$  be  $N$  datasets, each consisting of  $n$  samples, out of which  $K > \frac{N}{2}$  are sampled i.i.d. from a data distribution  $p_i$  that is  $\eta$ -close the distribution  $p$  in the sense of Section D.1. Assume that  $9\eta < \tau$ . For  $\frac{1}{2} < \beta \leq \frac{K}{N}$  and  $I = \text{FILTERSOURCES}(S_1, \dots, S_N; \beta)$  set  $S = \bigcup_{i \in I} S_i$ . Let  $\delta > 0$ . Then there exists a constant  $C = C(\delta, \tau, d, N, \eta)$ , such that for any  $n \geq C$ , the following inequalities hold with probability at least  $1 - \delta$  uniformly over all  $f \in \mathcal{H}$  and against any adversary:

$$\Gamma_S(f) \leq \Gamma_p(f) + 10\eta + \tilde{\mathcal{O}}\left(\sqrt{\frac{1}{n}}\right), \quad \mathcal{R}_S(f) \leq \mathcal{R}_p(f) + 9\eta + \tilde{\mathcal{O}}\left(\sqrt{\frac{1}{n}}\right). \quad (26)$$

*Proof.* First, we characterize a set of values into which the empirical risks and empirical deviation measures of the clean data sources falls with probability at least  $1 - \delta$ . Then we show that because the clean datasets cluster in such a way, any individual dataset that is accepted by the FILTERSOURCES algorithm provides good empirical estimates of the true risk and the unfairness measure. Finally, we show that the same holds for the union of these sets,  $S$ , which implies the inequalities (26). For the risk, the last step is a straightforward consequence of the second. For the fairness, however, a careful derivation is needed that crucially uses the disbalance measure as well.

**Step 1** Let  $G \subset [N]$  be the set of indexes  $i$ , such that  $S_i$  was not modified by the adversary. By definition,  $|G| = K$ . Now consider the following events that, as we will show, describe the likely values of the studied quantities on the clean datasets.

In particular, for all  $i \in G$ , let  $\mathcal{E}_i^{\mathcal{R}}$  be the event that:

$$\sup_{h \in \mathcal{H}} |\mathcal{R}_{S_i}(h) - \mathcal{R}_{p_i}(h)| \leq \Delta \left( \frac{\delta}{6N} \right), \quad (27)$$



let  $\mathcal{E}_i^\Gamma$  be the event that

$$\sup_{h \in \mathcal{H}} |\Gamma_{S_i}(h) - \Gamma_{p_i}(h)| \leq \Delta \left( \frac{\delta}{6N} \right), \quad (28)$$

let  $\mathcal{E}_i^{bin}$  be the event that

$$|c_i - n\tau_i| \leq n\Delta \left( \frac{\delta}{6N} \right) \quad (29)$$

and finally, let  $\mathcal{E}_i^{count}$  be the event that

$$0 < c_i < n. \quad (30)$$

Denote by  $(\mathcal{E}_i^{\mathcal{R}})^c, (\mathcal{E}_i^\Gamma)^c$  and  $(\mathcal{E}_i^{bin})^c, (\mathcal{E}_i^{count})^c$  the respective complements of these events. Then, by equations (20, 21, 22), we have:

$$\mathbb{P}((\mathcal{E}_i^{\mathcal{R}})^c) \leq \frac{\delta}{6N}, \quad \mathbb{P}((\mathcal{E}_i^\Gamma)^c) \leq \frac{\delta}{6N}, \quad \mathbb{P}((\mathcal{E}_i^{bin})^c) \leq \frac{\delta}{6N}, \quad \forall i \in G.$$

To bound the probability of the complement of  $\mathcal{E}_i^{count}$ , note that for any  $i \in G$

$$1 - \tau_i = \mathbb{P}_{(X,Y,A) \sim p_i}(A = 1) \leq \mathbb{P}_{(X,Y,A) \sim p}(A = 1) + \eta = 1 - \tau + \eta$$

and that  $1 - \tau + \eta < 1$  because of the assumption that  $\eta < \tau$ . Similarly,

$$\tau_i = \mathbb{P}_{(X,Y,A) \sim p_i}(A = 0) \leq \mathbb{P}_{(X,Y,A) \sim p}(A = 0) + \eta = \tau + \eta \leq 1 - \tau + \eta.$$

Now, for any  $i \in G$ , whenever  $n \geq C_1(\delta, \tau, d, N) = \frac{\log(\frac{4N}{\delta})}{\log(\frac{1}{1-\tau+\eta})} \geq \max \left\{ \frac{\log(\frac{4N}{\delta})}{\log(\frac{1}{1-\tau_i})}, \frac{\log(\frac{4N}{\delta})}{\log(\frac{1}{\tau_i})} \right\}$ , we have that

$$\begin{aligned} \mathbb{P}((\mathcal{E}_i^{count})^c) &= (1 - \tau_i)^n + \tau_i^n \\ &\leq \exp \left( -n \log \left( \frac{1}{1 - \tau_i} \right) \right) + \exp \left( -n \log \left( \frac{1}{\tau_i} \right) \right) \\ &\leq \frac{\delta}{4N} + \frac{\delta}{4N} = \frac{\delta}{2N}. \end{aligned}$$

Therefore, setting  $\mathcal{E} := (\bigwedge_{i \in G} \mathcal{E}_i^{\mathcal{R}}) \wedge (\bigwedge_{i \in G} \mathcal{E}_i^\Gamma) \wedge (\bigwedge_{i \in G} \mathcal{E}_i^{bin}) \wedge (\bigwedge_{i \in G} \mathcal{E}_i^{count})$  then by the union bound the probability of  $\mathbb{P}(\mathcal{E}^c) \leq K \frac{\delta}{6N} + K \frac{\delta}{6N} + K \frac{\delta}{6N} + K \frac{\delta}{2N} \leq 3 \frac{\delta}{6} + \frac{\delta}{2} = \delta$ .

Hence the probability of the event  $\mathcal{E}$  that all of (27, 28, 29, 30) hold is at least  $1 - \delta$ .

**Step 2** Now we show that under the event  $\mathcal{E}$ , the inequalities in (26) are fulfilled. Indeed, assume that  $\mathcal{E}$  holds. Fix any adversary  $\mathcal{A}$  and any  $h \in \mathcal{H}$ .

For any pair of clean sources  $i, j \in [N]$  the triangle law and the derivations in Section D.3 give:

$$\begin{aligned} \text{disc}(S_i, S_j) &= \sup_{h \in \mathcal{H}} |\mathcal{R}_{S_i}(h) - \mathcal{R}_{S_j}(h)| \\ &\leq \sup_{h \in \mathcal{H}} |\mathcal{R}_{S_i}(h) - \mathcal{R}_{p_i}(h)| + \sup_{h \in \mathcal{H}} |\mathcal{R}_{p_i}(h) - \mathcal{R}_p(h)| + \sup_{h \in \mathcal{H}} |\mathcal{R}_p(h) - \mathcal{R}_{p_j}(h)| + \sup_{h \in \mathcal{H}} |\mathcal{R}_{p_j}(h) - \mathcal{R}_{S_j}(h)| \\ &\leq 2\eta + 2\Delta \left( \frac{\delta}{6N} \right). \end{aligned}$$

Similarly,

$$\text{disp}(S_i, S_j) = \sup_{h \in \mathcal{H}} |\Gamma_{S_i}(h) - \Gamma_{S_j}(h)| \leq 4\eta + 2\Delta \left( \frac{\delta}{6N} \right)$$

and

$$\text{disb}(S_i, S_j) = \left| \frac{c_i}{n} - \frac{c_j}{n} \right| \leq 2\eta + 2\Delta \left( \frac{\delta}{6N} \right).$$

Therefore, for any pair of clean sources  $i, j \in [N]$ :

$$\text{disc}(S_i, S_j) + \text{disp}(S_i, S_j) + \text{disb}(S_i, S_j) \leq 8\eta + 6\Delta \left( \frac{\delta}{6N} \right). \quad (31)$$

It follows that, under  $\mathcal{E}$ , we have that  $q_i \leq 8\eta + 6\Delta \left( \frac{\delta}{6N} \right)$  for any clean  $i \in [N]$ . Since the fraction of clean sources is  $\frac{K}{N} \geq \beta$ , it follows that also  $q \leq 8\eta + 6\Delta \left( \frac{\delta}{6N} \right)$ , where  $q$  is the  $\beta$ -th quantile of the  $q_i$ 's.

Denote by  $I = \text{FILTERSOURCES}(S_1, \dots, S_N; \beta)$  the result of the filtering algorithm. Now for any  $i \in I$ , we have that  $q_i \leq q \leq 8\eta + 6\Delta \left( \frac{\delta}{6N} \right)$ . In addition, by the definition of  $q_i$ ,  $\text{disc}(S_i, S_j) \leq \text{disc}(S_i, S_j) + \text{disp}(S_i, S_j) + \text{disb}(S_i, S_j) \leq q_i$  for at least  $|I| = \beta N > \frac{N}{2}$  values of  $j \in [N]$ . Since  $K > \frac{N}{2}$ , this means that  $\text{disc}(S_i, S_j) \leq q_i \leq 8\eta + 6\Delta \left( \frac{\delta}{6N} \right)$  for at least 1 value  $j \in G$ . Therefore, we have:

$$\sup_{h \in \mathcal{H}} |\mathcal{R}_{S_i}(h) - \mathcal{R}_p(h)| \leq \sup_{h \in \mathcal{H}} |\mathcal{R}_{S_i}(h) - \mathcal{R}_{S_j}(h)| + \sup_{h \in \mathcal{H}} |\mathcal{R}_{S_j}(h) - \mathcal{R}_{p_j}(h)| + \sup_{h \in \mathcal{H}} |\mathcal{R}_{p_j}(h) - \mathcal{R}_p(h)| \quad (32)$$

$$\leq 8\eta + 6\Delta \left( \frac{\delta}{6N} \right) + \Delta \left( \frac{\delta}{6N} \right) + \eta \quad (33)$$

$$= 9\eta + 7\Delta \left( \frac{\delta}{6N} \right) \quad (34)$$

because  $\mathcal{E}$  holds. Similarly,

$$\sup_{h \in \mathcal{H}} |\Gamma_{S_i}(h) - \Gamma_p(h)| \leq 10\eta + 7\Delta \left( \frac{\delta}{6N} \right) \quad (35)$$

and

$$|c_i - n\tau| \leq 9\eta n + 7n\Delta \left( \frac{\delta}{6N} \right). \quad (36)$$

**Step 3** Finally, we study the risk and disparity measures based on all filtered data  $S = \cup_{i \in I} S_i$ .

Denote by  $\mathcal{R}_S(h)$  the empirical risk across the entire trusted dataset  $I$ :

$$\mathcal{R}_S(h) := \frac{1}{|I|} \sum_{i \in I} \mathcal{R}_{S_i}(h). \quad (37)$$

Then the triangle law gives:

$$|\mathcal{R}_S(h) - \mathcal{R}_p(h)| = \left| \frac{1}{|I|} \left( \sum_{i \in I} \mathcal{R}_{S_i}(h) - \mathcal{R}_p(h) \right) \right| \leq \frac{1}{|I|} \sum_{i \in I} |\mathcal{R}_{S_i}(h) - \mathcal{R}_p(h)| = 9\eta + 7\Delta \left( \frac{\delta}{6N} \right)$$

Since

$$3\Delta \left( \frac{\delta}{6N} \right) = 112 \sqrt{2 \frac{d \log \left( \frac{2en}{d} \right) + \log \left( \frac{144N}{\delta} \right)}{(\tau - \eta)n}} = \tilde{O} \left( \sqrt{\frac{d}{(\tau - \eta)n}} \right), \quad (38)$$

the bound on the risk follows.

Denote by  $\Gamma_S(h)$  the empirical estimate of demographic parity across the entire trusted dataset  $I$ :

$$\Gamma_S(h) := \left| \frac{\sum_{j \in I} \sum_{i=1}^n \mathbb{1}\{h(x_i^{(j)}) = 1, a_i^{(j)} = 0\}}{\sum_{j \in I} \sum_{i=1}^n \mathbb{1}\{a_i^{(j)} = 0\}} - \frac{\sum_{j \in I} \sum_{i=1}^n \mathbb{1}\{h(x_i^{(j)}) = 1, a_i^{(j)} = 1\}}{\sum_{j \in I} \sum_{i=1}^n \mathbb{1}\{a_i^{(j)} = 1\}} \right|. \quad (39)$$

For convenience, denote  $v_j = v_j(h) = \sum_{i=1}^n \mathbb{1}\{h(x_i^{(j)}) = 1, a_i^{(j)} = 0\}$  and  $w_j = w_j(h) = \sum_{i=1}^n \mathbb{1}\{h(x_i^{(j)}) = 1, a_i^{(j)} = 1\}$ , so that:

$$\Gamma_S(h) = \left| \frac{\sum_{j \in I} v_j}{\sum_{j \in I} c_j} - \frac{\sum_{j \in I} w_j}{\sum_{j \in I} (n - c_j)} \right|.$$

Note that because of equation (36) we have that:

$$\frac{\sum_{j \in I} v_j}{\sum_{j \in I} c_j} - \frac{\sum_{j \in I} w_j}{\sum_{j \in I} (n - c_j)} \leq \frac{\sum_{j \in I} v_j}{|I|(n\tau_1 - 7n\Delta(\frac{\delta}{6N}))} - \frac{\sum_{j \in I} w_j}{|I|(n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N}))}, \quad (40)$$

where  $\tau_1 = \tau - 9\eta > 0$ , whenever:

$$\min\{1 - \tau_1, \tau_1\} = \min\{1 - \tau + 9\eta, \tau - 9\eta\} = \tau - 9\eta > 7\Delta\left(\frac{\delta}{6N}\right) = 112\sqrt{2\frac{d \log\left(\frac{2en}{d}\right) + 2 \log\left(\frac{144N}{\delta}\right)}{n(\tau - \eta)}} \quad (41)$$

Since  $\Delta(\delta) = \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{n}}\right)$ , this holds whenever  $n > C_2$  for some constant  $C_2 = C_2(\delta, \tau, d, N)$ .

Specifically, it suffices to take:

$$n \geq \frac{4(112^2) \log\left(\frac{144N}{\delta}\right)}{(\tau - \eta)(\tau - 9\eta)^2} \implies \frac{\log\left(\frac{144N}{\delta}\right)}{n(\tau - \eta)} \leq \frac{(\tau - 9\eta)^2}{4(112^2)}$$

and

$$\log\left(\frac{2en}{d}\right) \leq \frac{(\tau - \eta)(\tau - 9\eta)^2}{4(112^2)d}n \implies \frac{d \log\left(\frac{2en}{d}\right)}{n(\tau - \eta)} \leq \frac{(\tau - 9\eta)^2}{4(112^2)}.$$

The second condition is fulfilled whenever

$$n \geq -W\left(-\frac{(\tau - \eta)(\tau - 9\eta)^2}{4(112^2)d}\right) / \left(\frac{(\tau - \eta)(\tau - 9\eta)^2}{4(112^2)d}\right) = -W\left(-\frac{(\tau - \eta)(\tau - 9\eta)^2}{8(112^2)e}\right) \frac{4(112^2)d}{(\tau - \eta)(\tau - 9\eta)^2},$$

where  $W(x)$  is the Lambert function. In summary, (41) and hence also (40) is fulfilled whenever:

$$n \geq C_2(\delta, \tau, d, N, \eta) := \max\left\{\frac{50176 \log\left(\frac{144N}{\delta}\right)}{(\tau - \eta)(\tau - 9\eta)^2}, -W\left(-\frac{(\tau - \eta)(\tau - 9\eta)^2}{100352e}\right) \frac{50176d}{(\tau - \eta)(\tau - 9\eta)^2}\right\}. \quad (42)$$

Now, under this assumption, note that:

$$\begin{aligned} \frac{\sum_{j \in I} v_j}{\sum_{j \in I} c_j} - \frac{\sum_{j \in I} w_j}{\sum_{j \in I} (n - c_j)} &\leq \frac{\sum_{j \in I} v_j}{|I|(n\tau_1 - 7n\Delta(\frac{\delta}{6N}))} - \frac{\sum_{j \in I} w_j}{|I|(n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N}))} \\ &= \frac{1}{|I|} \frac{(n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N})) \sum_{j \in I} v_j - (n\tau_1 - 7n\Delta(\frac{\delta}{6N})) \sum_{j \in I} w_j}{(n\tau_1 - 7n\Delta(\frac{\delta}{6N})) (n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N}))} \\ &= \frac{1}{|I|} \sum_{j \in I} \frac{(n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N})) v_j - (n\tau_1 - 7n\Delta(\frac{\delta}{6N})) w_j}{(n\tau_1 - 7n\Delta(\frac{\delta}{6N})) (n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N}))} \\ &= \frac{1}{|I|} \sum_{j \in I} \left( \frac{v_j}{n\tau_1 - 7n\Delta(\frac{\delta}{6N})} - \frac{w_j}{n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N})} \right) \\ &= \frac{1}{|I|} \sum_{j \in I} \left( \frac{v_j}{c_j} - \frac{w_j}{n - c_j} \right) \\ &\quad + \frac{1}{|I|} \sum_{j \in I} \left( \frac{v_j}{c_j} \left( \frac{c_j}{n\tau_1 - 7n\Delta(\frac{\delta}{6N})} - 1 \right) - \frac{w_j}{n - c_j} \left( \frac{n - c_j}{n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N})} - 1 \right) \right) \\ &\leq \frac{1}{|I|} \sum_{j \in I} \left( \frac{v_j}{c_j} - \frac{w_j}{n - c_j} \right) \\ &\quad + \frac{1}{|I|} \sum_{j \in I} \left( \frac{v_j}{c_j} \left( \frac{n\tau_1 + 7n\Delta(\frac{\delta}{6N})}{n\tau_1 - 7n\Delta(\frac{\delta}{6N})} - 1 \right) + \frac{w_j}{n - c_j} \left( 1 - \frac{n(1 - \tau_1) - 7n\Delta(\frac{\delta}{6N})}{n(1 - \tau_1) + 7n\Delta(\frac{\delta}{6N})} \right) \right) \end{aligned}$$

$$\begin{aligned}
&\leq \frac{1}{|I|} \sum_{j \in I} \left( \frac{v_j}{c_j} - \frac{w_j}{n - c_j} \right) \\
&\quad + \frac{1}{|I|} \sum_{j \in I} \left( \left( \frac{\tau_1 + 7\Delta(\frac{\delta}{6N})}{\tau_1 - 7\Delta(\frac{\delta}{6N})} - 1 \right) + \left( 1 - \frac{1 - \tau_1 - 7\Delta(\frac{\delta}{6N})}{1 - \tau_1 + 7\Delta(\frac{\delta}{6N})} \right) \right) \\
&\leq \frac{1}{|I|} \sum_{j \in I} \left( \frac{v_j}{c_j} - \frac{w_j}{n - c_j} \right) + \left( \frac{14\Delta(\frac{\delta}{6N})}{\tau_1 - 7\Delta(\frac{\delta}{6N})} + \frac{14\Delta(\frac{\delta}{6N})}{1 - \tau_1 + 7\Delta(\frac{\delta}{6N})} \right)
\end{aligned}$$

Similarly,

$$\frac{\sum_{j \in I} w_j}{\sum_{j \in I} (n - c_j)} - \frac{\sum_{j \in I} v_j}{\sum_{j \in I} c_j} \leq \frac{1}{|I|} \sum_{j \in I} \left( \frac{w_j}{n - c_j} - \frac{v_j}{c_j} \right) + \left( \frac{14\Delta(\frac{\delta}{6N})}{\tau_1 + 7\Delta(\frac{\delta}{6N})} + \frac{14\Delta(\frac{\delta}{6N})}{1 - \tau_1 - 7\Delta(\frac{\delta}{6N})} \right)$$

Denote:

$$T := \frac{14\Delta(\frac{\delta}{6N})}{\tau_1 + 7\Delta(\frac{\delta}{6N})} + \frac{14\Delta(\frac{\delta}{6N})}{1 - \tau_1 - 7\Delta(\frac{\delta}{6N})} + \frac{14\Delta(\frac{\delta}{6N})}{\tau_1 - 7\Delta(\frac{\delta}{6N})} + \frac{14\Delta(\frac{\delta}{6N})}{1 - \tau_1 + 7\Delta(\frac{\delta}{6N})} \leq \frac{56\Delta(\frac{\delta}{6N})}{\tau_1 - 7\Delta(\frac{\delta}{6N})}. \quad (43)$$

Then we have:

$$\begin{aligned}
\Gamma_S(h) &= \left| \frac{\sum_{j \in I} v_j}{\sum_{j \in I} c_j} - \frac{\sum_{j \in I} w_j}{\sum_{j \in I} (n - c_j)} \right| \\
&= \max \left\{ \frac{\sum_{j \in I} w_j}{\sum_{j \in I} (n - c_j)} - \frac{\sum_{j \in I} v_j}{\sum_{j \in I} c_j}, \frac{\sum_{j \in I} v_j}{\sum_{j \in I} c_j} - \frac{\sum_{j \in I} w_j}{\sum_{j \in I} (n - c_j)} \right\} \\
&\leq \max \left\{ \frac{1}{|I|} \sum_{j \in I} \left( \frac{w_j}{n - c_j} - \frac{v_j}{c_j} \right), \frac{1}{|I|} \sum_{j \in I} \left( \frac{v_j}{c_j} - \frac{w_j}{n - c_j} \right) \right\} + T \\
&\leq \frac{1}{|I|} \sum_{j \in I} \max \left\{ \left( \frac{w_j}{n - c_j} - \frac{v_j}{c_j} \right), \left( \frac{v_j}{c_j} - \frac{w_j}{n - c_j} \right) \right\} + T \\
&= \frac{1}{|I|} \sum_{j \in I} \left| \frac{w_j}{n - c_j} - \frac{v_j}{c_j} \right| + T \\
&= \frac{1}{|I|} \sum_{j \in I} \Gamma_{S_i}(h) + T \\
&\leq \Gamma_p(h) + 10\eta + 7\Delta \left( \frac{\delta}{6N} \right) + T.
\end{aligned}$$

Finally note that:

$$T \leq \frac{56\Delta(\frac{\delta}{6N})}{\tau_1 - 7\Delta(\frac{\delta}{6N})} = \frac{8}{\frac{\tau_1}{7\Delta(\frac{\delta}{6N})} - 1} = \frac{8}{\frac{(\tau - 9\eta)\sqrt{\tau - \eta}\sqrt{n}}{112\sqrt{2d \log(\frac{2en}{d}) + \log(\frac{144}{\delta})}} - 1} \leq \tilde{\mathcal{O}} \left( \frac{1}{\sqrt{n}} \right).$$

Hence, the theorem follows whenever:

$$\begin{aligned}
n &\geq C(\delta, \tau, d, N, \eta) = \max \{C_1, C_2\} \\
&= \max \left\{ \frac{\log \left( \frac{4N}{\delta} \right)}{\log \left( \frac{1}{1 - \tau} \right)}, \frac{50176 \log \left( \frac{144N}{\delta} \right)}{(\tau - \eta)(\tau - 9\eta)^2}, -W \left( -\frac{(\tau - \eta)(\tau - 9\eta)^2}{100352e} \right) \frac{50176d}{(\tau - \eta)(\tau - 9\eta)^2} \right\},
\end{aligned}$$

where  $W(x)$  is the Lambert function.  $\square$

## D.5 Proof of Lemma 2

Let  $S = \{(x_i, y_i, a_i)\}_{i=1}^n$ . For  $a \in \{0, 1\}$ , denote:

$$\gamma_S^a(h) = \frac{\sum_{i=1}^n \mathbb{1}\{h(x_i) = 1, a_i = a\}}{\sum_{i=1}^n \mathbb{1}\{a_i = a\}} \quad (44)$$

and

$$\gamma_p^a(h) = \mathbb{P}(h(X) = 1 | A = a), \quad (45)$$

so that  $\Gamma_S(h) = |\gamma_S^0(h) - \gamma_S^1(h)|$  and  $\Gamma_p(h) = |\gamma_p^0(h) - \gamma_p^1(h)|$ .

First we use a technique of Woodworth et al. (2017); Agarwal et al. (2018) for proving concentration results about conditional probability estimates to bound the probability of a large deviation of  $\Gamma_S(h)$  from  $\Gamma_p(h)$ , for a fixed hypothesis  $h \in \mathcal{H}$ . Our result is similar to the one in Woodworth et al. (2017), but for demographic parity, instead of equal odds.

**Lemma 3.** *Let  $h \in \mathcal{H}$  be a fixed hypothesis and  $p \in \mathcal{P}(\mathcal{X} \times A \times \mathcal{Y})$  be a fixed distribution. Let  $\tau = \min_{a \in \{0,1\}} \mathbb{P}_{(X,Y,A) \sim p}(A = a) \in (0, 0.5]$ . Then for any dataset  $S$ , drawn i.i.d. from  $p$ , of size  $n$  and for any  $\delta \in (0, 1)$  and any  $t > 0$ :*

$$\mathbb{P}(|\Gamma_S(h) - \Gamma_p(h)| > 2t) \leq 6 \exp\left(-\frac{t^2 \tau n}{8}\right). \quad (46)$$

*Proof.* Denote by  $S_a = \{i \in [n] : a_i = a\}$  the set of indexes of the points in  $S$  for which the protected group is  $a$ . Let  $c_a := |S_a|$  and  $P_a = \mathbb{P}_{(X,Y,A) \sim p}(A = a)$ , so that  $\tau = \min_a P_a$ . For both  $a \in \{0, 1\}$ , we have:

$$\begin{aligned} \mathbb{P}(|\gamma_S^a - \gamma_p^a| > t) &= \sum_{S_a} \mathbb{P}(|\gamma_S^a - \gamma_p^a| > t | S_a) \mathbb{P}(S_a) \\ &\leq \mathbb{P}\left(c_a \leq \frac{1}{2} P_a n\right) + \sum_{S_a : c_a > \frac{1}{2} P_a n} \mathbb{P}(|\gamma_S^a - \gamma_p^a| > t | S_a) \mathbb{P}(S_a) \\ &\leq \exp\left(-\frac{P_a n}{8}\right) + \sum_{S_a : c_a > \frac{1}{2} P_a n} 2 \exp(-2t^2 c_a) \mathbb{P}(S_a) \\ &\leq \exp\left(-\frac{P_a n}{8}\right) + 2 \exp(-t^2 P_a n) \\ &\leq 3 \exp\left(-\frac{t^2 \tau n}{8}\right). \end{aligned}$$

The triangle law gives:

$$||\gamma_S^0 - \gamma_S^1| - |\gamma_p^0 - \gamma_p^1|| \leq |\gamma_S^0 - \gamma_S^1 - \gamma_p^0 + \gamma_p^1| \leq |\gamma_S^0 - \gamma_p^0| + |\gamma_S^1 - \gamma_p^1|.$$

Combining the previous two results:

$$\begin{aligned} \mathbb{P}(|\gamma_S^0 - \gamma_S^1| - |\gamma_p^0 - \gamma_p^1| > 2t) &\leq \mathbb{P}(|\gamma_S^0 - \gamma_p^0| + |\gamma_S^1 - \gamma_p^1| > 2t) \\ &\leq \mathbb{P}((|\gamma_S^0 - \gamma_p^0| > t) \vee (|\gamma_S^1 - \gamma_p^1| > t)) \\ &\leq \mathbb{P}(|\gamma_S^0 - \gamma_p^0| > t) + \mathbb{P}(|\gamma_S^1 - \gamma_p^1| > t) \\ &\leq 6 \exp\left(-\frac{t^2 \tau n}{8}\right). \end{aligned}$$

□

Finally, we prove Lemma 2 by extending the previous result to hold uniformly over the whole hypothesis space, for any hypothesis space  $\mathcal{H}$  with a finite VC-dimension  $d := VC(\mathcal{H})$ . The extension is essentially identical to Konstantinov & Lampert (2022) and is included here for completeness.

**Lemma 2** (Uniform convergence for demographic parity). *Let  $d = VC(\mathcal{H}) \geq 1$  and let  $\tau = \min_{a \in \{0,1\}} \mathbb{P}_{(X,Y,A) \sim p}(A = a)$  for some constant  $\tau \in (0, 0.5]$ . Then for any dataset  $S$  of size  $n \geq \max \left\{ \frac{8 \log(\frac{8}{\delta})}{\tau}, \frac{d}{2} \right\}$  sampled i.i.d. from  $p$ ,  $\forall \delta \in (0, 1/2)$ :*

$$\mathbb{P}_S \left( \sup_{h \in \mathcal{H}} |\Gamma_S(h) - \Gamma_p(h)| \geq 16 \sqrt{2 \frac{d \log(\frac{2en}{d}) + \log(\frac{24}{\delta})}{n\tau}} \right) \leq \delta \quad (47)$$

*Proof.* To extend Lemma 3 to hold uniformly over  $\mathcal{H}$ , we first prove a version of the classic symmetrization lemma (Vapnik, 2013) for  $\Gamma$  and then proceed via a standard growth function argument.

1) Consider a ghost sample  $S' = \{(x'_i, a'_i, y'_i)\}_{i=1}^n$  also sampled i.i.d. from  $p$ . For any  $h \in \mathcal{H}$ , let  $\Gamma_{S'}(h)$  be the empirical estimate of  $\Gamma_p(h)$  based on  $S'$ .

We show the following symmetrization inequality for the  $\Gamma$  measure:

$$\mathbb{P}_S \left( \sup_{h \in \mathcal{H}} |\Gamma_S(h) - \Gamma_p(h)| \geq t \right) \leq 2 \mathbb{P}_{S,S'} \left( \sup_{h \in \mathcal{H}} |\Gamma_{S'}(h) - \Gamma_S(h)| \geq t/2 \right), \quad (48)$$

for any constant  $t \geq 8 \sqrt{\frac{2 \log(12)}{n\tau}}$ .

Indeed, let  $h^*$  be the hypothesis achieving the supremum on the left-hand side.<sup>9</sup> Then:

$$\mathbb{1}(|\Gamma_S(h^*) - \Gamma_p(h^*)| \geq t) \mathbb{1}(|\Gamma_{S'}(h^*) - \Gamma_p(h^*)| \leq t/2) \leq \mathbb{1}(|\Gamma_{S'}(h^*) - \Gamma_S(h^*)| \geq t/2).$$

Taking expectation with respect to  $S'$ :

$$\mathbb{1}(|\Gamma_S(h^*) - \Gamma_p(h^*)| \geq t) \mathbb{P}_{S'}(|\Gamma_{S'}(h^*, S') - \Gamma_p(h^*)| \leq t/2) \leq \mathbb{P}_{S'}(|\Gamma_{S'}(h^*) - \Gamma_S(h^*)| \geq t/2).$$

Now using Lemma 3:

$$\mathbb{P}_{S'}(|\Gamma_{S'}(h^*) - \Gamma_p(h^*)| \leq t/2) \geq 1 - 6 \exp \left( -\frac{t^2 \tau n}{128} \right) \geq 1 - \frac{1}{2} = \frac{1}{2},$$

where the second inequality follows from the condition  $t \geq 8 \sqrt{\frac{2 \log(12)}{n\tau}}$ . Therefore,

$$\frac{1}{2} \mathbb{1}(|\Gamma_S(h^*) - \Gamma_p(h^*)| \geq t) \leq \mathbb{P}_{S'}(|\Gamma_{S'}(h^*) - \Gamma_S(h^*)| \geq t/2).$$

Taking expectation with respect to  $S$ :

$$\begin{aligned} \mathbb{P}_S(|\Gamma_S(h^*) - \Gamma_p(h^*)| \geq t) &\leq 2 \mathbb{P}_{S,S'}(|\Gamma_{S'}(h^*) - \Gamma_S(h^*)| \geq t/2) \\ &\leq 2 \mathbb{P}_{S,S'} \left( \sup_{h \in \mathcal{H}} |\Gamma_{S'}(h) - \Gamma_S(h)| \geq t/2 \right). \end{aligned}$$

2) Next we use the symmetrization inequality (48) to bound the large deviation of  $\Gamma_S(h)$  uniformly over  $\mathcal{H}$ .

Specifically, given  $n$  points  $x_1, \dots, x_n \in \mathcal{X}$ , denote

$$\mathcal{H}_{x_1, \dots, x_n} \{ (h(x_1), \dots, h(x_n)) : h \in \mathcal{H} \}.$$

Then define the growth function of  $\mathcal{H}$  as:

$$G_{\mathcal{H}}(n) = \sup_{x_1, \dots, x_n} |\mathcal{H}_{x_1, \dots, x_n}|. \quad (49)$$

<sup>9</sup>If the supremum is not attained, the argument can be repeated for each element of a sequence of classifiers approaching the supremum

We will use that well-known Sauer's lemma (Vapnik, 2013), which states that whenever  $n \geq d$ ,  $G_{\mathcal{H}}(n) \leq \left(\frac{en}{d}\right)^d$ . Notice that given the two datasets  $S, S'$ , the values of  $\Gamma_S$  and  $\Gamma_{S'}$  depend only on the values of  $h$  on  $S$  and  $S'$  respectively. Therefore, for any  $t \geq 8\sqrt{\frac{2\log(12)}{\tau n}}$ ,

$$\mathbb{P}_S\left(\sup_{h \in \mathcal{H}} |\Gamma_S - \Gamma_p(h)| \geq t\right) \leq 2\mathbb{P}_{S,S'}\left(\sup_{h \in \mathcal{H}} |\Gamma_{S'}(h) - \Gamma_S(h)| \geq \frac{t}{2}\right) \quad (50)$$

$$\leq 2G_{\mathcal{H}}(2n)\mathbb{P}_{S,S'}\left(|\Gamma_{S'}(h) - \Gamma_S(h)| \geq \frac{t}{2}\right) \quad (51)$$

$$\leq 2G_{\mathcal{H}}(2n)\mathbb{P}_{S,S'}\left(\left(|\Gamma_S(h) - \Gamma_p(h)| \geq \frac{t}{4}\right) \vee \left(|\Gamma_{S'}(h) - \Gamma_p(h)| \geq \frac{t}{4}\right)\right) \quad (52)$$

$$\leq 4G_{\mathcal{H}}(2n)\mathbb{P}_S\left(|\Gamma_S(h) - \Gamma_p(h)| \geq \frac{t}{4}\right) \quad (53)$$

$$\leq 24G_{\mathcal{H}}(2n)\exp\left(-\frac{t^2\tau n}{516}\right) \quad (54)$$

$$\leq 24\left(\frac{2en}{d}\right)^d \exp\left(-\frac{t^2\tau n}{516}\right). \quad (55)$$

Here the second-to-last inequality is due to the same bound on the difference between  $\Gamma_S$  and  $\Gamma_p$  that was used in the previous lemma, and the last one follows from Sauer's lemma. Now if we use the threshold

$t = 16\sqrt{2\frac{d\log\left(\frac{2en}{d}\right) + \log\left(\frac{24}{\delta}\right)}{\tau n}} > 8\sqrt{\frac{2\log(12)}{\tau n}}$ , we get:

$$\mathbb{P}_S\left(\sup_{h \in \mathcal{H}} |\Gamma_S(h) - \Gamma_p(h)| \geq 16\sqrt{2\frac{d\log\left(\frac{2en}{d}\right) + \log\left(\frac{24}{\delta}\right)}{\tau n}}\right) < \delta. \quad (56)$$

□