

# Equal Opportunity in Online Classification with Partial Feedback

Yahav Bechavod\*   Katrina Ligett†   Aaron Roth‡   Bo Waggoner§   Zhiwei Steven Wu¶

February 5, 2019

## Abstract

We study an online classification problem with partial feedback in which individuals arrive one at a time from a fixed but unknown distribution, and must be classified as positive or negative. Our algorithm only observes the true label of an individual if they are given a positive classification. This setting captures many classification problems for which fairness is a concern: for example, in criminal recidivism prediction, recidivism is only observed if the inmate is released; in lending applications, loan repayment is only observed if the loan is granted. We require that our algorithms satisfy common statistical fairness constraints (such as equalizing false positive or negative rates — introduced as “equal opportunity” in [18]) *at every round*, with respect to the underlying distribution. We give upper and lower bounds characterizing the cost of this constraint in terms of the regret rate (and show that it is mild), and give an *oracle efficient* algorithm that achieves the upper bound.

## 1 Introduction

Many real-world prediction tasks in which fairness concerns arise — such as hiring, college admissions, lending, and recidivism prediction — are naturally modeled as online binary classification problems, but with an important twist: feedback is only received for one of the two classification outcomes. Worker performance is only observed for candidates who were actually hired; student success can only be tracked for students who were admitted; those who are denied a loan never have an opportunity to demonstrate that they would have repaid; only the detainees who were released from prison have an opportunity to commit new crimes. Applying standard techniques for enforcing statistical fairness constraints on the *gathered data* can thus lead to pernicious feedback loops that can lead to classifiers that badly violate these constraints on the underlying distribution. This kind of failure to “explore” has been highlighted as an important source of algorithmic unfairness — for example, in predictive policing settings [26, 13, 14].

---

\*School of Computer Science and Engineering, The Hebrew University. Email: yahav.bechavod@cs.huji.ac.il. Supported in part by Israeli Science Foundation (ISF) grant #1044/16, a subcontract on the DARPA Brandeis Project, and the Federmann Cyber Security Center in conjunction with the Israel national cyber directorate.

†School of Computer Science and Engineering, The Hebrew University. Email: katrina@cs.huji.ac.il. Supported in part by Israeli Science Foundation (ISF) grant #1044/16, a subcontract on the DARPA Brandeis Project, and the Federmann Cyber Security Center in conjunction with the Israel national cyber directorate. Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing.

‡Department of Computer and Information Sciences, University of Pennsylvania. Email: aaroth@cis.upenn.edu. Supported in part by NSF grants AF-1763307, CNS-1253345, a subcontract on the DARPA Brandeis Project, and an Amazon Research award.

§Microsoft Research New York. Email: bwag@colorado.edu.

¶Computer Science and Engineering Department, University of Minnesota. Email: zsw@umn.edu.

To avoid this problem, it is important to explicitly manage the exploration/exploitation tradeoff that characterizes learning in partial feedback settings, which is what we study in this paper. We ask for algorithms that enforce well-studied statistical fairness constraints across two protected populations (we focus on the “equal opportunity” constraint of [18], which enforces equalized false positive rates or false negative rates, but our techniques also apply to other statistical fairness constraints like “statistical parity” [12]). In particular, we ask for algorithms that satisfy these constraints (with respect to the unknown underlying distribution) at every round of the learning procedure. The result is that the fairness constraints restrict how our algorithms can *explore*, not just how they can *exploit*, which makes the problem of fairness-constrained online learning substantially different from in the batch setting. The main question that we explore in this paper is: “*how much does the constraint of fairness impact the regret bound of learning algorithms?*”

## 1.1 Our Model and Results

In our setting, there is an unknown distribution  $\mathcal{D}$  over *examples*, which are triples  $(\hat{x}, a, y) \in \mathcal{X} \times \{-1, 1\} \times \{-1, 1\}$ . Here  $\hat{x} \in \mathcal{X}$  represents a vector of features in some arbitrary *feature space*,  $a \in A = \{\pm 1\}$  is the *group* to which this example belongs (which we also call the *sensitive feature*), and  $y \in \mathcal{Y} = \{\pm 1\}$  is a binary label. We write  $x$  to denote a pair  $(\hat{x}, a)$  – the set of all features (including the sensitive one) that the learner has access to.

In each round  $t$ , our learner selects hypotheses from a *hypothesis class*  $\mathcal{H}$  consisting of functions  $h : \mathcal{X} \times A \rightarrow \mathcal{Y}$  recommending an action (or *label*) as a function of the features (potentially including the sensitive feature). We take the positive label to be the one that corresponds to observing feedback (hiring a worker, admitting a student, approving a loan, releasing an inmate, etc.) We allow algorithms which randomize over  $\mathcal{H}$ . Let  $\Delta(\mathcal{H})$  be the set of probability distributions over  $\mathcal{H}$ . We refer to a  $\pi \in \Delta(\mathcal{H})$  as a *convex combination of classifiers*.

**Definition 1.1 (False positive rate)** *For a fixed distribution  $\mathcal{D}$  on examples, we define the false positive rate (FPR) of a convex combination of classifiers  $\pi \in \Delta(\mathcal{H})$  on group  $j \in \{\pm 1\}$  to be*

$$FPR_j(\pi) = \mathbb{P}(\pi(x) = +1 | a = j, y = -1) = \mathbb{E}_{h \sim \pi} \left[ \mathbb{P}_{(x,y) \sim \mathcal{D}} (h(x) = +1 | a = j, y = -1) \right].$$

We denote the *difference* between false positive rates between populations as

$$\Delta_{FPR}(\pi) := FPR_1(\pi) - FPR_{-1}(\pi).$$

The fairness constraint we impose on our classifiers in this paper asks that false positive rates be approximately equalized across populations *at every round  $t$* . Throughout, analogous results hold for false negative rates. These constraints were called *equal opportunity* constraints in [18].

**Definition 1.2 ( $\gamma$ -equalized rates [18])** *Fix a distribution  $\mathcal{D}$ . A convex combination  $\pi \in \Delta(\mathcal{H})$  satisfies the  $\gamma$ -equalized false positive rate ( $\gamma$ -EFP) constraint if  $|\Delta_{FPR}(\pi)| \leq \gamma$ . We informally use the term  $\gamma$ -fair to refer to such a classifier or combination of classifiers.*

(As we will see in Definition 2.1, we will actually allow our algorithm a tiny probability of ever breaking the fairness constraint.)

Note that the fairness constraint is defined with respect to the true underlying distribution  $\mathcal{D}$ . One of the primary difficulties we face is that in early rounds, the learner has very little information about  $\mathcal{D}$ , and yet is required to satisfy the fairness constraint with respect to  $\mathcal{D}$ .

It is straightforward to see (and a consequence of a more general lower bound that we prove) that a  $\gamma$ -fair algorithm cannot in general achieve non-trivial regret to the set of  $\gamma$ -fair convex combinations of classifiers, because of ongoing statistical uncertainty about the fairness level for all non-trivial classifiers. Thus throughout this paper, our goal is to minimize our regret to the  $\gamma$ -fair convex combination of classifiers that has the lowest classification error on  $\mathcal{D}$ , while guaranteeing that our algorithm only deploys convex combinations of classifiers that guarantee fairness level  $\gamma'$  for some  $\gamma' > \gamma$ . Clearly, the optimal regret bound will necessarily be a function of the gap  $(\gamma' - \gamma)$ , and one of our aims is to characterize this tradeoff.

**An initial approach.** Even absent fairness constraints, the problem of learning from partial feedback is challenging and has been studied under the name “apple tasting”[19]. Via standard techniques, it can be reduced to a contextual bandit problem [4]. Therefore, an initial approach starts with the observation (Lemma B.1) that although the set of “fair distributions over classifiers” is continuously large, the “fair empirical risk minimization” problem only has a single constraint, and so we may without loss of generality consider distributions over hypotheses  $\mathcal{H}$  that have support of size 2. By an appropriate discretization, this allows us to restrict attention to a finite net of classifiers whenever  $\mathcal{H}$  itself is finite. From this observation, one could employ a simple strategy to obtain an information theoretic result: pair an “exploration” round in which all examples are classified as positive (so as to gather label information and rule out classifiers that substantially violate the fairness constraints on the gathered data), with an “exploitation” round in which a generic contextual bandits algorithm like a variant of EXP4 [3, 6] is run over the surviving finite (but exponentially large) number of empirically fair distributions from the net. This simple approach yields the following bound: For any parameter  $\alpha \in [0.25, 0.5]$ , there is an algorithm that obtains a regret bound of  $O(T^{2\alpha})$  to the best  $\gamma$  fair classifier while satisfying a  $\gamma'$ -fairness constraint at every round with a gap of  $(\gamma' - \gamma) = O(T^{-\alpha})$ .

**Our results.** We show that the tradeoff achieved by the inefficient algorithm is tight by proving a lower bound in Section 4. Given this result, the main focus will be on a computationally efficient upper bound. The main disadvantage of the simple bandits reduction above is that it is tremendously computationally expensive. In some sense, this is unavoidable, because we measure the regret of our learner with respect to 0/1 classification error, which is computationally hard to minimize, even for very simple classes  $\mathcal{H}$  (see, e.g., [22, 16, 11]). However, we can still hope to give an *oracle efficient algorithm* for our problem. This approach — which is common in the contextual bandits literature — assumes access to an “oracle” which can in polynomial time solve the empirical risk minimization problem over  $\mathcal{H}$  (absent fairness constraints), and is an attractive way to isolate the “hard part” of the problem that is often tractable in practice. Our main result, which we devote the body of the paper to, is to show that access to such an oracle is sufficient to give a polynomial time algorithm for the fairness-constrained learning problem, matching the simple information theoretically optimal bounds described above. To do this, we use two tools. At a high level, our strategy is to apply the oracle efficient stochastic contextual bandit algorithm from [2]. In order to do this, we need to supply it with an offline learning oracle for the set of classifiers that can with high probability be certified to satisfy our fairness constraints given the data so far. We construct an approximate oracle for this problem (given a learning oracle for  $\mathcal{H}$ ) using the oracle-efficient reduction for offline fair classification from [1]. We need to overcome a number of technical difficulties stemming from the fact that the fair oracle that we can construct is only an approximate empirical risk minimizer — whereas the oracle assumed in [2] is exact. Moreover, the algorithm from [2] assumes a finite hypothesis class, whereas we need to obtain no regret to a continuous family of distributions over hypotheses. The final result is an oracle-efficient algorithm trading off between regret and fairness, allowing for a regret bound of  $O(T^{2\alpha})$  to the best  $\gamma$  fair classifier while

satisfying  $\gamma'$ -fairness at every round, with a gap of  $\gamma' - \gamma = O(T^{-\alpha})$  for  $\alpha \in [0.25, 5]$ .

## 1.2 Additional Related Work

This paper builds upon two lines of work in the fair machine learning literature. The bulk of this literature studies batch classification problems under a variety of statistical fairness constraints that fix some statistic of interest and then ask that it be approximately equalized across protected sub-populations. Popular statistics include raw classification rates [8, 23, 15] (also known as *statistical parity* [12]), positive predictive value [24, 9], and false positive and false negative rates (also known as *equal opportunity*) [24, 9, 18]. See [5] for more examples. One of the main attractions of this family of constraints is that they can generally be enforced in the batch setting without the need to make any assumptions about the data distribution.

There is also a literature on fair online classification and regression in the contextual bandit setting [20, 21, 25]. These papers have studied the achievable regret when the learning algorithm must satisfy a fairness constraint *at every round*, as we require in this paper. However, previous work has demanded stringent *individual* fairness constraints that bind on particular pairs of individuals, rather than just the average behavior of the classifier over large groups (as statistical fairness constraints do). As a result, strong realizability assumptions had to be made in order to derive non-trivial regret bounds (and even in the realizable setting, simple concept classes like conjunctions were shown to necessitate an exponentially slower learning rate when paired with individual fairness constraints [20]). Our paper interpolates between these two literatures: we ask for *statistical* fairness constraints to be enforced at every round of a learning procedure, and show that in this case, even without any assumptions at all, the effect of the fairness constraint on the achievable regret bound is mild.

Recently, [7] considered the problem of enforcing statistical fairness in a full information online learning setting, but from a very different perspective. They showed that in the *adversarial* setting, it can be impossible to satisfy equalized false positive and false negative rate constraints averaged over history, even when the adversary is constrained so that each individual classifier in the hypothesis class individually satisfies the constraint. In contrast, they show that it is possible to do this for the equalized error rates constraint. Our setting is quite different: on the one hand, we require that our algorithm satisfy its fairness constraint *at every round*, not just on average over the history, and we work in a partial information setting. On the other hand, we assume that examples are drawn from a distribution, rather than being adversarially chosen.

## 2 Additional Preliminaries

Throughout the paper, we assume  $+1, -1 \in \mathcal{H}$ , where  $+1$  and  $-1$  are the two *constant* classifiers (that is,  $+1(x) = 1$  and  $-1(x) = -1$  for all  $x$ ). In some cases, we will additionally assume  $+a, -a \in \mathcal{H}$ , where  $+a$  and  $-a$  are the identity function (and its negation) on the sensitive feature (that is,  $+a(\hat{x}, a) = a$  and  $-a(\hat{x}, a) = -a$  for all  $\hat{x}, a$ ).

**The Online Setting:** The learner interacts with the environment as follows.

Online Learning in Our Partial Feedback Setting

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

Learner chooses a convex combination  $\pi_t \in \Delta(\mathcal{H})$ .

Environment draws  $(x_t, y_t) \sim \mathcal{D}$  independently; learner observes  $x_t$ .

Learner labels the point  $\hat{y}_t = h_t(x_t)$ , where  $h_t \sim \pi_t$ .

**if**  $\hat{y}_t = +1$  **then**

Learner observes  $y_t$ .

We measure a learner’s performance using 0-1 loss,  $\ell(\hat{y}_t, y_t) = \mathbb{1}[\hat{y}_t \neq y_t]$ . Given a class of distributions  $\mathcal{P}$  over  $H \subseteq \mathcal{H}$  and a sequence of  $T$  examples, the optimal convex combination of hypotheses from  $H$  in hindsight is

$$\pi^*(\mathcal{P}) = \operatorname{argmin}_{\pi \in \mathcal{P}} \sum_{t=1}^T \mathbb{E}_{h \sim \mathcal{P}} [\ell(h(x_t), y_t)].$$

A learner’s (pseudo)-regret with respect to  $\mathcal{P}$  is

$$\text{Regret} = \sum_{t=1}^T \mathbb{E}_{(x_t, y_t) \sim \mathcal{D}} [\ell(h(x_t), y_t)] - \sum_{t=1}^T \mathbb{E}_{(x_t, y_t) \sim \mathcal{D}} \mathbb{E}_{h \sim \pi^*(\mathcal{P})} [\ell(h(x_t), y_t)].$$

In particular, when  $\mathcal{P} = \{\pi \in \Delta(\mathcal{H}) : \pi \text{ satisfies } \gamma\text{-EFP}\}$ , we call this the learner’s  $\gamma$ -EFP regret.

Finally, we ask for online learning algorithms that satisfy the following notion of fairness:

**Definition 2.1 (A  $\gamma$ -EFP( $\delta$ ) online learning algorithm)** *An online learning algorithm is said to satisfy  $\gamma$ -EFP( $\delta$ ) fairness (for  $\delta \in [0, 1]$ ) if, with probability  $1 - \delta$  over the draw of  $\{(x_t, a_t, y_t)\}_{t=1}^T \sim \mathcal{D}^T$ , simultaneously for all  $t \in [T]$ :  $\pi_t$  satisfies  $\gamma$ -EFP.*

**Cost Sensitive Classification Algorithms:** In this paper, we aim to give oracle-efficient online learning algorithms — that is, algorithms that run in polynomial time per round, assuming access to an oracle which can solve the corresponding offline empirical risk minimization problem. Concretely, we assume oracles for solving *cost sensitive classification (CSC)* problems over  $\mathcal{H}$ , which are defined by a set of examples  $x_j$  and a set of weights  $c_j^{-1}, c_j^{+1} \in \mathbb{R}$  corresponding to the cost of a negative and positive classification respectively.

**Definition 2.2** *Given an instance of a CSC problem  $S = \{x_j, c_j^{-1}, c_j^{+1}\}_{j=1}^n$ , a CSC oracle  $\mathcal{O}$  for  $\mathcal{H}$  returns  $\mathcal{O}(S) \in \arg \min_{h \in \mathcal{H}} \sum_{j=1}^n c_j^{(h(x_j))}$ . From these oracles, we will construct  $\nu$ -approximate CSC oracles that may have restricted ranges  $\Pi \subseteq \Delta(\mathcal{H})$ . Such oracles return  $\mathcal{O}_\nu(S) = \pi \in \Pi$  such that  $\mathbb{E}_{h \sim \pi} [\sum_{j=1}^n c_j^{(h(x_j))}] \leq \arg \min_{\pi \in \Pi} \mathbb{E}_{h \sim \pi} [\sum_{j=1}^n c_j^{(h(x_j))}] + \nu$ .*

**From “Apple Tasting” to Contextual Bandits:** Online classification problems under the feedback model we study were first described as “Apple Tasting” problems [19]. The algorithm’s loss at each round accumulates according to the following loss matrix:

$$L = \begin{matrix} & y = +1 & y = -1 \\ \hat{y} = +1 & \begin{pmatrix} 0 & 1 \end{pmatrix} \\ \hat{y} = -1 & \begin{pmatrix} 1 & 0 \end{pmatrix} \end{matrix}$$

but feedback is only observed for positive classifications (when  $\hat{y} = +1$ ). This is a different feedback model than the more commonly studied *contextual bandits* setting. In that setting, feedback is instead observed for whatever action is taken: i.e. for either a positive or a negative classification. We briefly recall the contextual bandits setting below, for an arbitrary loss function:

Online Learning in the Contextual Bandits Setting

**for**  $t = 1, \dots, T$  **do**  
 Learner chooses a convex combination  $\pi_t \in \Delta(\mathcal{H})$ .  
 Environment draws  $(x_t, y_t) \sim \mathcal{D}$  independently, learner observes  $x_t$ .  
 Learner labels the point  $\hat{y}_t = h_t(x_t)$ , where  $h_t \sim \pi_t$ .  
 Learner observes loss  $\ell(\hat{y}_t, y_t) \in [0, 1]$ .

It is nevertheless straightforward to transform the apple tasting setting into the contextual bandits setting (similar observations have been previously made [4]).

**Proposition 2.3** *Given an algorithm for online learning in the contextual bandits setting with regret guarantee  $\text{Regret}(T)$  with probability  $1 - \delta$ , one can construct an algorithm for online learning in the apple tasting setting (our partial feedback model) guaranteeing regret  $2\text{Regret}(T)$  with probability  $1 - \delta$ .*

### 3 An Oracle-Efficient Algorithm

Our algorithm proceeds in two phases. First, during the first  $T_0$  rounds, the algorithm performs *pure exploration* and always predicts  $+1$  to collect labelled data. Because constant classifiers exactly equalize the false positive rates across populations, the exploration round satisfies our fairness constraint. The algorithm will then use the collected data set to form empirical false positive rate fairness constraints, which we use to define our construction of a fair CSC oracle, given a CSC oracle unconstrained by fairness. Then, in the remaining rounds, we will run an adaptive contextual bandit algorithm that minimizes cumulative regret, while satisfying the empirical fairness constraint at every round.

We make two mild assumptions to simplify our analysis and the statement of our final bounds:

**Assumption 3.1** *Negative examples from each of the two protected groups have constant probability mass:  $\Pr[a = 1, y = -1], \Pr[a = -1, y = -1] \in \Omega(1)$ .*

**Assumption 3.2** *The hypothesis class  $\mathcal{H}$  contains the two constant classifiers and the identity function and its negation on the protected attribute:  $\{+1, -1, +\mathbf{a}, -\mathbf{a}\} \subseteq \mathcal{H}$ .*

Our main theorem is as follows:

**Theorem 3.3** *For any  $\mathcal{H}$  satisfying Assumption 3.2 and data distribution satisfying Assumption 3.1, there exists an oracle-efficient algorithm that takes parameters  $\delta \in [0, \frac{1}{\sqrt{T}}]$  and  $\gamma \geq 0$  as input and satisfies  $(\gamma + O(\ln(|\mathcal{H}|T/\delta)/T^{1/4}))$ -EFP( $\delta$ ) fairness and has an expected regret at most  $\tilde{O}\left(\sqrt{T} \ln(|\mathcal{H}|/\delta)\right)$  with respect to the class of  $\gamma$ -EFP fair policies.*

**Remark 3.4** We state our theorem in what we believe is the most attractive parametric regime: when it can obtain a regret bound of  $O(\sqrt{T})$ . But it is straightforward, by modifying the length of the exploration round, to obtain a more general tradeoff—a regret bound of  $O(T^{2\alpha})$  with respect to the set of  $\gamma$ -EFP fair policies, while satisfying  $(\gamma + O(T^{-\alpha}))$ -EFP( $\delta$ ) fairness, for any  $\alpha \in [1/4, 1/2]$ . This tradeoff is tight, as we show in Section 4.

**Algorithm.** The outline of our algorithm is as follows.

1. Label the first  $T_0$  arrivals as  $\hat{y}_t = 1$ ; observe their true labels.
2. Based on this data, construct an efficient *FairCSC oracle*. The oracle will be given a cost-sensitive classification objective. It returns an approximately-optimal convex combination  $\pi$  of hypotheses subject to the linear constraint of  $(\gamma + T^{-1/4})$ -EFP on the empirical distribution of data. We show the algorithm can be implemented to always return a member of  $\Pi$ , defined to be the set of mixtures on  $\mathcal{H}$  with support size two whose empirical fairness on the exploration data is at most  $\gamma + \tilde{O}(T^{-1/4})$ .
3. Instantiate a bandit algorithm with policy class  $\Pi$ . The bandit algorithm, a modification of [2], is described in detail in the next sections. In order to select its hypotheses, the bandit algorithm makes calls to the FairCSC oracle we implemented above.
4. For the remaining rounds  $t > T_0$ , choose labels  $\hat{y}_t$  selected by the bandit algorithm and provide feedback to the bandit algorithm via the reduction given by Proposition 2.3.

**Analysis.** In the remainder of this section, we present our analysis in three main steps.

- First, we study the empirical fairness constraint given by the data collected during the exploration phase and give a reduction from a cost-sensitive classification problem subject to the fairness constraint to a standard cost-sensitive classification problem absent the constraint, based on [1]. We further provide a fair approximate CSC oracle that returns a sparse solution, a distribution over  $\mathcal{H}$  with support size of at most 2.
- Next, we present the algorithm run in the second phase: at each round  $t > T_0$ , the algorithm makes a prediction based on a randomized policy  $\pi_t \in \Delta(\mathcal{H})$ , which is a solution to a feasibility program given by [2]. We show how to rely on an approximate fair CSC oracle to solve this program efficiently. Consequently, we generalize the results of [2] to the setting in which the given oracle may only optimize the cost sensitive objective approximately. This may be of independent interest.
- Finally, we provide the regret analysis for the algorithm by bounding the deviation between the algorithm’s empirical regret and true expected regret. This in particular requires uniform convergence over the entire class of fair randomized policies, which we show by leveraging the sparsity of the fair distributions.

We now give the formal proof of Theorem 3.3, citing the results described above that will be proved in the next sections.

**Proof 3.5 (Proof of Theorem 3.3)** We set  $T_0 = \Theta(\sqrt{T \ln |\mathcal{H}| / \delta})$ . First, Lemma 3.6 shows that given our empirical EFP constraint, there exists an optimal policy of support size at most 2. Next, Lemma 3.7 shows that, with probability  $1 - \delta$  over arrivals  $1, \dots, T_0$ , all convex combinations  $\pi \in \Pi$  satisfy  $\hat{\gamma}$ -EFP for

$\hat{\gamma} = \gamma + \beta$ ,  $\beta := O\left(\sqrt{\ln(|\mathcal{H}|/\delta)}/T^{1/4}\right)$ . It also implies that the optimal  $\gamma$ -fair policy is in the class. Theorem 3.8 shows that, given a CSC oracle for  $\mathcal{H}$ , we can implement an efficient approximate CSC oracle for this class  $\Pi$ . Theorem 3.14 shows that, given an approximate CSC oracle for any class, there is an efficient bandit algorithm that plays from this class and achieves expected regret  $O\left(\ln(|\mathcal{H}|T/\delta)\sqrt{T}\right)$ .

**Fairness:** In the first  $T_0$  rounds we play  $+1$  which is 0-fair, and in the remaining rounds we play only policies from  $\Pi$ . With probability  $1 - \delta$  over the exploration data, every member of  $\Pi$  is  $(\gamma + \beta)$ -fair.

**Regret:** The algorithm's regret is at most  $T_0$  plus its regret, on rounds  $T_0 + 1, \dots, T$ , to the optimal policy in  $\Pi$ . By Proposition 2.3, this is at most twice the bandit algorithm's regret on those rounds. So our expected regret totals at most  $O\left(\ln(|\mathcal{H}|T/\delta)\sqrt{T}\right)$  to the best policy in  $\Pi$ . With probability  $1 - \delta$ ,  $\Pi$  contains the optimal  $\gamma$ -fair classifier; with the remaining probability, the algorithm's regret to the best  $\gamma$ -fair classifier can be bounded by  $T$ . Choosing  $\delta \leq \frac{1}{\sqrt{T}}$  gives the result.

### 3.1 Step 1: Constructing a Fair CSC Oracle From Exploration Data

Let  $S_E$  denote the set of  $T_0$  labeled examples  $\{z_i = (x_i, a_i, y_i)\}_{i=1}^{T_0}$  collected from the initial exploration phase, and let  $\mathcal{D}_E$  denote the empirical distribution over  $S_E$ . We will use  $\mathcal{D}_E$  as a proxy for the true distribution to form an *empirical fairness constraint*. To support the learning algorithm in the second phase, we need to construct an oracle that solves CSC problems subject to the empirical fairness constraint. Formally, an instance of the *FairCSC* problem for the class  $\mathcal{H}$  is given by a set of  $n$  tuples  $\{(x_j, c_j^{(-1)}, c_j^{(+1)})\}_{j=1}^n$  as before, along with a fairness parameter  $\gamma$  and an approximation parameter  $\nu$ . We would like solve the following fair CSC problem:

$$\min_{\pi \in \Delta(\mathcal{H})} \mathbb{E}_{h \sim \pi} \left[ \sum_{j=1}^n c_j^{(h(x_j))} \right] \quad \text{such that} \quad |\Delta_{FPR}(\pi, \mathcal{D}_E)| \leq \gamma \quad (1)$$

where  $\Delta_{FPR}(\pi, \mathcal{D}_E) = FPR_1(\pi, \mathcal{D}_E) - FPR_{-1}(\pi, \mathcal{D}_E)$  and each  $FPR_j(\pi, \mathcal{D}_E)$  denotes the false positive rate of  $\pi$  on distribution  $\mathcal{D}_E$ . We show a useful structural property that there always exists an optimal solution for any such FairCSC problem that has small support; the proof appears in Appendix B.1.

**Lemma 3.6** *There exists an optimal solution for the FairCSC that is a distribution over  $\mathcal{H}$  with support size no greater than 2.*

We therefore consider the set of sparse convex combinations:

$$\Pi = \{\pi \in \Delta(\mathcal{H}) \mid \text{Supp}(\pi) \leq 2, \quad |\Delta_{FPR}(\pi, \mathcal{D}_E)| \leq \gamma + \beta\}$$

and focus on algorithms that only play policies from  $\Pi$  and measure their performance with respect to  $\Pi$ . For any  $\pi \in \Pi$ , we will write  $\pi(h)$  to denote the probability  $\pi$  places on  $h$ . Applying a standard concentration inequality, we can show that each policy in  $\Pi$  is also approximately fair with respect to the underlying distribution.

**Lemma 3.7** *With probability  $1 - \delta$ , as long as  $T_0 \geq c\sqrt{T}\sqrt{\ln(|\mathcal{H}|/\delta)}$  for some universal constant  $c > 0$ , we have the following. First, every policy in  $\Pi$  satisfies  $\gamma + 2\beta$ -EFP, and second, every support-2  $\gamma$ -EFP policy is in  $\Pi$ , for  $\beta = O\left(\sqrt{\ln(|\mathcal{H}|/\delta)}/T^{1/4}\right)$ ,*



We provide a reduction from FairCSC problems to standard CSC problems as follows: 1) We first apply a standard transformation on the input CSC objective to derive an equivalent weighted classification problem, in which each example  $j$  has importance weight  $|c_j^{(-1)} - c_j^{(+1)}|$ . 2) We then run the fair classification algorithm due to [1] that solves the weighted classification problem *approximately* using a polynomial number of CSC oracle calls. 3) Finally, we follow an approach similar to that of [10] to shrink the support size of the solution returned by the fair classification algorithm down to at most 2, which can be done in polynomial time.

**Theorem 3.8 (Reduction from FairCSC to CSC)** *For any  $0 < \nu < \gamma/2$ , there exists an oracle-efficient algorithm that calls a CSC oracle for  $\mathcal{H}$  at most  $O(1/\nu^2)$  times and computes a solution  $\hat{\pi} \in \Delta(\mathcal{H})$  that has a support size of at most 2, satisfies  $\gamma$ -EFP, and has total cost*

$$\mathbb{E}_{h \sim \hat{\pi}} \left[ \sum_{j=1}^n c_j^{h(x_j, a_j)} \right] \leq \min_{\pi \in \Pi} \mathbb{E}_{h \sim \pi} \left[ \sum_{j=1}^n c_j^{h(x_j, a_j)} \right] + \epsilon$$

with  $\epsilon = 4\nu \sum_{j=1}^n |c_j^{(-1)} - c_j^{(+1)}|$ .

### 3.2 Step 2: The Adaptive Learning Phase

**Overview of bandit algorithm.** In the second phase, rounds  $t > T_0$ , we utilize a bandit algorithm to make predictions. We now describe the algorithm, which closely follows the ILOVETOCONBANDITS algorithm by [2] but with important modifications that are necessary to handle approximation error in the FairCSC oracle.

At each round  $t > T_0$ , the bandit algorithm produces a distribution  $Q_t$  over policies  $\pi$ . Each policy  $\pi$  is a convex combination of two classifiers in  $\mathcal{H}$  and satisfies approximate fairness. The algorithm then draws  $\pi$  from  $Q_t$ , draws  $h$  from  $\pi$ , and labels  $\hat{y}_t = h(x_t)$ . To choose  $Q_t$ , the algorithm places some constraints on  $Q$  and runs a short coordinate descent algorithm to find a  $Q$  satisfying those constraints. Finally, it mixes in a small amount of the uniform distribution over labels (which can be realized by mixing between  $+1$  and  $-1$ ). We will see that the constraints, called the feasibility program, correspond to roughly bounding the expected regret of the algorithm along with bounding the variance in regret of each possible  $\pi$ .

**Feasibility program.** To describe the feasibility program, we first introduce some notation. For each  $t$ , we will write  $p_t$  to denote the probability that prediction  $\hat{y}_t$  is selected by the learner, and for each policy  $\pi \in \Pi$ , let

$$\hat{L}_t(\pi) = \frac{1}{t} \sum_{s=1}^t \ell_s \frac{\Pr[\pi(x_s) = a_s]}{p_s}, \quad L(\pi) = \mathbb{E}_{(x, a, y) \sim \mathcal{D}} \left[ \mathbb{E}_{\pi} [\mathbf{1}[\pi(x) \neq y]] \right]$$

denote the estimated average loss given by the *inverse propensity score* (IPS) estimator and true expected loss for  $\pi$ , respectively. Similarly, let

$$\widehat{\text{Reg}}_t(\pi) = \hat{L}_t(\pi) - \min_{\pi' \in \Pi} \hat{L}_t(\pi'), \quad \text{Reg}(\pi) = L(\pi) - \min_{\pi' \in \Pi} L(\pi'),$$

denote the estimated average regret and the true expected regret. In order to bound the variance of the IPS estimators, we will ensure that the learner predicts each label with minimum probability  $\mu_t$  at each round  $t$ .

In particular, given a solution  $Q$  for the program and a minimum probability parameter  $\mu_t$ , the learner will predict according to the mixture distribution  $Q^{\mu_t}(\cdot | x)$ :

$$Q^{\mu_t}(\hat{y} | x) = \mu_t + (1 - 2\mu_t) \int_{\pi \in \Pi} Q(\pi) \Pr[\pi(x) = \hat{y}]$$

Note that this can be represented as a convex combination of classifiers from  $\mathcal{H}$  since we assume that  $\mathbf{1} \in \mathcal{H}$ . We define for each  $\pi \in \Pi$ ,  $b_t(\pi) = \frac{\widehat{\text{Reg}}_t(\pi)}{4(e-2)\mu_t \ln(T)}$ , and also initialize  $b_0(\pi) = 0$ .

We describe the feasibility problem solved at each step. The approach and analysis directly follow and extend that of [2]. In that work, the first step at each round is to compute the best policy so far, which lets us compute  $\widehat{\text{Reg}}_t(\pi)$  and  $b_t(\pi)$  for any policy  $\pi$ . Here, our FairCSC oracle only computes approximate solutions, and so we can only compute regret relative to the approximately best policy so far, which leads to corresponding approximations  $\widetilde{\text{Reg}}_t(\pi)$  and  $\tilde{b}_t(\pi)$ . Then, our algorithm solves the same feasibility program (although a few more technicalities must be handled): given history  $H_t$  (in the second phase) and minimum probability  $\mu_t$ , find a probability distribution  $Q$  over  $\Pi$  such that

$$\begin{aligned} \int_{\pi \in \Pi} Q(\pi) \tilde{b}_{t-1}(\pi) &\leq 4 && \text{(Low regret)} \\ \forall \pi \in \Pi : \mathbb{E}_{x \sim H_t} \left[ \frac{1}{Q^{\mu_t}(\pi(x) | x)} \right] &\leq 4 + \tilde{b}_{t-1}(\pi) && \text{(Low variance)} \end{aligned}$$

Intuitively, the first constraint ensures that the estimated regret (based on historical data) of the solution is at most  $\tilde{O}(1/\sqrt{t})$ . The second constraint bounds the variance of the resulting IPS loss estimator for policies in  $\Pi$ , which in turn allows us to bound the deviation between the empirical regret and the true regret for each policy over time. Importantly, we impose a tighter variance constraint on policies that have lower empirical regret so far, which prioritizes the regret estimation for the “good” policies thus far.

To solve the feasibility program using our FairCSC oracle, we will run a coordinate descent algorithm (full description in Section B.3). Over iterations, the algorithm maintains and updates a vector  $Q$  of non-negative weights that may sum to less than one; at the end, the remaining probability mass is placed on the empirically best policy  $\hat{\pi}^t$  (computed using a single call of FairCSC). At each iteration, the algorithm first checks whether the current  $Q$  violates the regret constraint; if so, the algorithm will shrink all the weights to meet the regret constraint. If the regret constraint is satisfied, the algorithm will then find the policy  $\pi$  such that its variance constraint is most violated, which can be identified using a single call of FairCSC oracle by the result of [2]. If the constraint violation is above 0, the algorithm increases the weight  $Q(\pi)$ . The algorithm halts when all of constraints are satisfied. Lastly, the distribution output by this computation is then mixed with a small amount of the uniform distribution  $\mu_t$  over labels.

Via a potential argument similar to the one of [2], we can show that the algorithm halts in a small number of iterations. We will also bound the additional error in the output solution due to the approximation in the FairCSC oracle. In the following, let  $\Lambda_0 = 0$  and for any  $t \geq 1$ ,

$$\Lambda_t := \frac{\nu}{4(e-2)\mu_t^2 \ln(T)}.$$

where  $\nu$  is the approximation parameter of the FairCSC oracle.

**Lemma 3.9** *Algorithm 1 halts in a number of iterations (and oracle calls) that is polynomial in  $\frac{1}{\mu_t}$ , and outputs a weight vector  $Q$ , then it is a probability distribution and the following hold:*

$$\int_{\pi \in \Pi} Q(\pi)(4 + b_{t-1}(\pi)) \leq 4 + \Lambda_t$$

$$\forall \pi \in \Pi : \mathbb{E}_{x \sim H_t} \left[ \frac{1}{Q^{\mu_t}(\pi(x) | x)} \right] \leq 4 + b_{t-1}(\pi) + \Lambda_t.$$

### 3.3 Step 3: Regret Analysis

The key step in our regret analysis is to establish a tight relationship between the estimated regret and the true expected regret and show that for any  $\pi \in \Pi$ ,  $\text{Reg}(\pi) \leq 2\widehat{\text{Reg}}(\pi) + \epsilon_t$ , with  $\epsilon_t = \tilde{O}(1/\sqrt{t})$ . The final regret guarantee then essentially follows from the guarantee of Lemma 3.9 that the estimated regret of our policy is bounded by  $\tilde{O}(1/t)$  with proper setting of  $\mu_t$ .

To bound the deviation between  $\text{Reg}(\pi)$  and  $\widehat{\text{Reg}}_t(\pi)$ , we need to bound the variance of our IPS estimators. Let us define the following for any probability distribution  $P$  over  $\Pi$ ,  $\pi \in \Pi$ ,

$$V(P, \pi, \mu) := \mathbb{E}_{x \sim \mathcal{D}} \left[ \frac{1}{P^\mu(\pi(x) | x)} \right] \quad \hat{V}_t(P, \pi, \mu) := \mathbb{E}_{x \sim H_t} \left[ \frac{1}{P^\mu(\pi(x) | x)} \right]$$

Recall that through the feasibility program, we can directly bound  $\hat{V}_t(Q_t, \pi, \mu_t)$  for each round. However, to apply a concentration inequality on the IPS estimator, we need to bound the population variance  $V(Q_t, \pi, \mu_t)$ . We do that through a deviation bound between  $\hat{V}_t(Q_t, \pi, \mu_t)$  and  $V(Q_t, \pi, \mu_t)$  for all  $\pi \in \Pi$ . In particular, we rely on the sparsity on  $\Pi$  and apply a covering argument. Let  $\Pi_\eta \subset \Pi$  denote an  $\eta$ -cover such that for every  $\pi$  in  $\Pi$ ,  $\min_{\pi' \in \Pi_\eta} \|\pi(h) - \pi'(h)\|_\infty \leq \eta$  for any  $h \in \mathcal{H}$ . Since  $\Pi$  consists of distributions with support size at most 2, we can take the cardinality of  $\Pi_\eta$  to be bounded by  $\lceil |\mathcal{H}|^2/\eta \rceil$ .

**Claim 3.10** *Let  $P$  be any distribution over the policy set  $\Pi$ , and let  $\pi$  be any policy in  $\Pi$ . Then there exists  $\pi' \in \Pi_\eta$  such that  $|V(P, \pi, \mu) - V(P, \pi', \mu)|, |\hat{V}_t(P, \pi, \mu) - \hat{V}_t(P, \pi', \mu)| \leq \frac{\eta}{\mu(\mu + \eta)}$ .*

**Lemma 3.11** *Suppose that  $\mu_t \geq \sqrt{\frac{\ln(2|\Pi_\eta|t^2/\delta)}{2t}}$ ,  $t \geq 8 \ln(2|\Pi_\eta|t^2/\delta)$ . Then with probability  $1 - \delta$ ,*

$$V(P, \pi, \mu_t) \leq 6.4\hat{V}_t(P, \pi, \mu_t) + 162.6 + \frac{2\eta}{\mu_t(\mu_t + \eta)}$$

Next we bound the deviation between the estimated loss and true expected loss for every all  $\pi \in \Pi$ .

**Lemma 3.12** *Assume that the algorithm solves the per-round feasibility program with accuracy guarantee of Lemma 3.9. With probability at least  $1 - \delta$ , we have for all  $t \in [T]$  all policies  $\pi \in \Pi$ ,  $\lambda \in [0, \mu_t]$ , and  $t \geq 8 \ln(2|\Pi_\eta|t^2/\delta)$ ,*

$$|L(\pi) - \hat{L}_t(\pi)| \leq (e - 2)\lambda \left( 188.2 + \frac{1}{t} \sum_{s=1}^t \left( 6.4b_{s-1}(\pi) + 6.4\Lambda_{s-1} + \frac{2\eta}{\mu_s(\mu_s + \eta)} \right) \right) + \frac{\ln\left(\frac{|\Pi_\eta|T}{\delta}\right)}{\lambda t}$$

To bound the difference between  $\text{Reg}(\pi)$  and  $\widehat{\text{Reg}}_t(\pi)$ , we will set  $\eta = 1/T^2$ ,  $\mu_t = \frac{3.2 \ln(|\Pi_\eta|T/\delta)}{\sqrt{t}}$  the approximation parameter  $\nu$  of FairCSC to be  $1/T$ .

**Lemma 3.13** *Assume that the algorithm solves the per-round feasibility program with the accuracy guarantee of Lemma 3.9. With probability at least  $1 - \delta$ , we have for all  $t \in [T]$  all policies  $\pi \in \Pi$ , and for all  $t \geq 8 \ln(2|\mathcal{H}|^2 T^3 / \delta)$ ,*

$$\text{Reg}(\pi) \leq 2\widehat{\text{Reg}}_t(\pi) + \epsilon_t, \quad \text{and} \quad \widehat{\text{Reg}}_t(\pi) \leq 2\text{Reg}(\pi) + \epsilon_t$$

with  $\epsilon_t = \frac{1000 \ln(|\mathcal{H}|^2 T^2 / \delta)}{\sqrt{t}}$ .

**Theorem 3.14** *The bandit algorithm, given access to an approximate-CSC oracle, runs in time polynomial in  $T$  and achieves expected regret at most  $O\left(\ln(|\mathcal{H}|T/\delta) \sqrt{T}\right)$ .*

**Proof 3.15** *The cumulative regret of the first  $T_1 = 8 \ln(2|\mathcal{H}|^2 T^3 / \delta)$  rounds is trivially bounded by  $O(\sqrt{T} \ln(|\mathcal{H}|T/\delta))$ . For each of the remaining rounds, we can use Lemma 3.13 to first bound the per-round regret of the sequence of  $Q_t$  as*

$$\int_{\pi \in \Pi} Q_t(\pi) \text{Reg}(\pi) \leq 2 \int_{\pi \in \Pi} Q_{t-1}(\pi) \widehat{\text{Reg}}(\pi) + \epsilon_{t-1}$$

*By the guarantee of Lemma 3.9, we can further bound the right hand side by  $(4(e-2)\mu_{t-1} \ln(T)) \Lambda_{t-1} \leq O(\ln(|\mathcal{H}|T/\delta)/\sqrt{t})$ . Summing over rounds, we see that the cumulative expected regret of the sequence of  $Q_t$ 's is bounded by  $O(\ln(|\mathcal{H}|T/\delta)\sqrt{T})$ . Finally, we need to take into account the  $\mu_t$  mixture of uniform prediction at each round, which incurs an additional cumulative regret of no more than  $O(\ln(|\mathcal{H}|T/\delta)\sqrt{T})$ .*

## 4 Lower Bound

In this section we show that the tradeoff that our algorithm exhibits between its regret bound and the ‘‘fairness gap’’  $\gamma' - \gamma$  (i.e. our algorithm is  $\gamma'$ -fair, but competes with the best  $\gamma$ -fair classifier when measuring regret) is optimal. We do this by constructing a lower bound instance consisting of two very similar distributions,  $\mathcal{D}_1$  and  $\mathcal{D}_2$  defined as a function of our algorithm’s fairness target  $\gamma$ . The instance is defined over a simple hypothesis class  $\mathcal{H}$ .  $\mathcal{H}$  contains the two constant classifiers ( $-1$  and  $+1$ ), and a pair of classifiers ( $h_1$  and  $h_2$ ) that each guarantee low error on both distributions, but only satisfy the 0-EFP constraint on one of the two.

Informally, we first prove that the two distributions cannot be distinguished for at least  $\Theta(\frac{1}{\gamma^2})$  rounds. We then argue that any algorithm satisfying our  $\gamma$ -EFP( $\delta$ ) constraint must play  $-1$  or  $+1$  with substantial probability over these initial rounds in order to guarantee that it does not violate its fairness guarantee on either  $\mathcal{D}_1$  or  $\mathcal{D}_2$ . But this implies incurring constant regret per round during this phase, which leads to our lower bound.

**Theorem 4.1** *Fix any  $\alpha \in (0, 0.5)$  and let  $T \geq \sqrt[3]{16}$ . Fix any  $\delta \leq 0.24$ . There exists a policy class  $\mathcal{H}$  containing  $\{\pm 1\}$  such that any algorithm satisfying a  $T^{-\alpha}$ -EFP( $\delta$ ) fairness constraint has expected regret with respect to the set of 0-EFP fair policies of  $\Omega(T^{2\alpha})$ .*

## Acknowledgements

We thank Nati Srebro for a conversation leading to the question we study here. We thank Michael Kearns for helpful discussions at an early stage of this work.

## References

- [1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna M. Wallach. A reductions approach to fair classification. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 60–69, 2018.
- [2] Alekh Agarwal, Daniel J. Hsu, Satyen Kale, John Langford, Lihong Li, and Robert E. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 1638–1646, 2014.
- [3] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- [4] Gábor Bartók, Dávid Pál, and Csaba Szepesvári. Toward a classification of finite partial-monitoring games. In *International Conference on Algorithmic Learning Theory*, pages 224–238. Springer, 2010.
- [5] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 0(0):0049124118782533, 2018.
- [6] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 19–26, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR.
- [7] Avrim Blum, Suriya Gunasekar, Thodoris Lykouris, and Nati Srebro. On preserving non-discrimination when combining expert advice. In *Advances in Neural Information Processing Systems*, pages 8386–8397, 2018.
- [8] Toon Calders and Sicco Verwer. Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292, 2010.
- [9] Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163, 2017.
- [10] Andrew Cotter, Heinrich Jiang, and Karthik Sridharan. Two-player games for efficient non-convex constrained optimization. *CoRR*, abs/1804.06500, 2018.
- [11] Amit Daniely, Nati Linial, and Shai Shalev-Shwartz. From average case complexity to improper learning complexity. *arXiv preprint arXiv:1311.2272*, 2013.
- [12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226. ACM, 2012.
- [13] Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. Runaway feedback loops in predictive policing. In *Conference on Fairness, Accountability and Transparency*, pages 160–171, 2018.

- [14] Danielle Ensign, Friedler Sorelle, Neville Scott, Scheidegger Carlos, and Venkatasubramanian Suresh. Decision making with limited feedback. In Firdaus Janoos, Mehryar Mohri, and Karthik Sridharan, editors, *Proceedings of Algorithmic Learning Theory*, volume 83 of *Proceedings of Machine Learning Research*, pages 359–367. PMLR, 07–09 Apr 2018.
- [15] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and removing disparate impact. In *KDD*, 2015.
- [16] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM Journal on Computing*, 41(6):1558–1590, 2012.
- [17] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, Jun 1981.
- [18] Moritz Hardt, Eric Price, Nati Srebro, et al. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, pages 3315–3323, 2016.
- [19] David P Helmbold, Nicholas Littlestone, and Philip M Long. Apple tasting. *Information and Computation*, 161(2):85–139, 2000.
- [20] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. In *Advances in Neural Information Processing Systems*, pages 325–333, 2016.
- [21] Matthew Joseph, Michael J. Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. Meritocratic fairness for infinite and contextual bandits. In Jason Furman, Gary E. Marchant, Huw Price, and Francesca Rossi, editors, *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, AIES 2018, New Orleans, LA, USA, February 02-03, 2018*, pages 158–163. ACM, 2018.
- [22] Adam Tauman Kalai, Adam R Klivans, Yishay Mansour, and Rocco A Servedio. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6):1777–1805, 2008.
- [23] Toshihiro Kamishima, Shotaro Akaho, and Jun Sakuma. Fairness-aware learning through regularization approach. In *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*, pages 643–650. IEEE, 2011.
- [24] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.
- [25] Yang Liu, Goran Radanovic, Christos Dimitrakakis, Debmalya Mandal, and David C Parkes. Calibrated fairness in bandits. *arXiv preprint arXiv:1707.01875*, 2017.
- [26] Kristian Lum and William Isaac. To predict and serve? *Significance*, 13(5):14–19, 2016.

## A Proof of Proposition 2.3

**Proof A.1 (Proof of Proposition 2.3)** Consider the following transformed loss matrix:

$$\tilde{L} = \begin{matrix} & y = +1 & y = -1 \\ \hat{y} = +1 & \begin{pmatrix} 0 & 2 \end{pmatrix} \\ \hat{y} = -1 & \begin{pmatrix} 1 & 1 \end{pmatrix} \end{matrix}$$

Given an online learning with partial feedback problem, we instantiate the bandit algorithm and always play the action it recommends. We then provide the algorithm with its feedback  $\tilde{L}_{\hat{y}^t, y^t}$ . This is possible because if  $\hat{y}^t = +1$  then we observe  $y^t$ , and if not then the feedback is 1 regardless of the unobserved value of  $y^t$ . For a sequence of arrivals  $S = \{(x^t, y^t)\}_{t=1}^T$ , let  $m(S)$  be the number of arrivals with  $y^t = -1$ . Let  $L(\pi, S) = \sum_{t=1}^T L_{\hat{y}^t, y^t}$  and similarly for  $\tilde{L}(\pi, S)$ . Then we have for all  $\pi, S$  that  $\tilde{L}(\pi, S) = L(\pi, S) + m(S)$ . In other words, on each round where  $y^t = 1$ , a prediction experiences the same loss under  $L$  and under  $\tilde{L}$ ; and on each round where  $y^t = -1$ , the loss is exactly one larger in the bandit setting. This difference does not depend on the prediction of the hypothesis, therefore every policy's total loss under the bandit loss is exactly  $m(S)$  larger than under the original loss.

It follows that our algorithm's regret is exactly equal to the bandit algorithm's. Finally, a bookkeeping note: in order that losses be bounded in  $[0, 1]$ , we must repeat the above argument using  $0.5\tilde{L}$  in place of  $\tilde{L}$ , which simply scales the bandit algorithm's regret by 0.5 relative to our algorithm's.

## B Missing Proofs for Section 3

### B.1 Proof of Lemma B.1

In this subsection, we establish a useful structural property for the general problem minimizing linear loss function subject to fairness constraints. This in turn provides a proof for Lemma B.1. In particular, given a hypothesis class  $\mathcal{H}$ , and a training set of labelled samples  $S$ , vectors  $a, b \in \mathbb{R}^{|\mathcal{H}|}$ , consider the problem:

$$\begin{aligned} \min_{x \in \Delta(|\mathcal{H}|)} \quad & a^\top x \\ \text{subject to} \quad & b^\top x \leq \gamma \\ & b^\top x \geq -\gamma \end{aligned}$$

Note that both the problem of weighted classification or cost-sensitive classification can be viewed as an instantiation of the linear program defined above. The sparsity in the solution will be useful in our analysis.

**Theorem B.1** *In the linear program above, there exists an optimal solution that is a distribution over  $\mathcal{H}$  with support size no greater than 2.*

**Proof B.2** Consider the following embedding of  $\mathcal{H}$  in  $\mathbb{R}^2$ :  $\forall h \in \mathcal{H} : \phi(h) = (a_h, b_h)$ . Let  $A = \{\phi(h) \mid \pi \in \mathcal{H}\}$ . Then the optimization problem can be written as the following problem over the convex hull  $\text{conv}(A)$ :

$$\begin{aligned} \text{minimize}_{(z_1, z_2) \in \text{conv}(A)} \quad & z_1 \\ \text{subject to} \quad & z_2 \leq \gamma \\ & z_2 \geq -\gamma \end{aligned}$$

Note there exists an optimal solution  $z^*$  that lies on an edge of the polytope defined by  $\text{conv}(A)$ . This means  $z^*$  is either a vertex of  $\text{conv}(A)$  or can be written as a convex combination of two vertices of  $\text{conv}(A)$ , say  $z'$  and  $z''$ . In the former case,  $z^*$  can be induced by a single hypothesis  $h^* \in \mathcal{H}$ , and in the latter case we know there exist  $h', h'' \in \mathcal{H}$  such that  $z' = \phi(h')$  and  $z'' = \phi(h'')$ . This means the optimal solution  $z^*$  can be induced by a convex combination of hypotheses.

Then the result of Lemma B.1 follows immediately.

## B.2 Proof of Theorem 3.8

Recall that we collect a set of  $T_0$  labeled examples  $\{z_i = (x_i, a_i, y_i)\}_{i=1}^{T_0}$  during the initial exploration phase, and let  $\mathcal{D}_E$  denote the corresponding empirical distribution. Recall that  $\mathcal{H}$  is a hypothesis class defined over both the features and the protected group memberships. We assume that  $\mathcal{H}$  contains a constant classifier (which implies that there is at least one fair classifier to be found, for any distribution). To simplify notation, we consider hypotheses that labels each example with either 0 or 1.

Suppose that we are given an instance of cost-sensitive classification instance  $(X_j, C_j^1, C_j^0)$ . We would like to compute a distribution over classifiers from  $\mathcal{H}$  that minimizes total cost subject to the false positive rate fairness constraint. In particular, consider the following *fair cost-sensitive classification (CSC)* problem:

$$\min_{\pi \in \Delta(\mathcal{H})} \mathbb{E}_{h \sim \pi} \left[ \sum_{j=1}^n (C_j^1 h(X_j) + C_j^0 (1 - h(X_j))) \right] \quad (2)$$

$$\text{such that } \forall j \in \{\pm 1\} \quad \text{FPR}_j(\pi) - \text{FPR}_{-j}(\pi) \leq \gamma. \quad (3)$$

$\text{FPR}_j(\pi) = \mathbb{E}_{h \sim \pi} [\text{FPR}_j(h)]$ . We will write  $\text{OPT}_C$  to denote the objective value at the optimum for the problem, that is the minimum cost achieved by a  $\gamma$ -EFP policy over distribution  $\mathcal{D}_E$ .

Equivalently, we can consider optimizing the following objective function:

$$\min_{\pi \in \Delta(\mathcal{H})} \mathbb{E}_{h \sim \pi} \left[ \sum_{j=1}^n W_j \mathbf{1}\{h(X_j) \neq Y_j\} \right]$$

where each  $W_j = |C_j^0 - C_j^1|$ ,  $Y_j = 1$  if  $C_j^0 > C_j^1$  and  $Y_j = 0$  otherwise. To reduce the problem further to the same formulation of [1], we consider objective with normalized weights

$$\min_{\pi \in \Delta(\mathcal{H})} \mathbb{E}_{h \sim \pi} \left[ \sum_{j=1}^n w_j \mathbf{1}\{h(X_j) \neq Y_j\} \right]$$

such that each  $w_j = W_j / (\sum_j W_j)$ . To simplify notation, we will write  $\text{err}(h, \mathcal{P}) = \sum_{j=1}^n w_j \mathbf{1}\{h(X_j) \neq Y_j\}$ , and  $\text{OPT}$  to denote optimal objective subject to  $\gamma$ -EFP.

For each of the fairness constraint in (3), we will introduce a dual variable  $\lambda_j \geq 0$ . This allows us to define the partial Lagrangian of the problem:

$$\mathcal{L}(\pi, \lambda) = \mathbb{E}_{h \sim \pi} [\text{err}(h, \mathcal{P})] + \sum_{j \in \{\pm 1\}} \lambda_j (\text{FPR}_j(\pi) - \text{FPR}_{-j}(\pi) - \gamma)$$

By strong duality, we have

$$\text{OPT} = \min_{g \in \Delta(\mathcal{H})} \max_{\lambda \in \mathbb{R}_+^2} \mathcal{L}(g, \lambda) = \max_{g \in \Delta(\mathcal{H})} \min_{\lambda \in \mathbb{R}_+^2} \mathcal{L}(g, \lambda).$$



where OPT is the optimal objective value of the ERM problem.

[1] provide an oracle-efficient algorithm for finding a  $\nu$ -approximate saddle point  $(\hat{g}, \hat{\lambda})$  of the Lagrangian:

$$\begin{aligned}\mathcal{L}(\hat{\pi}, \hat{\lambda}) &\leq \mathcal{L}(g, \hat{\lambda}) + \nu && \text{for all } g \in \Delta(\mathcal{H}) \\ \mathcal{L}(\hat{\pi}, \hat{\lambda}) &\geq \mathcal{L}(\hat{\pi}, \lambda) - \nu && \text{for all } \lambda \in \Lambda\end{aligned}$$

In their result, the algorithm restricts the dual space to be  $\Lambda = \{\|\lambda\|_1 \leq B \mid \lambda \in \mathbb{R}_+^2\}$  for some sufficiently large constant  $B$ . Their convergence rate and approximation parameter both depend on such  $C$ . We show that under the assumption that  $\mathcal{H}$  contains the two classifiers  $\mathbf{1}[a = j]$  for all  $j \in \{\pm 1\}$ , it is sufficient to set  $C = 2$ , and thus restrict the dual space to be

$$\Lambda = \{\|\lambda\|_1 \leq 2 \mid \lambda \in \mathbb{R}_+^2\}$$

Consequently, we can use their algorithm to find a  $\nu$ -approximate saddle point with only  $O\left(\frac{1}{\nu^2}\right)$  number of calls to the oracle  $\text{CSC}(\mathcal{H})$ .

**Lemma B.3 (Follows from Theorem 1 of [1])** *There is an oracle-efficient algorithm that computes a  $\nu$ -approximate saddle point for the restricted Lagrangian with  $\Lambda\{\|\lambda\|_1 \leq 2 \mid \lambda \in \mathbb{R}_+^2\}$ , using  $O(1/\nu^2)$  calls to a CSC oracle over  $\mathcal{H}$ .*

Moreover, the approximate saddle point provides an approximate solution for our problem.

**Lemma B.4** *Suppose that the class  $\mathcal{H}$  contains the two classifiers  $\mathbf{1}[a = j]$  for all  $j$  and that  $(\hat{\pi}, \hat{\lambda})$  is a  $\nu$ -approximate saddle point of the Lagrangian. Then the distribution  $\hat{\pi}$  satisfies*

$$\text{err}(\hat{\pi}, \mathcal{P}) \leq \text{OPT} + 2\nu, \quad \text{and} \quad \forall j \in \{\pm 1\} \quad \text{FPR}_j(\hat{g}) - \text{FPR}_{-j}(\hat{g}) \leq \gamma + 2\nu.$$

**Proof B.5** *Let  $\pi^*$  be the optimal feasible solution for the fair ERM problem. First, by the definition of approximate saddle point, we know that*

$$\begin{aligned}\text{err}(\hat{\pi}, \mathcal{P}) &= \mathcal{L}(\hat{\pi}, \mathbf{0}) \\ &\leq \max_{\lambda \in \Lambda} \mathcal{L}(\hat{\pi}, \lambda) \\ &\leq \mathcal{L}(\hat{\pi}, \hat{\lambda}) + \nu \\ &\leq \min_{\pi \in \Delta(\mathcal{H})} \mathcal{L}(\pi, \hat{\lambda}) + 2\nu \\ &\leq \mathcal{L}(\pi^*, \hat{\lambda}) + 2\nu = \text{OPT} + 2\nu\end{aligned}$$

where the equality follows from the fact that  $\mathcal{L}(\pi^*, \hat{\lambda}) = \text{OPT}$ .

Next, we will bound the fairness constraint violations. Suppose without loss of generality that the following fairness constraint is violated:  $\text{FPR}_1(\hat{\pi}) - \text{FPR}_{-1}(\hat{\pi}) = \gamma + \alpha$  for some  $\alpha \geq 0$ . Let  $\lambda' \in \Lambda$  such that  $\lambda'_1 = 2$ . Then

$$\mathcal{L}(\hat{\pi}, \hat{\lambda}) + \nu \geq \mathcal{L}(\hat{\pi}, \lambda') = \text{err}(\hat{\pi}, \mathcal{P}) + 2\alpha$$

Thus, by the assumption of approximate saddle point,

$$\text{err}(\hat{\pi}, \mathcal{P}) \leq \mathcal{L}(\hat{\pi}, \hat{\lambda}) + \nu - 2\alpha \leq \mathcal{L}(\pi^*, \hat{\lambda}) + 2\nu - \alpha = \text{OPT} + 2\nu - \alpha.$$

Now consider a distribution  $\pi'$  that is defined as the mixture of

$$\pi' = (1 - \alpha)\hat{\pi} + \alpha\mathbf{1}[a = -1].$$

This means

$$\begin{aligned} \text{FPR}_1(\pi') &= (1 - \alpha)\text{FPR}_1(\hat{\pi}) + \alpha\text{FPR}_1(\mathbf{1}[a = -1]) = (1 - \alpha)\text{FPR}_1(\hat{\pi}) \\ \text{FPR}_{-1}(\pi') &= (1 - \alpha)\text{FPR}_{-1}(\hat{\pi}) + \alpha\text{FPR}_{-1}(\mathbf{1}[a = -1]) = (1 - \alpha)\text{FPR}_{-1}(\hat{\pi}) + \alpha \end{aligned}$$

It follows that

$$\text{FPR}_1(\pi') - \text{FPR}_{-1}(\pi') = (1 - \alpha)(\gamma + \alpha) - \alpha \leq \gamma$$

which implies that  $\pi'$  is a feasible solution for the fair ERM problem. This implies that

$$\text{err}(\hat{\pi}, \mathcal{P}) \geq \text{err}(\pi', \mathcal{P}) - \alpha \geq \text{OPT} - \alpha$$

Thus, we have  $\text{OPT} + 2\nu - 2\alpha \geq \text{OPT} - \alpha$ , which implies that  $\alpha \leq 2\nu$ . This completes the proof.

To facilitate our analysis, we would like a solution  $\hat{\pi}$  that satisfies the fairness constraint without any violation. To achieve this, we simply tighten the constraint by an amount of  $2\nu$  and compute the  $\nu$ -approximate saddle point for the tightened Lagrangian, replacing  $\gamma$  with  $\gamma' = \gamma - 2\nu$ . We also need ensure such tightening of the constraint does not severely increase the resulting error.

**Lemma B.6 (Bound on additional error from tightening.)** *Suppose that  $\gamma > 2\nu$ . Let  $\text{OPT}'$  be the objective value at the optimum for the tightened optimization problem:*

$$\begin{aligned} & \min_{\pi \in \Delta(\mathcal{H})} \mathbb{E}_{h \sim \pi} \left[ \sum_{j=1}^n w_j \mathbf{1}\{h(X_j) \neq Y_j\} \right] \\ & \text{such that } \forall j \in \{\pm 1\} \quad \text{FPR}_j(\pi) - \text{FPR}_{-j}(\pi) \leq \gamma - 2\nu \end{aligned}$$

Then as long as that the class  $\mathcal{H}$  contains the two classifiers  $\mathbf{1}[a = j]$  for both  $j \in \{\pm 1\}$ ,  $\text{OPT}' - \text{OPT} \leq 2\nu$ .

**Proof B.7** *Let  $\pi^*$  be an optimal solution to the original (un-tightened) problem. Suppose without loss of generality that the following fairness constraint is violated:  $\text{FPR}_1(\hat{\pi}) - \text{FPR}_{-1}(\hat{\pi}) \geq 0$ . Now consider a distribution  $\pi'$  that is defined as the mixture of*

$$\pi' = (1 - 2\nu)\pi^* + 2\nu\mathbf{1}[a = -1].$$

Consequently, we can write

$$\begin{aligned} \text{FPR}_1(\pi') &= (1 - 2\nu)\text{FPR}_1(\pi^*) + 2\nu\text{FPR}_1(\mathbf{1}[a = -1]) = (1 - 2\nu)\text{FPR}_1(\pi^*) \\ \text{FPR}_{-1}(\pi') &= (1 - 2\nu)\text{FPR}_{-1}(\pi^*) + 2\nu\text{FPR}_{-1}(\mathbf{1}[a = -1]) = (1 - 2\nu)\text{FPR}_{-1}(\pi^*) + 2\nu \end{aligned}$$

It follows that

$$\text{FPR}_1(\pi') - \text{FPR}_{-1}(\pi') = (1 - 2\nu)(\gamma + 2\nu) - 2\nu \leq \gamma$$

which implies that  $\pi'$  is a feasible solution for the fair ERM problem. This implies that

$$\text{err}(\hat{\pi}, \mathcal{P}) \geq \text{err}(\pi', \mathcal{P}) - 2\nu \geq \text{OPT} - 2\nu$$

This completes the proof.

Next, we translate the approximation guarantee for the normalized weighted classification problem to the original cost-sensitive classification problem. This leads to our guarantee stated below.

**Lemma B.8** *For any  $0 < \nu < \gamma/2$ , there exists an oracle-efficient algorithm that calls CSC oracle over  $\mathcal{H}$  at most  $O(1/\nu^2)$  times and computes a solution  $\hat{\pi}$  that satisfies  $\gamma$ -EFP and has total cost*

$$\mathbb{E}_{h \sim \hat{\pi}} \left[ \sum_{j=1}^n (C_j^1 h(X_j) + C_j^0 (1 - h(X_j))) \right] \leq \text{OPT}_C + \epsilon$$

with  $\epsilon = 4\nu \sum_{j=1}^n |C_j^1 - C_j^0|$ .

The result of Lemma B.8 shows a computationally efficient algorithm that returns an approximate CSC solution with support size at most  $O(1/\nu^2)$ . Finally, we will shrink the support of the solution. To derive a sparse-support solution, we consider a linear program that computes a probability distribution over the support of  $\hat{\pi}$ . Then we will compute a basic solution obtain the final sparse solution (e.g. by running a variant of the ellipsoid algorithm [17]).

### B.3 Missing Details and Proofs in Section 3.2

---

**Algorithm 1:** Coordinate descent algorithm for solving the feasibility program

---

1 **Input:** history  $H_t$  from previous rounds; minimum probability  $\mu_t$ ; target accuracy parameter  $\nu$

2 **Initialize:**  $Q = \mathbf{0}$ ; Call FairCSC( $\nu$ ) to compute the policy  $\pi_0$  that approximately minimizes  $\hat{L}_t(\pi)$  (up to error  $\nu$ ).

**for**  $\pi \in \Pi$  **do**

Let  $\widetilde{\text{Reg}}_t(\pi) = \max\{\hat{L}_t(\pi) - \hat{L}_t(\pi_0), 0\}$ ,  $\tilde{b}_t(\pi) = \frac{\widetilde{\text{Reg}}_t(\pi)}{4(e-2)\mu_t \ln(T)}$

**end**

**for**  $\pi \in \Pi$  **do**

$V_\pi(Q) = \mathbb{E}_{x \sim H_t} [1/Q^{\mu_t}(\pi(x) | x)]$

$S_\pi(Q) = \mathbb{E}_{x \sim H_t} [1/Q^{\mu_t}(\pi(x) | x)^2]$

$\tilde{D}_\pi(Q) = V_\pi(Q) - (4 + \tilde{b}_{t-1}(\pi))$

**end**

3 **if**  $\int_{\pi \in \Pi} Q(\pi)(4 + \tilde{b}_\pi) > 4$  **then**

Replace  $Q$  by  $cQ$  with

$c = \frac{4}{\int_{\pi} Q(\pi)(4 + \tilde{b}_{t-1}(\pi))} < 1$

**end**

4 **if** calling FairCSC( $\nu$ ) for  $\pi$  approximating  $\max_{\pi'} \tilde{D}_{\pi'}(Q)$ , we have  $\tilde{D}_\pi(Q) > 0$  **then**

Add the following (positive) quantity to  $Q(\pi)$  while keeping all other weights unchanged:

$\alpha_\pi(Q) = \frac{V_\pi(Q) + \tilde{D}_\pi(Q)}{2(1 - 2\mu_t)S_\pi(Q)}$

**end**

5 **else**

Halt. If the sum of the weights  $Q$  is smaller than 1, let  $Q$  place the remaining weight on  $\pi_0$ .

Output  $Q$  (note the algorithm will draw from  $Q^{\mu_t}$ ).

**end**

---

**Proof B.9 (Proof of lemma 3.9)** *The first oracle call is used to approximately solve*

$$\begin{aligned} \arg \min_{\pi} \hat{L}_t(\pi) &= \arg \min_{\pi} \frac{1}{t} \sum_{s=1}^t \ell_s \frac{\Pr[\pi(x_s) = a_s]}{Q_s(a_s | x_s)} \\ &= \frac{1}{\mu_t} \arg \min_{\pi} \sum_{s=1}^t \frac{\mu_t \ell_s}{t} \frac{\Pr[\pi(x_s) = a_s]}{Q_s(a_s | x_s)} \end{aligned}$$

where, because  $Q_s(a|x)$  is constrained to at least  $\mu_s$  (which is decreasing in  $s$ ), the argmin now has weights summing to at most 1. Therefore the oracle, given  $\nu$ , returns  $\tilde{\pi}$  such that

$$\min_{\pi} \hat{L}_t(\pi) \leq \hat{L}_t(\tilde{\pi}) \leq \min_{\pi} \hat{L}_t(\pi) + \frac{\nu}{\mu_t}.$$

This implies that, for all  $\pi$ ,

$$\widehat{\text{Reg}}_t(\pi) \geq \widetilde{\text{Reg}}_t(\pi) \geq \widehat{\text{Reg}}_t(\pi) - \frac{\nu}{\mu_t}.$$

This gives

$$b_t(\pi) \geq \tilde{b}_t(\pi) \geq b_t(\pi) - \Lambda_t.$$

If the first condition is met and the algorithm halts, then  $\int Q(\pi)(4 + \tilde{b}_t(\pi)) \leq 4$ , implying that the sum of  $Q$ 's weights is at most 1 (since  $\tilde{b}_t(\pi) \geq 0$ ), and implying that  $\int Q(\pi)(4 + b_t(\pi)) \leq 4 + \Lambda_t$ , which is the first inequality.

Next, the oracle is called once per loop to request

$$\arg \max_{\pi} \tilde{D}_{\pi}(Q) = \arg \max_{\pi} \sum_{s=1}^t \frac{1}{tQ_s^{\mu_s}(a_s | x_s)} - (4 + \tilde{b}_{t-1}(\pi))$$

There are two cases, where  $\widetilde{\text{Reg}}_t(\pi) = 0$  and otherwise. If 0, then we again obtain an additive  $\frac{\nu}{\mu_t}$  approximation. Otherwise, after dropping terms that don't depend on  $\pi$ , we have

$$\arg \max_{\pi} \sum_{s=1}^t \frac{1}{tQ_s^{\mu_t}(a_s | x_s)} - \frac{\ell_s \Pr[\pi(x_s) = a_s]}{4(e-2)\mu_t t \ln(T) Q_s(a_s | x_s)}$$

Scaling each term by  $4(e-2)\ln(T)\mu_t^2$  ensures that the sum of the weights is at most 1, implying that the approximation we get is again an additive  $\Lambda_t$ . So if  $\pi$  is chosen by the algorithm, then  $\max_{\pi'} \tilde{D}_{\pi'}(Q) \geq \tilde{D}_{\pi}(Q) \geq \tilde{D}_{\pi'}(Q) - \Lambda_t$ . Plugging in the guarantee for  $b_t$ , if we let  $D_{\pi}(Q) = V_{\pi}(Q) - (4 + b_{t-1}(\pi))$ , then we get

$$\max_{\pi^*} D_{\pi^*}(Q) + \Lambda_t \geq \tilde{D}_{\pi}(Q) \geq \max_{\pi^*} D_{\pi^*}(Q) - \Lambda_t.$$

So if the algorithm halts after obtaining  $\pi$  from the oracle with  $\tilde{D}_{\pi}(Q) \leq 0$ , then  $\max_{\pi^*} D_{\pi^*}(Q) \leq \Lambda_t$ , which implies the second guarantee.

To show convergence of the algorithm, consider the following potential function

$$\Phi(Q) = \frac{\mathbb{E}_{H_t} [\text{RE}(\mathcal{U}_2 \| Q^{\mu_t}(\cdot | x))]}{1 - 2\mu_t} + \frac{\int_{\pi \in \Pi} Q(\pi) \tilde{b}_{t-1}(\pi)}{4}$$

where  $\mathcal{U}_2$  denotes the uniform distribution over the two predictions and  $\text{RE}(p||q)$  denotes the unnormalized relative entropy between two nonnegative vectors  $p$  and  $q$  in  $\mathbb{R}^2$  (over the two predictions):

$$\text{RE}(p||q) = \sum_{\hat{y} \in \{\pm 1\}} (p_{\hat{y}} \ln(p_{\hat{y}}/q_{\hat{y}}) + q_{\hat{y}} - p_{\hat{y}}).$$

First, we note that any renormalization step does not increase potential, i.e. letting  $c = 4 / \int_{\pi} Q(\pi)(4 + \tilde{b}_t(\pi))$ , if  $c < 1$  (which is equivalent to the update condition) then  $\Phi(cQ) \leq \Phi(Q)$ . This is directly proven in

Lemma 6 of [2] and we do not re-prove it. The only difference is that where we used  $\tilde{b}_t(\pi)$  in the definition of  $c$  and  $\Phi$  [2] uses  $b_{t-1}(\pi)$ ; but the proof does not use any property of  $b_{t-1}(\pi)$  except nonnegativity.

Second, we note that a renormalization step can only occur once in a row; after that, either the algorithm halts, or the other condition ( $\tilde{D}_\pi(Q) > 0$ ) is triggered.

Third, when the other condition is triggered, the potential decreases significantly, specifically, by at least  $\frac{1}{4(1-2\mu_t)}$ . This is also directly proven in Lemma 7 of [2].<sup>1</sup> The only difference is that the proof in that paper uses  $\tilde{b}_t(\pi)$  instead of  $b_{t-1}(\pi)$ , which yields  $\tilde{D}_\pi(Q)$  rather than  $D_\pi(Q)$ . However, the only property of  $D_\pi(Q)$  used in the proof is  $D_\pi(Q) > 0$ , which is satisfied by  $\tilde{D}_\pi(Q)$  as well.

The potential begins with  $Q = \mathbf{0}$  at  $\Phi(Q) \leq \frac{\ln \frac{1}{1-2\mu_t}}{1-2\mu_t}$ , and remains nonnegative by definition, so after a polynomial number of steps, the algorithm satisfies both conditions and halts.

## B.4 Missing Proofs in Section 3.3

**Proof B.10 (Proof of Claim 3.10)** Let  $\pi$  be any policy in  $\Pi$ . Note, in particular, that  $-\mathbf{1} \in \Pi$  and let  $\pi'_{-1} \in \Pi_\eta$  such that

$$\min_{\pi' \in \Pi_\eta} \|-\mathbf{1} - \pi'_{-1}\|_\infty \leq \eta$$

Then, we can see that

$$\begin{aligned} |V(P, \pi, \mu) - V(P, \pi', \mu)| &\leq |V(P, -\mathbf{1}, \mu) - V(P, \pi'_{-1}, \mu)| \leq \frac{1}{\mu} - \frac{1}{\mu + \eta} = \frac{\eta}{\mu(\mu + \eta)} \\ |\hat{V}_t(P, \pi, \mu) - \hat{V}_t(P, \pi', \mu)| &\leq |\hat{V}_t(P, -\mathbf{1}, \mu) - \hat{V}_t(P, \pi'_{-1}, \mu)| \leq \frac{1}{\mu} - \frac{1}{\mu + \eta} = \frac{\eta}{\mu(\mu + \eta)} \end{aligned}$$

The following lemma follows directly from Lemma 10 of [2].

**Lemma B.11 (Full version of Lemma 3.11)** Fix any  $\mu \in [0, 1/2]$  and any  $\delta \in (0, 1)$ . Then, with probability  $1 - \delta$ ,

$$V(P, \pi, \mu) \leq 6.4\hat{V}_t(P, \pi, \mu) + \frac{75(1-2\mu) \ln |\Pi_\eta|}{\mu_t^2 t} + \frac{6.3 \ln(2|\Pi_\eta|^2 t^2 / \delta)}{\mu_t t} + \frac{2\eta}{\mu_t(\mu_t + \eta)}$$

for all probability distributions  $P$  over  $\Pi$ , all  $\pi \in \Pi$ , and for all  $t$ . In particular, if

$$\mu_t \geq \sqrt{\frac{\ln(2|\Pi_\eta| t^2 / \delta)}{2t}}, t \geq 8 \ln(2|\Pi_\eta| t^2 / \delta)$$

then,

$$V(P, \pi, \mu_t) \leq 6.4\hat{V}_t(P, \pi, \mu_t) + 162.6 + \frac{2\eta}{\mu_t(\mu_t + \eta)}$$

We will make use of the following concentration inequality.

**Lemma B.12 (Freedman's inequality [6])** Let  $Z_1, \dots, Z_n$  be a martingale difference sequence with  $Z_i \leq R$  for all  $i$ . Let  $V_n = \sum_{i=1}^n \mathbb{E}[Z_i^2 | Z_1, \dots, Z_{i-1}]$ . For any  $\delta \in (0, 1)$  and any  $\lambda \in [0, 1/R]$ , with probability at least  $1 - \delta$

$$\sum_{i=1}^n Z_i \leq (e-2)\lambda V_n + \frac{\ln(1/\delta)}{\lambda}$$

<sup>1</sup>In that paper the potential function is scaled by a factor of  $\tau \mu_t$  compared to here, where  $\tau > 0$ .

**Proof B.13 (Proof of Lemma 3.12)** By applying the Freedman's inequality and union bound, we know that with probability  $1 - \delta'$ , for all  $t \in [T]$ ,  $\pi \in \Pi$  and  $\lambda \in [0, 1/\mu_t]$ ,

$$|L(\pi) - \hat{L}_t(\pi)| \leq (e - 2)\lambda \left( \frac{1}{t} \sum_{s=1}^t V(Q_t, \pi, \mu_t) \right) + \frac{\ln(|\Pi_\eta|T/\delta')}{t\lambda} \quad (4)$$

By the result of Lemma 3.11, we know that with probability  $1 - \delta'$ , for all  $P \in \Pi$ , for any  $\mu_t$  and  $t \geq 8 \ln(2|\Pi_\eta|t^2/\delta')$ ,

$$V(P, \pi, \mu) \leq 6.4\hat{V}_t(P, \pi, \mu_t) + 162.6 + \frac{2\eta}{\mu_t(\mu_t + \eta)} \quad (5)$$

We will condition on events of (4) and (5) for the remainder of the proof, which occurs with probability at least  $1 - 2\delta'$ . Then we can further rewrite

$$|L(\pi) - \hat{L}_t(\pi)| \leq (e - 2)\lambda \left( \frac{1}{t} \sum_{s=1}^t \left( 6.4\hat{V}_t(Q_t, \pi, \mu_t) + 162.6 + \frac{2\eta}{\mu_t(\mu_t + \eta)} \right) \right) + \frac{\ln(|\Pi_\eta|T/\delta')}{\lambda t}$$

Recall that by the accuracy guarantee of Lemma 3.9, we know for all  $\pi \in \Pi$ ,

$$\hat{V}_t(Q_t, \pi, \mu_t) \leq 4 + b_{t-1}(\pi) + \Lambda_{t-1}$$

Thus, we can further bound

$$\begin{aligned} & |L(\pi) - \hat{L}_t(\pi)| \\ & \leq (e - 2)\lambda \left( \frac{1}{t} \sum_{s=1}^t \left( 6.4(4 + b_{s-1}(\pi) + \Lambda_{s-1}) + 162.6 + \frac{2\eta}{\mu_t(\mu_t + \eta)} \right) \right) + \frac{\ln(|\Pi_\eta|T/\delta')}{\lambda t} \\ & \leq (e - 2)\lambda \left( 188.2 + \frac{1}{t} \sum_{s=1}^t \left( 6.4b_{s-1}(\pi) + 6.4\Lambda_{s-1} + \frac{2\eta}{\mu_t(\mu_t + \eta)} \right) \right) + \frac{\ln(|\Pi_\eta|T/\delta')}{\lambda t} \end{aligned}$$

To complete the proof, we will set  $\delta' = \delta/2$ .

**Proof B.14 (Proof of Lemma 3.13)** To simplify notation, let

$$C_t = (e - 2)\lambda \left( 188.2 + \frac{1}{t} \sum_{s=1}^t \left( 6.4\Lambda_{s-1} + \frac{\eta}{\mu_s(\mu_s + \eta)} \right) \right) + \frac{\ln(|\Pi_\eta|T/\delta)}{t\lambda}$$

Recall that

$$\Lambda_t := \frac{\nu}{4(e - 2)\mu_t^2 \ln(T)}.$$

Then as long as we have  $\nu \leq 1/T$  and  $\eta \leq 1/T^2$ , we have

$$C_t \leq 190(e - 2)\lambda + \frac{\ln(|\Pi_\eta|T/\delta)}{t\lambda}$$

We will prove our result by induction. First, the base case holds trivially given our choice of  $\epsilon_t$ . Next, we will show  $\text{Reg}(\pi) \leq 2\widehat{\text{Reg}}_t(\pi) + \epsilon_t$ , and  $\text{Reg}(\pi) \leq 2\text{Reg}_t(\pi) + \epsilon_t$  follows analogously. Observe that for any policy  $\pi$ , we can first decompose the regret difference as

$$\text{Reg}(\pi) - \widehat{\text{Reg}}_t(\pi) \leq (L(\pi) - \hat{L}_t(\pi)) - (L(\pi^*) - \hat{L}_t(\pi^*))$$

where  $\pi^*$  denotes the optimal policy in  $\Pi$ . Then using the result of Lemma 3.12, we can further bound the regret difference as follows: for any  $\lambda \in [0, \mu_t]$ ,

$$\begin{aligned}
& \text{Reg}(\pi) - \widehat{\text{Reg}}_t(\pi) \\
& \leq \frac{6.4(e-2)\lambda}{t} \left( \sum_{s=1}^t b_{s-1}(\pi) + b_{s-1}(\pi^*) \right) + 2C_t \\
& = \frac{1.6\lambda}{\mu_t \ln(T)t} \left( \sum_{s=1}^t \widehat{\text{Reg}}_{s-1}(\pi) + \widehat{\text{Reg}}_{s-1}(\pi^*) \right) + 2C_t \\
& \leq \frac{3.2\lambda}{\mu_t \ln(T)t} \left( \sum_{s=1}^t \text{Reg}(\pi) + \text{Reg}(\pi^*) + \epsilon_{s-1} \right) + 2C_t && \text{(Induction hypothesis)} \\
& \leq \frac{3.2\lambda}{\mu_t \ln(T)t} \left( t \text{Reg}(\pi) + \sum_{s=1}^t \epsilon_{s-1} \right) + 2C_t && \text{(Reg}(\pi^*) = 0) \\
& \leq \frac{3.2\lambda}{\mu_t \ln(T)} \text{Reg}(\pi) + \frac{3.2\lambda}{\mu_t \ln(T)t} \left( \sum_{s=1}^t \epsilon_{s-1} \right) + 2C_t
\end{aligned}$$

We will set  $\lambda = \mu_t/3.2$ , which allows us to simplify the bound

$$\text{Reg}(\pi) - \widehat{\text{Reg}}(\pi) \leq \frac{\text{Reg}(\pi)}{\ln(T)} + \frac{1}{\ln(T)t} \left( \sum_{s=1}^t \epsilon_{s-1} \right) + 2C_t$$

Since  $(1 - 1/\ln(T)) > 1/2$  and  $\mu_t = \frac{3.2 \ln(|\Pi_\eta|T/\delta)}{\sqrt{t}}$ , it follows that

$$\begin{aligned}
\text{Reg}(\pi) & \leq 2\widehat{\text{Reg}}(\pi) + \frac{2}{\ln(T)t} \left( \sum_{s=1}^t \epsilon_{s-1} \right) + 4C_t \\
& \leq 2\widehat{\text{Reg}}(\pi) + \frac{2}{\ln(T)t} \left( \sum_{s=1}^t \epsilon_{s-1} \right) + 4 \left( \frac{190(e-2) \ln(|\Pi_\eta|T/\delta)}{\sqrt{t}} + \frac{1}{\sqrt{t}} \right) \\
& \leq 2\widehat{\text{Reg}}(\pi) + \frac{2}{\ln(T)t} \left( \sum_{s=1}^t \epsilon_{s-1} \right) + \frac{560 \ln(|\Pi_\eta|T/\delta)}{\sqrt{t}}
\end{aligned}$$

Observe that  $\sum_{s=1}^t \epsilon_{s-1} = 1000(\ln(|\Pi_\eta|T/\delta)) \sum_{s=1}^{t-1} \frac{1}{s} \leq 1000(\ln(|\Pi_\eta|T/\delta))\sqrt{t}$ . This means

$$\text{Reg}(\pi) \leq 2\widehat{\text{Reg}}(\pi) + \frac{2000}{\ln(T)\sqrt{t}}(\ln(|\Pi_\eta|T/\delta)) + \frac{560 \ln(|\Pi_\eta|T/\delta)}{\sqrt{t}} \leq 2\widehat{\text{Reg}}(\pi) + \epsilon_t$$

where the last inequality holds as long as  $\ln(T) \geq 5$ .

## C Lower Bound Proof

In this section, we prove theorem 4.1. We make use of a couple of standard tools:



**Lemma C.1** (*Pinsker's Inequality*) Let  $\mathcal{D}_1, \mathcal{D}_2$  be probability distributions. Let  $A$  be any event. Then:

$$|\mathcal{D}_1(A) - \mathcal{D}_2(A)| \leq \sqrt{\frac{1}{2}KL(\mathcal{D}_1||\mathcal{D}_2)}$$

The following is a simple corollary that follows from the additivity of KL-divergence over product distributions.

**Corollary C.2** Let  $t \in \mathbb{N}$ . Consider the product distributions  $\mathcal{D}_1^t, \mathcal{D}_2^t$ . For any event  $A$ ,

$$|\mathcal{D}_1^t(A) - \mathcal{D}_2^t(A)| \leq \sqrt{\frac{1}{2}t \cdot KL(\mathcal{D}_1||\mathcal{D}_2)}$$

Next, for any algorithm  $\mathcal{A}$ , round  $t$ , hypothesis  $h$ , and distribution  $\mathcal{D}$ , let

$$P_t(h, \mathcal{D}) = \mathbb{P}[\mathcal{A} \text{ plays } h \text{ on round } t]$$

when given inputs from  $\mathcal{D}$ . We say an algorithm  $(\beta, t, h)$ -distinguishes distributions  $\mathcal{D}_1$  and  $\mathcal{D}_2$  if

$$|P_t(h, \mathcal{D}_1) - P_t(h, \mathcal{D}_2)| > \beta.$$

**Lemma C.3** Let  $\mathcal{D}_1, \mathcal{D}_2$  be two probability distributions. No algorithm can  $(\beta, t, h)$ -distinguish  $\mathcal{D}_1$  and  $\mathcal{D}_2$  for any  $h$  and  $t \leq \frac{2\beta^2}{KL(\mathcal{D}_1||\mathcal{D}_2)}$ .

**Proof C.4** Assume for contradiction that there exists an algorithm that  $(\beta, t, h)$ -distinguishes  $\mathcal{D}_1$  and  $\mathcal{D}_2$  for some  $h$  and  $t \leq \frac{2\beta^2}{KL(\mathcal{D}_1||\mathcal{D}_2)}$ . This defines an event  $A$  such that

$$|\mathcal{D}_1^t(A) - \mathcal{D}_2^t(A)| > \beta \geq \sqrt{\frac{1}{2}tKL(\mathcal{D}_1||\mathcal{D}_2)}$$

which contradicts corollary C.2.

With these tools in hand, we are ready to prove the lower bound (Theorem 4.1):

**Proof C.5 (Proof of Theorem 4.1)** Fix any  $\alpha \in (0, 0.5)$  and let  $T \geq \sqrt[3]{16}$ . Denote  $\gamma = T^{-\alpha}$ . Fix any  $\delta \leq 0.24$ .

Define the following distributions over  $(X, A, Y)$ :

$\mathcal{D}_1$  given by:

		$x_1$	$x_2$	$x_3$	$x_4$
$A = -1$	$\mathbb{P}[(x, a)]$	1/8	1/8	1/8	1/8
	$\mathbb{P}[y = 1 (x, a)]$	$0.5 + 4\gamma$	$0.5 - 4\gamma$	1	0
$A = +1$	$\mathbb{P}[(x, a)]$	1/8	1/8	1/8	1/8
	$\mathbb{P}[y = 1 (x, a)]$	$0.5 - 4\gamma$	$0.5 + 4\gamma$	1	0

$\mathcal{D}_2$  given by:

		$x_1$	$x_2$	$x_3$	$x_4$
$A = -1$	$\mathbb{P}[(x, a)]$	1/8	1/8	1/8	1/8
	$\mathbb{P}[y = 1 (x, a)]$	$0.5 + 4\gamma$	$0.5 - 4\gamma$	1	0
$A = +1$	$\mathbb{P}[(x, a)]$	1/8	1/8	1/8	1/8
	$\mathbb{P}[y = 1 (x, a)]$	$0.5 + 4\gamma$	$0.5 - 4\gamma$	1	0

The available hypotheses  $\mathcal{H} = \{-1, +1, h_1, h_2\}$  are defined as:

		$x_1$	$x_2$	$x_3$	$x_4$
$-1$	$A = -1$	-1	-1	-1	-1
	$A = +1$	-1	-1	-1	-1
$+1$	$A = -1$	+1	+1	+1	+1
	$A = +1$	+1	+1	+1	+1
$h_1$	$A = -1$	+1	-1	+1	-1
	$A = +1$	+1	-1	+1	-1
$h_2$	$A = -1$	+1	-1	+1	-1
	$A = +1$	-1	+1	+1	-1

The performance of the hypotheses in  $\mathcal{H}$  on the two distributions is given by: On  $\mathcal{D}_1$ :

	$L_{\mathcal{D}}^{0-1}(h)$	$\Delta_{FPR}(h)$
$-1$	0.5	0
$+1$	0.5	0
$h_1$	0.25	$4\gamma$
$h_2$	$0.25 - 2\gamma$	0

On  $\mathcal{D}_2$ :

	$L_{\mathcal{D}}^{0-1}(h)$	$\Delta_{FPR}(h)$
$-1$	0.5	0
$+1$	0.5	0
$h_1$	$0.25 - 2\gamma$	0
$h_2$	0.25	$4\gamma$

Note that on both distributions,  $h_1$  and  $h_2$  both have substantially lower error than the two constant classifiers, but only one of them satisfies the  $\gamma$ -fairness constraint — and which one of them it is depends on whether the underlying distribution is  $\mathcal{D}_1$  or  $\mathcal{D}_2$ . Note also that one of them always satisfies a 0-fairness constraint, and so sets the benchmark for 0-EFP regret. The main fact driving our lower bound is that until the algorithm can reliably distinguish  $\mathcal{D}_1$  from  $\mathcal{D}_2$ , it must place substantial weight on the constant classifiers, incurring high regret.

We first establish that the two distributions are hard to distinguish by showing that the KL-divergence

between  $\mathcal{D}_1$ ,  $\mathcal{D}_2$  is bounded by  $O(\gamma^2)$ :

$$\begin{aligned}
KL(\mathcal{D}_1||\mathcal{D}_2) &= \frac{2}{8} \left( \frac{1+8\gamma}{2} \ln \left( \frac{1+8\gamma}{1-8\gamma} \right) + \frac{1-8\gamma}{2} \ln \left( \frac{1-8\gamma}{1+8\gamma} \right) \right) \\
&= \gamma \ln \left( \frac{\frac{1+8\gamma}{1-8\gamma}}{\frac{1-8\gamma}{1+8\gamma}} \right) \\
&= \gamma \ln \left( \left( \frac{1+8\gamma}{1-8\gamma} \right)^2 \right) \\
&= 2\gamma \ln \left( \frac{1+8\gamma}{1-8\gamma} \right) \\
&= 2\gamma \ln \left( 1 + \frac{16\gamma}{1-8\gamma} \right) \\
&\leq 2\gamma \frac{16\gamma}{1-8\gamma} \\
&= \frac{64\gamma^2}{2(1-8\gamma)} \\
&\leq 64\gamma^2
\end{aligned}$$

Let  $\mathcal{A}$  be a  $\gamma$ -EFP( $\delta$ ) algorithm. Let  $K = \frac{0.01^2}{32\gamma^2}$  (and note that, for  $\alpha \in (0, 0.5)$ ,  $K = \frac{0.01^2}{32\gamma^2} = \frac{0.01^2 T^{2\alpha}}{32} < \frac{0.01^2 T}{32} \leq T$ ). Let  $t \leq K$  (note that the number of samples observed by time  $t$  is  $t' \leq t$ ); then by lemma C.3,

$$P_t(h_1, \mathcal{D}_2) \leq P_t(h_1, \mathcal{D}_1) + 0.01$$

$$P_t(h_2, \mathcal{D}_1) \leq P_t(h_2, \mathcal{D}_2) + 0.01$$

Observe that any convex combination  $\pi$  of classifiers played under  $\mathcal{D}_1$  fails to satisfy the  $\gamma$ -EFP constraint unless it puts weight less than  $1/4$  on  $h_1$ . Similarly, any convex combination  $\pi$  of classifiers played under  $\mathcal{D}_2$  fails to satisfy the  $\gamma$ -EFP constraint unless it puts weight less than  $1/4$  on  $h_2$ . Since by definition, and  $\gamma$ -EFP( $\delta$ ) algorithm plays only  $\gamma$ -EFP hypotheses on any distribution it is played on except with probability  $\delta$ , we have that for all  $t \in [T]$

$$P_t(h_1, \mathcal{D}_1) \leq \frac{1}{4} + \delta$$

$$P_t(h_2, \mathcal{D}_2) \leq \frac{1}{4} + \delta$$

And thus

$$P_t(h_1, \mathcal{D}_2) \leq P_t(h_1, \mathcal{D}_1) + 0.01 = 0.25 + 0.01 + \delta = 0.26 + \delta$$

$$P_t(h_2, \mathcal{D}_1) \leq P_t(h_2, \mathcal{D}_2) + 0.01 = 0.25 + 0.01 + \delta = 0.26 + \delta$$

Hence on either distribution, we have,

$$\mathbb{P}[\mathcal{A} \text{ plays } +1 \text{ or } -1 \text{ on round } t] \geq 1 - (0.25 + \delta) - (0.26 + \delta) = 0.49 - 2\delta$$

The best performing 0-EFP policy on  $\mathcal{D}_1$  is  $h_2$ , while on  $\mathcal{D}_2$  it is  $h_1$ . Both of these induce expected per-round loss of less than  $\frac{1}{4}$ . Since the expected per round loss of either  $+1$  or  $-1$  is  $\frac{1}{2}$  on both distributions, if

*+1 or -1 are played with constant probability, the expected per-round regret incurred is a constant bounded away from zero. As a result, the expected 0-EFP regret of  $\mathcal{A}$  is at least  $\Omega(K) = \Omega\left(\frac{1}{\gamma^2}\right)$ .*

*The result is that any  $T^{-\alpha}$ -EFP( $\delta$ ) algorithm must have expected 0-EFP regret of  $\Omega(T^{2\alpha})$ .*