

Group Fairness in Bandits with Biased Feedback

Candice Schumann
University of Maryland
College Park, USA
candiceschumann@gmail.com

Nicholas Mattei
Tulane University
New Orleans, USA
nsmattei@gmail.com

Zhi Lang
University of Maryland
College Park, USA
willylang1996@gmail.com

John P. Dickerson
University of Maryland
College Park, USA
johnd@umd.edu

ABSTRACT

We propose a novel formulation of group fairness with biased feedback in the contextual multi-armed bandit (CMAB) setting. In the CMAB setting, a sequential decision maker must, at each time step, choose an arm to pull from a finite set of arms after observing some context for each of the potential arm pulls. In our model, arms are partitioned into two or more sensitive groups based on some protected feature(s) (e.g., age, race, or socio-economic status). Initial rewards received from pulling an arm may be distorted due to some unknown societal or measurement bias. We assume that in reality these groups are equal despite the biased feedback received by the agent. To alleviate this, we learn a societal bias term which can be used to both find the source of bias and to potentially fix the problem outside of the algorithm. We provide a novel algorithm that can accommodate this notion of fairness for an arbitrary number of groups, and provide a theoretical bound on the regret for our algorithm. We validate our algorithm using synthetic data and two real-world datasets for intervention settings wherein we want to allocate resources fairly across groups.

KEYWORDS

Group fairness; fair bandits; contextual bandits; human collaboration

ACM Reference Format:

Candice Schumann, Zhi Lang, Nicholas Mattei, and John P. Dickerson. 2022. Group Fairness in Bandits with Biased Feedback. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 9 pages.

Knowing that one may be subject to bias is one thing; being able to correct it is another.

Jon Elster

1 INTRODUCTION

In many online settings, a computational or human agent must sequentially select an item from a slate, receive feedback on that selection, and then use that feedback to learn how to select the best items in the following rounds. Within computer science, economics,

and operations research circles, this is typically modeled as a *multi-armed bandit (MAB)* problem [46]. Examples include algorithms for selecting what advertisements to display to users on a webpage [28], systems for dynamic pricing [29], and content recommendation services [24]. Indeed, such decision-making systems continue to expand in scope, making ever more important decisions in our lives such as setting bail [13], making hiring decisions [6, 37], and policing [35]. Thus, the study of the properties of these algorithms is of paramount importance as highlighted by Chouldechova and Roth [11] motivating priorities for fairness research in machine learning.

In the basic MAB setting, there are n arms, each associated with a fixed but unknown reward probability distribution [2, 23]. At each time step $t \in T$, an agent pulls an arm and receives a reward that is independent of any previous action and follows the selected arm's probability distribution. The goal of the agent is to maximize the total collected reward over time. A generalization of MAB is the contextual multi-armed bandit (CMAB) where the agent observes a d -dimensional *context* along with the observed rewards to choose a new arm. In the CMAB problem, the agent learns the relationship between contexts and rewards and selects the best arm [1].

Yet, the use of MAB- and CMAB-based systems often results in behavior that is societally repugnant. Sweeney [47] noted that queries for public records on Google resulted in different contextual advertisements based on whether the query target had a traditionally African American or Caucasian name; in the former case advertisements were more likely to contain text relating to criminal incidents. Following that initial report similar instances continue to be observed, both in the bandit setting and in the general machine learning world [31]. In lockstep, the academic community has begun developing approaches to tackling issues of (un)fairness in learning settings. We have an opportunity to identify and understand why the data we have may be *causing* the bias.

A Computing Community Consortium (CCC) report on fairness in ML identified that most studies of fairness are focused on classification problems [11]. These works define a statistical notion of fairness, typically a notion of equal treatment of equals [33], and propose algorithms to abide by these constraints. Two issues identified by Chouldechova and Roth [11] that we address in this paper are extensions to notions of *group fairness* and looking at fairness in *online dynamic systems*, e.g., CMABs. We address these gaps by formalizing and providing an algorithm for fairness with biased feedback when the arms of the bandit can be partitioned into groups. Direct applications of our work including scenarios

discussed within the AAMAS community like aiding the allocation of human resources in talent sourcing [38].

The recent AI100 study [26], whose goal is to take a broad and long term look at the opportunities and pitfalls for AI researchers, has highlighted the need to develop systems that work *with* humans, providing oversight, transparency, and explanation. Our bandit formulation is one step towards creating more human-centered AI [41], a new area of study that seeks to understand and balance computer automation and autonomy with the level of human control in a given system. Many of the negative applications of MAB based systems we have discussed so far too often occur because there is too much autonomy given to the system, and it optimizes away from what humans or society considers desirable. By explicitly modeling the underlying bias term, we hope to improve computer aided decision making by understanding and mitigating the dangers that can occur when there are excessive levels of human control or excessive levels of computer autonomy; leading to systems that are more transparent, auditable, and trustworthy.

Running Example. As a running example throughout the paper, imagine the position of an agent at a bank or a lender on a micro-lending site. Here, the agent must sequentially pick loans to fund. In many cases, such as the micro-lending site Kiva, a user is presented with a slate of potential loans they may fund when they log in and this slate is generated by a recommender system [45]. Each of these loans, i.e. arms, has a context which includes attributes of the applicant (e.g., personal statement, repayment history, business plan). The loans can also be partitioned into sets of m sensitive attributes, e.g. location, race, or gender. In the simplest case, assume we have two female applicants and two male applicants on the slate at a given time. We also assume that when pulling an arm from, for example, a female applicant, there is some societal bias introduced into the reward. Yet, in many settings (and, as we assume in this work), the average *true* (i.e., unbiased) reward across groups is equal. We want to balance the number of times the agent selects women versus men given this societal bias built into the feedback.

While we use loans as our running example, our notion of regret could be extended to a number of other areas including recent work in MAB problems on hiring situations [39], including the recent AAMAS Blue Sky Paper by Schumann et al. [38] specifically calling for the community to contribute to fair hiring. One could imagine a situation where hiring decisions are made w.r.t. a short-term reward signal that is biased,¹ versus a longer-term reward of performance which is less biased, e.g., via an end-of-year review that is based on a more quantitative metric such as on-the-job performance. A similar argument can be made about school admissions or matching workers to online tasks in a crowdwork setting.

Contributions. We propose a novel formulation of group fairness in the contextual multi-armed bandit (CMAB) setting. In our model, arms are partitioned into two or more sensitive groups based on some protected feature, e.g., race. Despite the fact that there may be differences in expected payout between the groups, we may wish to ensure some form of fairness between picking arms from the

various groups. Our goal is to capture the phenomena where we want to balance the arms being pulled from both groups and (learn to) ignore societal bias generated by sensitive group membership. We define two novel notions of reward and regret to capture implicit societal bias: proportional parity and equal group parity. We provide a novel algorithm that can accommodate these notions of fairness for an arbitrary number of groups, learn the societal bias term itself, and provide bounds on the regret for our algorithm. We validate our algorithms using synthetic data and real-world datasets for intervention settings wherein we want to allocate resources fairly across protected groups.²

2 RELATED WORK

Fairness in machine learning has become one of the most active topics in computer science [11]. The idea of using formal notions of fairness, i.e. axioms or properties, to design decision schemes has a long history in economics and political economy [33, 53]. Typically within ML research, fairness is operationalized using the Rawlsian idea that similar individuals should be treated similarly; formally extended to the classification setting by Dwork et al. [15], who provided algorithms to ensure individual fairness at the cost of the utility of the overall system. Their work underscores that in many cases statistical parity is not sufficient to ensure individual fairness, as we may treat groups fairly but in doing so may be very unfair to some specific individual. Determining when, how, and if to define fairness is an ongoing discussion with roots well before the time of computer science [44]; indeed, it is known that many natural conditions for fairness cannot be achieved in tandem [21]. Still, group fairness is found in many fielded systems [4, 51], and we focus on it here.

The study of fairness in MAB was initiated by Joseph et al. [19], who showed for both MAB and CMAB one can implement a fairness definition where within a given pool of applicants, e.g., college admission or mortgages, a worse applicant is not favored over a better one, despite a learning algorithm’s uncertainty over the true payoffs. However, Joseph et al. [19] only focus on individual fairness, and do not formally treat the idea of group fairness. Individual fairness is, in some sense, group fairness taken to an extreme, where every arm is its own singleton group; it offers strong guarantees, but under strong assumptions [5, 20].

Celis et al. [9] propose a bandit-based approach to personalization where arm pulls are constrained to fit some probability distribution defined by a fairness metric such as demographic parity. For example, when recommending news articles, their algorithm provides personalized articles from both left and right sources. Their formulation is perhaps closest in the literature to our formulation as it deals with group fairness, however it does not explicitly assume biased feedback. Instead, it enforces a fair probability distribution without learning about the bias present in the data.

There are a number of other recent studies of fairness in the MAB literature. Chen et al. [10] investigate a task allocation setting with a fairness constraint that captures a minimum rate at which

¹Class-based bias presents itself within seconds of an in-person interview; see <https://news.yale.edu/2019/10/21/yale-study-shows-class-bias-hiring-based-few-seconds-speech>.

²A full version of this paper, complete with an appendix containing proofs and additional experiments, can be found at <https://arxiv.org/abs/1912.03802>. We will also periodically update this work should typos or other errors be found; if you see any, please feel free to reach out! Code to reproduce the experiments is also available at <https://github.com/candiceschumann/groupfairtreatment>.

a task is assigned to a particular arm; their model is quite general and captures the adversarial and some non-stationary settings. Liu et al. [27] look at fairness between arms under the assumption that arm reward distributions are similar (another interpretation of equal treatment of equals). Patil et al. [32] define fairness such that each arm must be pulled for a predetermined required fraction over the total available rounds. Claire et al. [12] use the MAB framework to distribute resources amongst teammates in human-robot interaction settings; again, fairness is defined as a pre-configured minimum rate that each arm must be pulled. Hossain et al. [17] take a more theory-oriented approach to a similar setting, proposing a multi-agent variant of a stochastic MAB setting with a Nash social welfare definition of fairness.

Since preliminary versions of this work were presented [38, 40] there have been several papers that have investigated similar problems. Wang et al. [50] look at fairness of exposure in CMAB base systems, specifically focusing on similarity of merit, which is more in line with individual rather than the group fairness we consider here. Tang et al. [49] consider a setting inspired by liver transplantation where the objective is to trade off a more egalitarian, max-min, policy in allocating opportunities for surgeons to gain experience in liver transplant training. Finally, Ron et al. [34] investigate a setting of allocating opportunities to sub-populations in a corporate decision making setting where each arm needs to be pulled at least a budgeted number of times, but where the cost of allocating an opportunity to a non-optimal arm is known in advance. Interestingly, their algorithm also achieves a $T^{2/3}$ regret, similar to our results.

One needs to be careful when appealing to purely statistical metrics for ensuring fairness. As argued by Corbett-Davies and Goel [13], simply setting our sights on a form of classification parity, i.e., forcing that some statistical measure be normalized across groups, we may miss bigger picture issues. Specifically, by only focusing on the statistics of the data we have, we miss an opportunity to identify and understand why the data we have may be *causing* the bias. Later, we will argue that our novel formalization of regret allows us to actually learn particular sources of bias that may exist in our data.

3 PRELIMINARIES

We follow the standard CMAB setting and assume that we are attempting to maximize a measure over a series of time steps $t \in T$. We assume that there is a d -dimensional domain for the context space, $\mathcal{X} = \mathbb{R}^d$. The agent is presented with a set A of arms from which to select and we have $|A| = n$ total arms. Each of these arms is associated with a, possibly disjoint, context space $\mathcal{X}_i \subseteq \mathcal{X}$. Additionally, we assume that we have m sensitive groups and that the arms are partitioned into these sensitive groups such that $P_1 \cap \dots \cap P_m = \emptyset$, $P_1 \cup \dots \cup P_m = A$, and $\forall i \in m, |P_i| > 1$. For exposition's sake, we assume a binary sensitive attribute with $m = 2$ for the remaining of the paper. However, we show the generality of our results to any number of groups in Section 4.

Each arm i has a true linear reward function $f_i : \mathcal{X} \rightarrow \mathbb{R}$ such that $f_i(x) = \beta_i \cdot x$ where β_i is a vector of coefficients that is unknown to the agent. During each round $t \in T$, a context $x_{t,i} \in \mathcal{X}_i$ is given for each arm i . One arm is pulled per round. When arm i is pulled during round t , a reward is returned: $r_{t,i} = f_i(x_{t,i}) + e_{t,i}$ where

$e_{t,i} \sim \mathcal{N}(0, 1)$. The goal of the agent is to minimize the regret over all timesteps in T . Formally, the regret of the agent at timestep t is the difference between the arm selected and the best arm that could have been selected. Let i^* denote the optimal arm that could be selected and a be the selected arm. Then, the regret at T is

$$R(T) = \sum_{t=1}^T f(x_{i^*,t}) - f(x_{a,t}). \quad (1)$$

In this paper we compare our proposed algorithm against three other algorithms: TOPINTERVAL, a variation of LinUCB from Li et al. [24] with additional annotations to track group membership and treatment of arms, NAIVEFAIR which randomly picks a sensitive group and then applies TOPINTERVAL to that group,³ and INTERVALCHAINING, an individually fair algorithm from Joseph et al. [18]. All algorithms use ordinary least squares (OLS) estimators of the arm coefficients $\hat{\beta}_i$ with a confidence variable $w_{i,t}$ such that the true utility lies within $[\hat{\beta}_i \cdot x_{i,t} - w_{i,t}, \hat{\beta}_i \cdot x_{i,t} + w_{i,t}]$ with probability $1 - \delta$. NAIVEFAIR implements a naive version of demographic parity without explicitly looking at societal bias. TOPINTERVAL either explores by pulling an arm uniformly at random or exploits by pulling the arm with the highest upper confidence $\hat{\beta}_i \cdot x_{i,t} + w_{i,t}$. To ensure individual fairness, INTERVALCHAINING either explores by choosing an arm uniformly at random or exploits by pulling arms that have overlapping confidence intervals with the arm with the highest upper confidence.

3.1 Regret with Societal Bias

As mentioned before, ground truth rewards for sensitive groups can be noisy due to societal or measurement bias. We now formalize this bias in terms of multi-armed bandits. For ease of exposition we assume two groups, but we generalize this in our results. Again, we assume that n arms can be partitioned into two sets P_1 and P_2 such that $P_1 \cap P_2 = \emptyset$ and $P_1 \cup P_2 = [n]$. We consider P_1 , with $|P_1| > 1$ as the sensitive set, or set with some societal bias. In this situation, each arm i has another true utility function $f^*(x_{i,t}) = \beta_i \cdot x_{i,t}$ where β_i is a vector of coefficients; if arm i is pulled at timestep t the following reward is returned:

$$r_{i,t} = \beta_i \cdot x_{i,t} + \mathbb{1}[i \in P_1] \psi_{P_1} \cdot x_{i,t} + \mathcal{N}(0, 1), \quad (2)$$

where $\mathbb{1}[i \in P_1] = 1$ when $i \in P_1$ and 0 otherwise, and ψ_{P_1} is a societal or systematic bias against group P_1 . Note that ψ_{P_2} is a zero vector for the non-sensitive group. Hence, the underlying *biased* utility function can be written as $f(x_{i,t}) = \beta_i \cdot x_{i,t} + \mathbb{1}[i \in P_1] \psi_{P_1} \cdot x_{i,t}$.

Using our running example, let's assume that the down payment reward received has some bias against the male applicants compared to the female applicants, while the final repayment does not. Note that the final repayment is not measured after accepting a loan and is only measured much later. The loan agency should then take the bias into account while learning what 'good' applications look like. Or, in a hiring setting, an applicant may have a biased interview (initial reward) while their true performance is measured only after working for a year (later true reward).

³See Section 5.1 for more information

We define true regret for pulling an arm a at time T as

$$R^*(T) = \sum_{t=1}^T f^*(x_{i^*,t}) - f^*(x_{a,t}) \quad (3)$$

where i^* is the optimal arm to pull at timestep t and $f^*(x_{i,t})$ is the true reward with no bias terms $\psi_{P_1} \cdot x_{i,t}$. We also assume that the average true reward (with no bias) for group P_1 should be the same as the average reward for group P_2 . Compare this to Equation 1, which would return the regret on the biased reward function $f(x_{i,t})$. In the loan agency example, this real regret $f^*(x_{i,t})$ would measure the regret of the final repayments instead of the biased down payment regret.

One can view the societal bias term ψ_i that we learn for some group i as our algorithm learning how to automatically identify and adjust for anti-discrimination for group i compared to all other groups. Anti-discrimination is the practice of identifying a relevant feature in data and adjusting it to provide fairness under that measure [13]. One example of this, discussed by Dwork et al. [15], Joseph et al. [19], and in the official White House algorithmic decision making statement [30], comes up in college admissions. Given other factors, specifically income level, some colleges weight SAT scores *less* in wealthy populations due to the presence of tutors while increasing the weight of working-class populations [3]. While in these admissions settings the adjustments may be ad-hoc, we learn our bias term from data. Past work has compared the vector β learned for each arm as akin to adjusting for these biases [15]. While this is true at an *individual* level, our explicit modeling of bias allows us to discover these adjustments at a *group* level.

4 GROUP FAIR CONTEXTUAL BANDITS

In this section, given our new definition of reward (Equation 2) and corresponding new definition of regret (Equation 3), we present the algorithm GROUPFAIRTOPINTERVAL (Algorithm 1) which takes societal bias into account. We also give a bound on its regret in this new reward and regret setting. Subsequently, we briefly describe the algorithm.

In GROUPFAIRTOPINTERVAL, each round t is randomly chosen with probability $\frac{1}{t^{1/3}}$ to be an exploration round. The exploration round randomly chooses an arm to pull.

The remaining rounds become exploitation rounds, where linear estimates are used to pull arms. GROUPFAIRTOPINTERVAL learns two different types of standard OLS linear estimators [22]. The first is a coefficient vector $\hat{\beta}_{i,t}$ for each arm i (line 7). Additionally, GROUPFAIRTOPINTERVAL learns a group coefficient vector $\hat{\psi}_{P_j,t}$ for each group P_j (lines 4 and 5). To calculate these coefficient vectors, the algorithm keeps track of previous arm pull rewards for each arm i at every timestep t in a vector $Y_{i,t}$, and the corresponding contexts for each arm pull in a matrix $X_{i,t}$. A similar vector $\mathcal{Y}_{P_j,t}$ and matrix $\mathcal{X}_{P_j,t}$ is kept for both groups P_j . As mentioned previously, we treat P_1 as the sensitive group of arms. An arm i in the non-sensitive group P_2 has a reward estimation of $\hat{\beta}_{i,t} \cdot x_{i,t}$, while an arm i in the sensitive group P_1 has a bias corrected reward estimation of $\hat{\beta}_{i,t} \cdot x_{i,t} - \hat{\psi}_{P_1,t} + \hat{\psi}_{P_2,t}$.

For each arm i , the algorithm calculates confidence intervals $w_{i,t}$ around the linear estimates $\hat{\beta}_{i,t} \cdot x_{i,t}$ using a Quantile function

Q (line 9). This means that the true utility (including some bias) falls within $[\hat{\beta}_{i,t} \cdot x_{i,t} - w_{i,t}, \hat{\beta}_{i,t} \cdot x_{i,t} + w_{i,t}]$ with probability $1 - \delta$ at every arm i and every timestep t . Similarly, for each group P_j and context $w_{i,t}$ for a given arm i at timestep t , the algorithm calculates a confidence interval $b_{P_j,i,t}$ using a Quantile function Q (lines 4 and 5). This means that the true *group* utility (or true average group utility) falls within $[\hat{\psi}_{P_j,i,t} \cdot x_{i,t} - b_{P_j,i,t}, \hat{\psi}_{P_j,i,t} \cdot x_{i,t} + b_{P_j,i,t}]$ with probability $[1 - \delta]$. Using the confidence intervals $w_{i,t}$ and $b_{P_j,i,t}$, and the linear estimates $\hat{\beta}_{i,t} \cdot x_{i,t}$ and $\hat{\psi}_{P_j,i,t} \cdot x_{i,t}$, we calculate the upper bound of the estimated reward for each arm i (lines 15 and 17), pulling the arm with the highest upper bound (line 18).

Algorithm 1 GROUPFAIRTOPINTERVAL

Require: δ, P_1, P_2

```

1: for  $t = 1 \dots T$  do
2:   With probability  $\frac{1}{t^{1/3}}$ , play  $i_t \in_R \{1, \dots, n\}$  and observe re-
     ward  $y_{i_t,t}$ 
3:   otherwise:
4:      $\hat{\psi}_{P_1,t} := (\mathcal{X}_{P_1,t}^T \mathcal{X}_{P_1,t})^{-1} \mathcal{X}_{P_1,t}^T \mathcal{Y}_{P_1,t-1}$ 
5:      $\hat{\psi}_{P_2,t} := (\mathcal{X}_{P_2,t}^T \mathcal{X}_{P_2,t})^{-1} \mathcal{X}_{P_2,t}^T \mathcal{Y}_{P_2,t-1}$ 
6:     for  $i = 1 \dots n$  do
7:        $\hat{\beta}_{i,t} := (X_{i,t}^T X_{i,t})^{-1} X_{i,t}^T Y_{i,t-1}$ 
8:        $F_{i,t} := \mathcal{N}\left(0, \sigma^2 x_{i,t} \left(X_{i,t}^T X_{i,t}\right)^{-1} x_{i,t}^T\right)$ 
9:        $w_{i,t} := Q_{F_{i,t}}\left(\frac{\delta}{2nI}\right)$ 
10:      if  $i \in P_1$  then
11:         $\mathcal{F}_{P_1,i,t} := \mathcal{N}\left(0, \sigma^2 x_{i,t} \left(\mathcal{X}_{P_1,t}^T \mathcal{X}_{P_1,t}\right) x_{i,t}^T\right)$ 
12:         $\mathcal{F}_{P_2,i,t} := \mathcal{N}\left(0, \sigma^2 x_{i,t} \left(\mathcal{X}_{P_2,t}^T \mathcal{X}_{P_2,t}\right) x_{i,t}^T\right)$ 
13:         $b_{P_1,i,t} := Q_{\mathcal{F}_{P_1,i,t}}\left(\frac{\delta}{2|P_1|T}\right)$ 
14:         $b_{P_2,i,t} := Q_{\mathcal{F}_{P_2,i,t}}\left(\frac{\delta}{2|P_2|T}\right)$ 
15:         $\hat{u}_{i,t} := \hat{\beta}_{i,t} \cdot x_{i,t} + w_{i,t} - \hat{\psi}_{P_1,t} \cdot x_{i,t} + b_{P_1,i,t} + \hat{\psi}_{P_2,t} \cdot$ 
            $x_{i,t} + b_{P_2,i,t}$ 
16:      else
17:         $\hat{u}_{i,t} := \hat{\beta}_{i,t} \cdot x_{i,t} + w_{i,t}$ 
18:      Play  $\operatorname{argmax}_i \hat{u}_{i,t}$  and observe reward  $y_{i_t,t}$ 
```

Returning to our running example, using GROUPFAIRTOPINTERVAL, the loan agency would learn a down payment reward function for each of the arms, i.e., a coefficient vector β_i where $i \in [\text{young female arm, young male arm, older female arm, older male arm}]$, as well as the group average coefficients for the gender-grouped arms, ψ_{P_j} , for male and female. Using the gender-grouped coefficients, expected rewards for male arms are reweighted to account for the bias in down payment.

Standard algorithms like TOPINTERVAL⁴ would choose an arm $i = \operatorname{argmax}(\hat{\beta} \cdot x_{i,t} + w_{i,t})$, ignoring societal bias (Equation 2, leading to a larger true regret (Equation 3)). Note that GROUPFAIRTOPINTERVAL can be extended to multiple groups by defining an overall average reward.

⁴A variant of the contextual bandit LinUCB by Li et al. [24]

GROUPFAIRTOPINTERVAL is fair—in the context of our group fairness definitions—and satisfies the following theorem. Appendix B of the full paper provides a detailed, complete proof.

THEOREM 1. *For two groups P_1 and P_2 , where P_1 has a bias offset in rewards, GROUPFAIRTOPINTERVAL has regret*

$$R^*(T) = O\left(\sqrt{\frac{dn \ln \frac{2nT}{\delta}}{l}} T^{2/3} + \left(\frac{dnL}{l} \left(\ln^2 \frac{2nT}{\delta} + \ln d\right)\right)^{2/3}\right).$$

PROOF SKETCH. We start by proving two lemmas. The first of which states that with probability at least $1 - \delta$:

$$\left| \hat{\beta}_{i,t} \cdot x_{i,t} - (\beta_i \cdot x_{i,t} + \mathbb{1}[i \in P_1] \psi_{P_1} \cdot x_{i,t}) \right| \leq w_{i,t} \quad (4)$$

holds for any i at time t . Similarly, the second states that with probability at least $1 - \delta$:

$$\left| \hat{\beta}_{i,t} \cdot x_{i,t} - \beta_i \cdot x_{i,t} \right| \leq w_{i,t} \quad (5)$$

holds for any group P_j , any arm i , and at any timestep t . By combining these two lemmas, we can see that arms should be treated fairly.

The regret for GROUPFAIRTOPINTERVAL can be broken down into three terms:

$$R^*(T) = \sum_{t: t \text{ is an explore round}} \text{regret}(t) \quad (6)$$

$$+ \sum_{t: t \text{ is an exploit round and } t < T_1} \text{regret}(t) \quad (7)$$

$$+ \sum_{t: t \text{ is an exploit round and } t \geq T_1} \text{regret}(t).$$

First, for any t we have:

$$\sum_{t' < t} \frac{1}{t'^{1/3}} = \Theta(t^{2/3}). \quad (8)$$

We then show that the number of rounds T_1 after which we have sufficient samples such that the estimators are well concentrated is:

$$T_1 = \Theta\left(\min_a \left(\frac{dnL}{\lambda_{\min_{a,d}}} \left(\ln^2 \frac{2}{\delta} + \ln d \right) \right)^{3/2}\right). \quad (9)$$

Finally, we bound the third term in Equation 6 as follows:

$$\sum_{t: t \text{ is an exploit round and } t \geq T_1} \text{regret}(t) \quad (10)$$

$$\leq O\left(\sqrt{dn \frac{\ln \frac{2nT}{\delta}}{\min_i \lambda_{\min_{i,d}}}} T^{2/3} + \delta' T\right). \quad (11)$$

Combining Equations 6, 8, 9, and 10, we have Theorem 1. \square

Note that we can extend Algorithm 1 to m groups. In this setting, we make the strong assumption that true rewards are centered about ρ defined by the user.⁵ In this adaption of the algorithm, we set the upper bound radius for arm i as:

$$\hat{u}_{i,t} = \hat{\beta}_{i,t} \cdot x_{i,t} + w_{i,t} + \rho - \hat{\psi}_{P_j,t} \cdot x_{i,t} + b_{P_j,i,t}$$

where $i \in P_j$. We then have the following theorem for multiple groups:

⁵See Appendix B.2 for further details.

THEOREM 2. *For m groups P_1, \dots, P_m , GROUPFAIRTOPINTERVAL (MULTIPLE GROUPS) has regret*

$$R^*(T) = O\left(\sqrt{\frac{dn \ln \frac{2nT}{\delta}}{l}} T^{2/3} + \left(\frac{dnmL}{l} \left(\ln^2 \frac{2nT}{\delta} + \ln d\right)\right)^{2/3}\right).$$

where $l = \min_i \lambda_{\min_{i,d}}$, with $\lambda_{\min_{i,d}}$ the smallest eigenvalue of $X_{i,t}^T X_{i,t}$; and $L > \max_t \lambda_{\max}(x_{i,t}^T x_{i,t})$.

5 EXPERIMENTS

To empirically evaluate GROUPFAIRTOPINTERVAL, we perform experiments on synthetic data to demonstrate the effects of various parameters, and on real datasets to demonstrate how GROUPFAIRTOPINTERVAL performs in the wild. In each of these sections we compare to TOPINTERVAL, due to Li et al. [24], NAIVEFAIR (See Section 5.1), and INTERVALCHAINING, due to Joseph et al. [19].

5.1 NAIVEFAIR

One popular definition of group fairness in classification is the notion of demographic parity. Formally, given a protected demographic group A , we want:

$$\Pr(\hat{Y} = 1|A = 0) = \Pr(\hat{Y} = 1|A = 1), \quad (12)$$

where the probability of assigning a classification label $\hat{Y} = 1$ does not change based on the sensitive attribute class A . Demographic parity is important when ground truth classes Y are extremely noisy for sensitive groups due to some societal or measurement bias. Assume that we have a classifier that predicts whether an individual should receive a loan where our sensitive attribute A is binary gender. Demographic parity states that the probability of getting a loan should be the same for males ($A = 0$) and females ($A = 1$).

In converting this definition of demographic parity to the multi-armed bandit setting, we alter the definition to be that the probability of pulling an arm a does not change based on group membership P_j :

$$\Pr(\text{pull } a|a \in P_0) = \Pr(\text{pull } a|a \in P_1). \quad (13)$$

Continuing our running example, assume we are a loan agency. The loan agency receives 4 applications at every timestep t : an applicant from a young female, an applicant from a young male, an applicant from an older female, an applicant from an older male; we must choose one application to grant at each timestep. After granting a loan the loan agency receives a down payment on that loan as reward. This reward is then used to update the estimates of whether or not a “good” loan application was received for the pulled arm. Assume that the loan agency wants to act fairly using the binary sensitive attribute of gender. Then, the probability that the loan agency chooses a female applicant at timestep t should be the same as the probability of choosing a male applicant.

A naive algorithm to enforce this definition of fairness is defined in Algorithm 2. We first pick from the groups uniformly at random, and then apply a regular CMAB algorithm like TOPINTERVAL⁶ or

⁶TOPINTERVAL is a variant of LinUCB by Auer et al. [2].

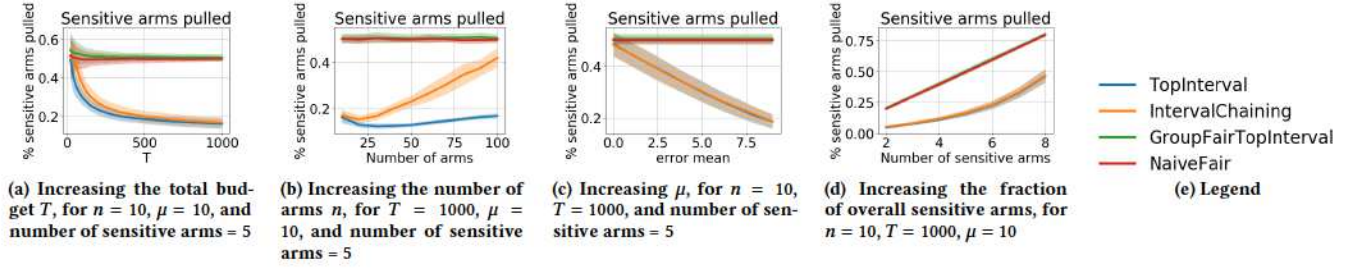


Figure 1: Percentage of total arm pulls that were pulled using sensitive arms.




Figure 2: Regret for synthetic experiments. The solid lines are regret given the rewards received from pulling the arms (including the group bias). The dashed lines is the true regret (without the group bias).

Algorithm 2 NAIVEGROUPFAIR**Require:** δ, P_1, P_2

- 1: **for** $t = 1 \dots T$ **do**
- 2: $P \leftarrow$ Randomly choose group P_1 or P_2 .
- 3: Pull arm in P based on TOPINTERVAL

CONTEXTUALTHOMPSONSAMPLING [1] to choose which arm to pull within the group. Using our running example, NAIVEGROUPFAIR would randomly pick between male or female, and then choose the best applicant between the younger and older pair.

5.2 Synthetic Experiments

In each synthetic experiment, we generate true coefficient vectors β_i by choosing coefficients uniformly at random for each arm i . Contexts at each timestep t are chosen randomly for each arm i . Bias coefficients ψ_1 are set uniformly at random with mean $\mu = 10$. Seeds are set at the beginning of each experiment to keep arms consistent between algorithms.

We run four different types of experiments:⁷ a) Varying the total budget for pulling arms (T) while setting number of arms $n = 10$, error mean $\mu = 10$, number of sensitive arms equal to 5, and context dimension $d = 2$ (Figures 2a and 1a). b) Varying the total number of arms n while setting total budget $T = 1000$, error mean $\mu = 10$, ratio of sensitive arms to 50%, and context dimension $d = 2$ (Figures 2b and 1b). c) Varying the error mean μ while setting total budget $T = 1000$, number of arms $n = 10$, number of sensitive arms equal to 5, and context dimension $d = 5$ (Figures 2c and 1c). d) Varying the number of sensitive arms while setting total budget $T = 1000$, number of arms $n = 10$, error mean $\mu = 10$, and context dimension $d = 2$ (Figures 2d and 1d).

⁷Additional experiments can be found in Appendix C.

The plots in Figure 1 show the percentage of times an algorithm pulled a sensitive arm over the full budget T . In order to be fair, the percentage of sensitive arms pulled should be proportional to the number of sensitive arms, i.e., when there are 5 sensitive arms out of 10 total, the percentage of sensitive arms pulled is roughly 50%. Figure 2 shows the perceived regret that includes bias ψ as solid lines, and real regret that corrects bias (see Equations 2 and 3) as dashed lines. Algorithms with low real regret are considered ‘good’.

Figure 1a shows that once exploration is over, GROUPFAIRTOPINTERVAL pulls sensitive arms roughly 50% of the time, matching the 50% of sensitive arms. Figure 2a shows that GROUPFAIRTOPINTERVAL performs comparably on real regret as TOPINTERVAL performs on biased regret. This means GROUPFAIRTOPINTERVAL should be used over TOPINTERVAL in contexts where bias is anticipated. NAIVEFAIR performs poorly in the context of societal bias.

Figure 1b illustrates that INTERVALCHAINING becomes more group fair as the number of arms increase. This is because many arms are chained together and therefore, arms are chosen uniformly at random. Figure 2b illustrates this random picking of arms as real regret and biased regret increases dramatically for INTERVALCHAINING.

As expected, Figure 1c illustrates that when the error mean μ is large, both INTERVALCHAINING and TOPINTERVAL choose fewer sensitive arms. This leads to a high real regret as shown in Figure 2c. Following Kleinberg et al. [21], Figure 2c also suggests that one cannot have both individual *and* group fairness in a scenario with high mean error. The randomness in NAIVEFAIR leads to a very high regret for both perceived regret and real regret.

Figure 1d demonstrates the fairness property of proportionality. The percentage of sensitive arms pulled by GROUPFAIRTOPINTERVAL matches the number of sensitive arms. As shown in Figure 2d,




Figure 3: Results of running contextual bandit algorithms on the family income and expenditure dataset (Figures 3a and 3b), as well as the COMPAS dataset (Figures 3c and 3d). Figures 3a and 3c show the percentage of pulls that were of sensitive arms. Figures 3b and 3d show the biased regret for each of the algorithms. Note that the “real” regret like that shown in the synthetic experiments cannot be calculated.

the number of sensitive arms does not affect the real regret of GROUPFAIRTOPINTERVAL.

5.3 Experiments on Real-World Data

After exploring GROUPFAIRTOPINTERVAL on synthetic data, we move on to using both the Philippines family income and expenditure dataset on Kaggle⁸ and the ProPublica COMPAS dataset.⁹ When one looks at the gender and age breakdown in the family income dataset, one can see that quite often female heads of households make more money than males in the Philippines. This is most likely due to the large number of Filipino women who work out of the country; it is estimated that up to 20% of the GDP of the Philippines is actually remittances from these overseas—primarily female—workers.¹⁰ In fact, almost 60% of overseas workers are women and 75% of these women are between the ages of 25 and 44.¹¹ In the COMPAS dataset ProPublica observed a societal bias over recidivism risk scores for African-Americans.¹²

Experimental Setup. Given the skew of high income coming from female head of households in the family income dataset, we treat the binary ‘Household Head Sex’ feature as the sensitive attribute. To create arms, we split up households based on ‘Household Head Age’ bucketed into the following five groups: (8, 27], (27, 45], (45, 63], (63, 81], and (81, 99]. We then have 10 different arms (for example, two arms would be Female head of household between 8 and 27, and Male head of household between 8 and 27).

Similarly, we treat African-American individuals from the COMPAS dataset as the sensitive attribute. We create arms by splitting up households based on the three age categories found in the data. We therefore have six different arms.

At each timestep t , we randomly select an individual from each arm. The context vector is the remaining features where any nominal features are transformed into integers. After an arm is pulled, a reward of the household income (for the family income dataset) or violent decile score (for COMPAS) is returned. We use these datasets for illustrative purposes.

Results. We see the same behavior of arm pulls in the real world data. Figures 3a and 3c show that after a period of exploration, the percentage of sensitive arms (male-grouped arms) pulled gets very close to 50%, matching the proportion of sensitive-grouped arms.

Figures 3b and 3d are perhaps more interesting. Since we cannot measure the “real” regret without the bias we assumed from the sensitive-grouped arms, we consider the gap between GROUPFAIRTOPINTERVAL and TOPINTERVAL as the price of fairness. The gap in regret is small compared to the increase in percentage of sensitive arms pulled. However, the gap in regret for NAIVEFAIR is large in comparison. This suggests that explicitly learning a societal bias term will help in biased settings with low price to perceived regret. Note that there is a difference between regret scales for the two different datasets. This is due to the family income and expenditure dataset reporting regret in income, while the COMPAS dataset reports regret in recidivism score.

6 GENERAL DISCUSSION & ETHICAL IMPLICATIONS OF THE WORK

This work was directly motivated by research into bias found in machine learning models. There have also been calls to action for more research to be done on bias mitigation in online learning settings [7, 11], specifically in multi-armed bandit settings [38] and related areas such as recommender systems [8, 43, 45]. Directly addressing these calls, in this work we propose a method of alleviating societal or measurement bias introduced into reward feedback. Using our CMAB model should help mitigate biased behaviors found in bandit systems currently in use [48].

Additionally, as noted by O’Neil [31], models can provide a biased feedback loop. We hope that by incorporating a societal bias term we can learn something about the bias that is being introduced. The coefficient vector will show which features are incorporating bias into the model. This allows users to address these features outside of the model and potentially find the sources of the societal bias. We do note that addressing societal bias and fixing the solution is a nontrivial task, the societal bias term provides the initial step of measurement.

On the other hand, as noted by Schumann et al. [38], Shneiderman [42], and many others, humans should still be active participants in decision making. If models such as our CMAB model are used to replace more and more human decision makers, this could

⁸<https://www.kaggle.com/grosvnpaul/family-income-and-expenditure>

⁹<https://www.kaggle.com/danofer/compass>

¹⁰<https://www.nationalgeographic.com/magazine/2018/12/filipino-workers-return-from-overseas-philippines-celebrates/>

¹¹<https://psa.gov.ph/content/2017-survey-overseas-filipinos-results-2017-survey-overseas-filipinos>

¹²<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

have unintended and potentially negative medium- and long-term side effects. All models should be monitored for biased feedback loops in the given contexts that they are being used [31].

Choosing a particular definition of fairness—conditioned on deciding that it is even appropriate to formally define a notion of fairness in the first place—is a morally-laden decision. We note that, as machine learning practitioners, in many societally-relevant applications it is paramount that we maintain an open dialogue with stakeholders. In this work, we analyze a sequential decision-making system under one particular definition, group fairness; it is certainly not the case that this is a one-size-fits-all solution that would be deployable without receiving input from that larger set of stakeholders. Indeed, recent research [36] shows that non-expert users may have vastly varying degrees of comprehension of different definitions of fairness, and that the degree of comprehension may be a function in part of education level and other features that may correlate with measures of marginalization; this hints that the consequences of incorporating fairness definitions into machine-learning-based systems may not be uniformly understood by participants, and indeed that those participants who may be impacted the most by that change could comprehend that potential impact the least. In an allocative system like the one we describe in the main paper, nuanced considerations must be considered.

Our new definitions of reward (Equation 2) and regret (Equation 3) for the MAB setting provide an opportunity to look at biased data in a new light. In many cases, ground truths provided during learning are noisy with respect to sensitive groups. Additionally, debiased ground truths may be very expensive to receive or may take a long time to acquire. For instance, if looking at loans, true rewards of repayment may take years to receive. Or, for example, in hiring—the true reward of hiring an individual may take over a year to estimate, while the initial estimate may be influenced by a hiring team’s unconscious bias over features such as ethnicity, gender, or orientation.

Our proposed algorithm, GROUPFAIRTOPINTERVAL, learns societal bias in the data while still being able to differentiate between individual arms. Previous solutions relied on setting ad-hoc thresholds, requiring some form of quota, or choosing groups uniformly at random. While it is true that GROUPFAIRTOPINTERVAL can easily be extended to a case where we know that the average of a group is a constant offset from the other group. That being said, defining such offsets raises a host of other ethical questions. For instance, in the US, the EEOC (Equal Employment Opportunity Commission) posited that the ratio of the most favored group compared to the least favored group must not be less than 0.8. Meanwhile, this type of comparison is forbidden in some countries [25]. In any case, these prior solutions either lead to high regret, or require a large amount of domain knowledge for the chosen application.

7 CONCLUSION & FUTURE RESEARCH

This paper explores group fairness in the contextual multi-armed bandit (MAB) setting. Our main contributions are: (1) we provide a new definition of reward and regret which captures societal bias; (2) we provide an algorithm that learns and corrects for that definition of societal bias; and (3) we empirically explore the effects different CMAB algorithms have in the setting of societal bias.

Future work could expand GROUPFAIRTOPINTERVAL to enforce individual fairness within groups. Intersectional group fairness is also important to look at in the MAB setting where more than one type of sensitive attribute needs to be protected. Additionally, other group fairness definitions such as Equalized Opportunity should be converted to the MAB setting [16]. Another interesting direction for future work is to mix ideas from the study of budget constrained bandits [14, 52] with our fairness definitions. We have also assumed individual arms have fixed group membership; generalizing to a setting where memberships in protected groups may change at every timestep t would fit more real world applications.

ACKNOWLEDGMENTS

Dickerson and Schumann were supported by NSF CAREER Award IIS-1846237, NSF Award CCF-1852352, NSF Award SMA-2039862, NIST MSE Award #20126334, DARPA GARD #HR00112020007, DARPA SI3-CMD #S4761, DoD WHS Award #HQ003420F0035, and ARPA-E DIFFERENTIATE Award #1257037. Mattei was supported by NSF Awards IIS-RI-2007955 and IIS-III-2107505. We thank Aviva Prins and Aravind Srinivasan for helpful comments on earlier versions of this paper, and the anonymous reviewers for helpful comments and constructive critiques.

REFERENCES

- [1] Shipra Agrawal and Navin Goyal. 2013. Thompson Sampling for Contextual Bandits with Linear Payoffs. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*. 127–135.
- [2] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning* 47, 2-3 (2002), 235–256.
- [3] Douglas Belkin. 2019. SAT to give students ‘adversity score’ to capture social and economic background. *The Wall Street Journal* (2019).
- [4] Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, A Mojsilović, et al. 2019. AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development* (2019).
- [5] Reuben Binns. 2020. On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT*)*. 514–524.
- [6] Miranda Bogen. 2019. All the Ways Hiring Algorithms Can Introduce Bias. <https://hbr.org/2019/05/all-the-ways-hiring-algorithms-can-introduce-bias> [2019-08-22].
- [7] M. Bogen and A. Rieke. 2018. *Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias*. Technical Report. Upturn.
- [8] Robin Burke, Amy Voda, Nicholas Mattei, and Nasim Sonboli. 2020. Algorithmic Fairness, Institutional Logics, and Social Choice. In *Harvard CRCS Workshop on AI for Social Good at 29th International Joint Conference on Artificial Intelligence (IJCAI 2020)*.
- [9] L. Elisa Celis, Sayash Kapoor, Farnood Salehi, and Nisheeth Vishnoi. 2019. Controlling Polarization in Personalization: An Algorithmic Framework. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*)*. 160–169.
- [10] Yifang Chen, Alex Cuellar, Haipeng Luo, Jignesh Modi, Heramb Nemlekar, and Stefanos Nikolaidis. 2020. Fair contextual multi-armed bandits: Theory and experiments. In *Conference on Uncertainty in Artificial Intelligence (UAI)*. PMLR, 181–190.
- [11] Alexandra Chouldechova and Aaron Roth. 2018. The Frontiers of Fairness in Machine Learning. *CoRR abs/1810.08810* (2018). arXiv:1810.08810 <http://arxiv.org/abs/1810.08810>
- [12] Houston Claire, Yifang Chen, Jignesh Modi, Malte Jung, and Stefanos Nikolaidis. 2019. Reinforcement Learning with Fairness Constraints for Resource Distribution in Human-Robot Teams. *arXiv preprint arXiv:1907.00313* (2019).
- [13] Sam Corbett-Davies and Sharad Goel. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023* (2018).
- [14] Wenkui Ding, Tao Qin, Xu-Dong Zhang, and Tie-Yan Liu. 2013. Multi-armed bandit with budget constraint and variable costs. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*.
- [15] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in*

- Theoretical Computer Science Conference*. ACM, 214–226.
- [16] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*. 3315–3323.
 - [17] Safwan Hossain, Evi Micha, and Nisarg Shah. 2021. Fair Algorithms for Multi-Agent Multi-Armed Bandits. In *Neural Information Processing Systems (NeurIPS)*.
 - [18] Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2016. Rawlsian fairness for machine learning. *arXiv preprint arXiv:1610.09559* (2016).
 - [19] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. 2016. Fairness in Learning: Classic and Contextual Bandits. In *Advances in Neural Information Processing Systems (NeurIPS)* 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). 325–333.
 - [20] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In *International Conference on Machine Learning*. 2569–2577.
 - [21] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. *Proceedings of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems* (2016).
 - [22] Chung-Ming Kuan. 2004. Classical least squares theory. http://homepage.ntu.edu.tw/~ckuan/pdf/et_ch3_Fall2009.pdf
 - [23] Tallai Andherbertrobbins Lai and Herbert Robbins. 1985. Asymptotically Efficient Adaptive Allocation Rules. *Advances in Applied Mathematics* 6, 1 (1985), 4–22.
 - [24] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web (WWW)*. ACM, 661–670.
 - [25] Robert Lieberman. 2001. A tale of two countries: the politics of color blindness in France and the United States. *French Politics, Culture & Society* 19, 3 (2001), 32–59.
 - [26] Michael L. Littman, Ifeoma Ajunwa, Guy Berger, Craig Boutilier, Morgan Currie, Finale Doshi-Velez, Gillian Hadfield, Michael C. Horowitz, Charles Isbell, Hiroaki Kitano, Karen Levy, Terah Lyons, Melanie Mitchell, Julie Shah, Steven Sloman, Shannon Vallor, and Toby Walsh. 2021. Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report. *Stanford University* (2021).
 - [27] Yang Liu, Goran Radanovic, Christos Dimitrakakis, Debmalaya Mandal, and David C Parkes. 2017. Calibrated fairness in bandits. In *Workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT-ML)*.
 - [28] Jérémie Mary, Romaric Gaudel, and Philippe Preux. 2015. Bandits and Recommender Systems. In *Machine Learning, Optimization, and Big Data*. 325–336.
 - [29] Kanishka Misra, Eric M Schwartz, and Jacob Abernethy. 2019. Dynamic Online Pricing with Incomplete Information Using Multiarmed Bandit Experiments. *Marketing Science* 38, 2 (2019), 226–252.
 - [30] Executive Office of the President, Cecilia Munoz, Domestic Policy Council Director, Megan (US Chief Technology Officer Smith (Office of Science, Technology Policy)), DJ (Deputy Chief Technology Officer for Data Policy, Chief Data Scientist Patil (Office of Science, and Technology Policy)). 2016. *Big data: A report on algorithmic systems, opportunity, and civil rights*. Executive Office of the President.
 - [31] Cathy O’Neil. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.
 - [32] Vishakha Patil, Ganesh Ghalme, Vineet Nair, and Y Narahari. 2019. Achieving fairness in the stochastic multi-armed bandit problem. *arXiv preprint arXiv:1907.10516* (2019).
 - [33] J. Rawls. 1971. *A Theory of justice*. Harvard University Press.
 - [34] Tom Ron, Omer Ben-Porat, and Uri Shalit. 2021. Corporate Social Responsibility via Multi-Armed Bandits. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FACT)*. 26–40.
 - [35] Cynthia Rudin. 2013. Predictive policing using machine learning to detect patterns of crime. *Wired Magazine* (2013).
 - [36] Debjani Saha, Candice Schumann, Duncan C. McElfresh, John P. Dickerson, Michelle L Mazurek, and Michael Carl Tschantz. 2020. Measuring Non-Expert Comprehension of Machine Learning Fairness Metrics. In *Proceedings of the International Conference on Machine Learning (ICML)*.
 - [37] Candice Schumann, Samsara N. Counts, Jeffrey S. Foster, and John P. Dickerson. 2019. The Diverse Cohort Selection Problem. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 601–609.
 - [38] Candice Schumann, Jeffrey S Foster, Nicholas Mattei, and John P Dickerson. 2020. We Need Fairness and Explainability in Algorithmic Hiring. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1716–1720.
 - [39] Candice Schumann, Zhi Lang, Jeffrey S. Foster, and John P. Dickerson. 2019. Making the Cut: A Bandit-based Approach to Tiered Interviewing. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*.
 - [40] Candice Schumann, Zhi Lang, Nicholas Mattei, and John P. Dickerson. 2019. Group Fairness in Bandit Arm Selection. In *NeurIPS Workshop on Machine Learning and Causal Inference for Improved Decision Making*.
 - [41] Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction* 36, 6 (2020), 495–504.
 - [42] Ben Shneiderman. 2020. Human-centered Artificial Intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction* 36, 6 (2020), 495–504.
 - [43] Ashudeep Singh, Yoni Halpern, Nithum Thain, Konstantina Christakopoulou, Ed H. Chi, Jilin Chen, and Alex Beutel. 2020. Building Healthy Recommendation Sequences for Everyone: A Safe Reinforcement Learning Approach. In *FAccT Rec Workshop on Responsible Recommendation at RecSys-20*.
 - [44] Adam Smith. 1776. *The Wealth of Nations*. New York: *The Modern Library* (1776).
 - [45] Nasim Sonboli, Robin Burke, Nicholas Mattei, Farzad Eskandarian, and Tian Gao. 2020. “And the Winner Is...”: Dynamic Lotteries for Multi-group Fairness-Aware Recommendation. *CoRR abs/2009.02590* (2020). [arXiv:2009.02590](https://arxiv.org/abs/2009.02590)
 - [46] Richard S. Sutton and Andrew Barto. 2017. *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
 - [47] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Commun. ACM* 56, 5 (2013), 44–54.
 - [48] Latanya Sweeney. 2013. Discrimination in online ad delivery. *arXiv preprint arXiv:1301.6822* (2013).
 - [49] Yanhan Savannah Tang, Alan Andrew Scheller-Wolf, and Sridhar R Tayur. 2021. Generalized Bandits with Learning and Queueing in Split Liver Transplantation. *Available at SSRN 3855206* (2021).
 - [50] Lequn Wang, Yiwei Bai, Wen Sun, and Thorsten Joachims. 2021. Fairness of Exposure in Stochastic Bandits. In *International Conference on Machine Learning (ICML)*.
 - [51] James Wexler, Mahima Pushkarna, Tolga Bolukbasi, Martin Wattenberg, Fernanda Viégas, and Jimbo Wilson. 2019. The What-If Tool: Interactive probing of machine learning models. *IEEE transactions on visualization and computer graphics* (2019).
 - [52] Huasen Wu, Rayadurgam Srikant, Xin Liu, and Chong Jiang. 2015. Algorithms with logarithmic or sublinear regret for constrained contextual bandits. In *Advances in Neural Information Processing Systems*. 433–441.
 - [53] H Peyton Young. 1995. *Equity: in theory and practice*. Princeton University Press.