Online Combinatorial Optimization with Group Fairness Constraints

Negin Golrezaei[1] , **Rad Niazadeh**[2] , **Kumar Kshitij Patel**[3] and **Fransisca Susan**[1]
[1]Sloan School of Management, MIT
[2]Booth School of Business, University of Chicago
[3]Toyota Technological Institute at Chicago (TTIC)
golrezae@mit.edu, rad.niazadeh@chicagobooth.edu, kkpatel@ttic.edu, fsusan@mit.edu

Abstract

As digital marketplaces and services continue to expand, it is crucial to maintain a safe and fair environment for all users. This requires implementing fairness constraints into the sequential decision-making processes of these platforms to ensure equal treatment. However, this can be challenging as these processes often need to solve NP-complete problems with exponentially large decision spaces at each time step. To overcome this, we propose a general framework incorporating robustness and fairness into NP-complete problems, such as optimizing product ranking and maximizing submodular functions. Our framework casts the problem as a max-min game between a primal player aiming to maximize the platform's objective and a dual player in charge of group fairness constraints. We show that one can trace the entire Pareto fairness curve by changing the thresholds on the fairness constraints. We provide theoretical guarantees for our method and empirically evaluate it, demonstrating its effectiveness.

1 Introduction

Online combinatorial optimization problems, such as optimizing product rankings, assortment planning, or supply-demand matching for a sequence of arriving input instances, are prevalent in digital marketplaces. From a pure *utilitarian* perspective, an economically justified goal in several of these applications is to make combinatorial decisions across time steps that maximize the expected value of a global objective function representing the *market share*, i.e., the total/average user engagement for the entire population. However, these marketplaces typically have users with heterogeneous preferences based on demographics such as race, gender, and age [Hitsch *et al.*, 2010]. Given these demographic groups, a more *egalitarian* decision maker might want to maximize the expected average market share while ensuring each group's share of this quantity is *"large enough"* based on some predetermined threshold, thus fostering inclusivity and fairness for all minority groups. This practice is not only ethically important but also beneficial for business, as it attracts new customers, retains existing ones, and contributes to a healthy ecosystem by fostering a positive reputation.

While personalization is effective in achieving the above goal like, say, by displaying a unique assortment of products to each demographic group, the use of sensitive information for personalization can be restricted due to legal issues and privacy concerns that digital marketplaces face, which in turn affects the effectiveness and scope of personalized experiences [California Legislative Service, 2020]. Also, in some contexts, treating seemingly similar individuals differently can potentially be against the equal-protection doctrine and perceived as unlawful. This has motivated a concerted effort to introduce the notion of demographic group fairness guarantees into optimization problems by various mathematical semantics. For example, there is a growing literature on formulating the problem as a group distributional robust optimization (DRO) [Ben-Tal *et al.*, 2013; Duchi *et al.*, 2021; Sagawa *et al.*, 2019; Soma *et al.*, 2022], or incorporating parity among demographic groups by adding group fairness constraints [Kleinberg *et al.*, 2018; Chouldechova and Roth, 2018; Chen *et al.*, 2022; Rahmattalabi *et al.*, 2019].

However, most existing approaches are limited to continuous and offline decision-making settings where the decision-maker can efficiently solve the (unfair) offline optimization problem. Hence, these approaches do not provide general solutions for online combinatorial optimization problems with *discrete* and *exponential size* decision spaces at each time step, where the decision-maker should *learn* and *compute* through explorations/exploitation. To complicate things further, the offline version of many of these combinatorial optimization problems, e.g., assortment planning under a large class of consumer choice models or maximizing market-share through sub-modular function maximization, are NP-Hard [Goyal *et al.*, 2016; Nemhauser *et al.*, 1978; Asadpour *et al.*, 2022; Niazadeh *et al.*, 2021] and thus any robust/fair approach must employ existing approximation algorithms for these problems. We bridge this gap in this paper by providing a general learning and computing framework for online combinatorial optimization with group fairness constraints, where our algorithms can employ (an online learning variant of) algorithms for the unfair version of the problem in a black-box oracle fashion. We now describe our problem.

The full version of our paper, including all proofs, can be accessed at https://ssrn.com/abstract=4824251

1.1 Problem Formulation and Examples

Let $\mathcal{F}$ be a conic-closed[1] function class defined over (an exponentially large) domain $\mathcal{D}$ of actions. At every time step $t \in [T]$, there are L functions $\{f_t^i \in \mathcal{F}\}_{i\in[L]}$ such that $f_t^i : \mathcal{D} \to [0, f_{\max}]$ for all $i \in [L]$, where L is the number of demographic groups and f_t^i is the reward function of group i at time t. Furthermore, let $\{\tau^i\}_{i\in[L]}$ be a set of non-negative fairness thresholds in $[0, f_{\max}]$, and $\mathcal{C} \subseteq \mathcal{D}$ be a constrained domain space. A learner, denoted by ALG, is aware of these fairness thresholds, and at time step $t \in [T]$, ALG picks an action $x_t \in \mathcal{C}$. Then,

1. either the L functions $\{f_t^i \in \mathcal{F}\}_{i\in[L]}$ (in the **full-information setting**);
2. or just the rewards $\{f_t^i(x_t) \in \mathcal{F}\}_{i\in[L]}$ (in the **bandit-information setting**),

are revealed to the learner ALG, and it receives a reward $\sum_{i\in[L]} f_t^i(x_t)$. The action x_t is chosen based on the knowledge of all the functions up to time $t-1$. In contrast, the functions $\{f_t^i\}_{t\in[L]}$ can depend on the knowledge of the previous functions $\{f_k^i\}_{i\in[L],k\in[t-1]}$ and actions $\{x_k\}_{k\in[t-1]}$, s.t.[2],

$$\mathbb{E}[f_t^i(x)|\mathcal{H}_t, x] = f^i(x),$$

where we define a sequence of filtrations for $t \in [T+1]$ as,

$$\mathcal{H}_t := \sigma\left(\{f_k^i\}_{i\in[L],k\in[t-1]}, \{x_k\}_{k\in[t-1]}\right).$$

In general ALG might select a distribution of actions $P \in \Delta(\mathcal{C})$, where $\Delta(\mathcal{C})$ is the set of all distributions over $\mathcal{C}$. Then, the online maximization benchmark with group fairness constraints, ONLINE-OPT, which we compete against, is,

$$\begin{aligned}
&\max_{P\in\Delta(C)} \frac{1}{T}\sum_{t,i} \mathbb{E}_{x_t\sim P}\left[f_t^i(x_t)|\mathcal{H}_t\right] \\
&\quad \text{s.t.} \quad \frac{1}{T}\sum_{t} \mathbb{E}_{x_t\sim P}\left[f_t^i(x_t)|\mathcal{H}_t\right] \ge \tau^i,\ \forall i \in [L], \\
&= \max_{P\in\Delta(C)} \sum_{i} \mathbb{E}_{x\sim P}\left[f^i(x)\right] \\
&\quad \text{s.t.} \quad \mathbb{E}_{x\sim P}\left[f^i(x)\right] \ge \tau^i,\ \forall i \in [L].
\end{aligned} \tag{1}$$

Note that the online optimization problem (1) [3] considers *ex-ante fairness* constraints because we only require the expected value $\mathbb{E}_{x\sim P}\left[f^i(x)\right]$ to be greater than τ^i. When $\tau_i = 0$ for all $i \in [L]$, we are only concerned about maximizing the total market share without any fairness considerations,

$$\max_{P\in\Delta(C)} \sum_{i} \mathbb{E}_{x\sim P}\left[f^i(x)\right]. \tag{2}$$

[1] If $f^1, f^2 \in \mathcal{F}$, then $c_1 f^1 + c_2 f^2 \in \mathcal{F}$ for any $c_1, c_2 \ge 0$.

[2] This stochastic adversarial/martingale setup is common in reinforcement learning [Ribeiro, 2002; Besbes *et al.*, 2014; Jafarnia-Jahromi *et al.*, 2021]. The i.i.d. setting where for all $i \in [L]$ the function f_t^i for group i is picked independently and identically from distribution $\mathcal{D}_i \in \Delta(\mathcal{F})$ which is fixed in advance, is a special case.

[3] As is usual in optimization literature, we will refer to the optimization problem and its optimal value interchangeably.

On the other hand, is the following "robust" problem,

$$\max_{P\in\Delta(C)} \min_{i\in[L]} \mathbb{E}_{x\sim P}\left[f^i(x)\right], \tag{3}$$

which is known as the group-DRO problem and has been studied extensively [Ben-Tal *et al.*, 2013; Duchi *et al.*, 2021; Sagawa *et al.*, 2019] in machine learning literature. Suppose we have access to the value $\tau^\star$ of the robust problem (3). In that case, we can recover group-DRO by solving problem (1) with fairness constraints $\tau_i = \tau^\star$ for all $i \in [L]$. More generally, problem (1) interpolates between the utilitarian objective when $\tau_i = 0$ for all $i \in [L]$ as we are only concerned about maximizing the total market share and the *egalitarian* objective when $\tau_i = \tau^\star$ for all $i \in [L]$, where we are maximizing the market share of the worst group. Notably, there are many ways to interpolate between these objectives. With a symmetric threshold $\tau^i = \tau$ that varies in the interval $[0, \tau^\star]$, this approach recovers the *Pareto frontier of fairness-efficiency*, which is a canonical object of study at the intersection of fairness and decision making, for e.g., see [Bertsimas *et al.*, 2011; Bertsimas *et al.*, 2012]. Using constraints to trace the utility-fairness curve is related to scalarization and ϵ-constraint approaches [Miettinen, 1999] in vector optimization.

Remark 1 (Offline Problem)**.** *A very special case of our problem is the following offline problem (see the full version of our paper for more discussion):*

$$\begin{aligned}
&\max_{P\in\Delta(C)} \frac{1}{L}\sum_{i\in[L]} \mathbb{E}_{x\sim P}\left[f^i(x)\right] \\
&\quad \textit{s.t.} \quad \mathbb{E}_{x\sim P}\left[f^i(x)\right] \ge \tau^i, \qquad \forall i \in [L].
\end{aligned} \tag{4}$$

Compared to the online problem (1)*, the functions* $\{f^i\}^{i\in[L]}$ *are known in advance for the offline problem. We will refer to the optimal value of problem* (4) *by* OFFLINE-OPT.

Running Examples. We will consider two running examples to illustrate the generality of our framework.

Example 1 (**Online Shortest Path Problem**)**.** *Consider routing traffic through a city in real-time using a navigation service. Different neighborhoods (some of which might be marginalized) would be affected differently [Park and Kwan, 2020] due to factors such as exposure to air pollution, traffic congestion, road accidents, economic opportunities, etc. This problem can be reduced to an "online shortest path" problem—one of the core combinatorial prediction problems in online learning. Consider a graph* $G = (V, E)$*, with vertices* V *connected by a set of directed edges* $E \subseteq V^2$*. Let* $u \in V$ *and* $v \in V$ *be the source and sink vertex, respectively. Let* $\mathcal{C}$ *be the set of all cycle-free paths from* u *to* v*, where each edge only appears once. Note that* $\mathcal{C}$ *is potentially exponential in* $|E|$*. At time* $t \in [T]$*, for* $i \in [L]$*,* $f_t^i : E \to [0, 1]$*, and* $f_t^i(e)$ *measures the reward for picking an edge* $e \in E$ *for group* i*. Then, for a given path* $x_t \in \mathcal{C}$*, the total reward for group* i *is measured as* $f_t^i(x_t) = \sum_{e\in x_t} f_t^i(e)$ *which is a noisy measurement of the mean reward* $f^i(x_t)$*. This repeats each day, and at the end of the day* t *the reward functions are revealed for all the edges,* $\{f_t^i(e) : e \in E, i \in [L]\}$*.*

Example 2 (**Online Sub-modular Maximization**). *The notion of sub-modularity is commonly used to describe the diminishing return property in discrete and continuous domains. Many optimization problems that arise in the real world such as viral marketing in a social network [Kempe* et al.*, 2003], ranking an assortment of products [Asadpour* et al.*, 2022], welfare maximization [Dobzinski and Schapira, 2006] and speeding up satisfiability solvers [Streeter and Golovin, 2008] can be expressed as maximizing a sub-modular function under constraints. In our setting, assume that $f_t^i : \mathcal{C} \to [0,1]$ is a set function defined on some constrained set of sets $\mathcal{C} \subseteq 2^{[n]}$, such that, for all $S, T \in [n]$, $f_t^i(S \cup T) + f_t^i(S \cap T) \leq f_t^i(S) + f_t^i(T)$, i.e., f_t^i is sub-modular. A common set constraint is the cardinality constraint which only allows sets of size at most $k \leq n$. It is well-known that sub-modular function maximization (SMM) with cardinality constraints is an NP-hard problem that admits the classic $\left(1 - \frac{1}{e}\right)$-approximation greedy algorithm [Nemhauser* et al.*, 1978]. Many applications of SMM have fairness considerations in the presence of multiple groups. For instance, recommendations on a digital marketplace should not be unreasonably worse for any gender.*

Our proposed framework in this paper can accommodate NP-hard problems such as sub-modular function maximization. Naturally, we can only hope to attain the value ONLINE-OPT of problem (1) up to a multiplicative factor $\gamma \in (0, 1)$. We will choose γ as the best approximation factor for the offline version of the problem of maximizing the individual function for each group. Thus it will arise from the inherent hardness of the problem. For instance, for SMM with cardinality constraints, we will pick $\gamma = \left(1 - \frac{1}{e}\right)$. Furthermore, since the optimization problem (1) is equivalent to solving a series of feasibility problems, we can only satisfy the fairness constraints up to this multiplicative factor γ. In particular, in section 3, we propose Algorithm 1 to minimize the following approximate regret metric often used in the literature [Niazadeh *et al.*, 2021; Kakade *et al.*, 2007; Dudik *et al.*, 2020],

$$\gamma \cdot \text{ONLINE-OPT} - \frac{1}{T} \sum_{i \in [L], t \in [T]} \mathbb{E}_{x_t, f_t^i} \left[f_t^i(x_t) | \mathcal{H}_t \right], \quad (5)$$

while also γ-approximately satisfying the fairness constraints for all $i \in [L]$,

$$\gamma \cdot \tau_i \approx \frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{x_t, f_t^i} \left[f_t^i(x_t) | \mathcal{H}_t \right]. \quad (6)$$

We will refer to the regret in (5) as **γ-average regret**.

1.2 Our Contributions

The contributions of our paper are as follows,

- We provide a general framework to incorporate group fairness constraints to any online combinatorial optimization problem that we know how to solve with a single group. Our approach is easy to implement and accommodates both full and bandit information feedback.
- We provide theoretical guarantees for both the regret and fairness of our approach (see algorithm 1) with black-box access to an online optimization oracle (see definition 1) for the combinatorial problem with a single group. In particular, we show vanishing γ-average regret while being approximately fair (violating the fairness constraint by, at most, some small error $\delta > 0$) for large T. Thus, we extend prior work on primal-dual algorithms for online optimization [Shalev-Shwartz and Singer, 2007] to our setting with group fairness.
- We present an empirical evaluation of our algorithm using parameters derived from the widely used MovieLens ratings dataset [Harper and Konstan, 2015]. Our results demonstrate that our online algorithm achieves sub-linear regret when comparing it to the approximate optimal benchmark. Additionally, we investigate the effects of fairness constraints on the total expected market share, providing insights into the trade-offs between group fairness and assortment optimization.

2 Related Work

Recently, machine learning algorithms have increasingly entered various fields, sparking a growing interest in algorithmic fairness in optimization, including combinatorial optimization. Researchers have studied a wide range of problems such as matching [Chierichetti *et al.*, 2019], bandit optimization [Liu *et al.*, 2022; Joseph *et al.*, 2016], resource allocation [Bertsimas *et al.*, 2012; Manshadi *et al.*, 2021], assortment planning [Singh and Joachims, 2018; Biega *et al.*, 2018; Chen *et al.*, 2022], and reserve pricing in search ad markets [Deng *et al.*, 2022]. Algorithmic fairness generally falls into three categories: 1) individual fairness, where algorithms make comparable predictions for similar individuals [Dwork *et al.*, 2012; Loi *et al.*, 2019; Chen *et al.*, 2022], 2) group fairness, which ensures equal treatment of distinct groups, often in resource use or performance in combinatorial optimization [Singh and Joachims, 2018; Balseiro *et al.*, 2021], and 3) subgroup fairness, which combines elements of both [Kearns *et al.*, 2018; Kearns *et al.*, 2019].

For fairness in assortment planning, previous research mainly targets individual fairness for sellers on online platforms. [Chen *et al.*, 2022] use pairwise fairness to ensure that items of similar quality get similar visibility. They solve their constrained assortment optimization problem by framing it as a linear program and using the Ellipsoid method to find a nearly optimal solution. Fairness in ranking represents another related field. Essentially, assortment planning acts as a specialized case of ranking: items in the top K positions get full visibility, while the rest remain invisible. [Garcia-Soriano and Bonchi, 2021] work on minimizing individual unfairness and imposing group fairness constraints in offline ranking. They propose a polynomial-time algorithm to solve the problem. [Celis *et al.*, 2018] focus on group fairness across demographics, formulate their problem as a linear program, prove its computational hardness, and introduce an LP-based approximation algorithm with theoretical guarantees. However, they do not explore online scenarios.

Our paper differs from previous literature in several ways.

We focus on group fairness in general combinatorial optimization problems, unlike most previous works that center on individual fairness for specific problems. We also tackle optimization problems that might be unsolvable as a whole, such as NP-hard problems, and that do not have objectives that are linear in actions. This approach enables us to address various problems relevant to online platforms, from sub-modular function maximization to assortment planning. Our offline results most closely relate to concurrent work by [Tang and Yuan, 2023], who employ the Ellipsoid method to solve offline combinatorial optimization problems with fairness constraints. However, our work has two main advantages. First, we use a simple primal-dual framework that is easy to implement, unlike the numerically unstable and more complex Ellipsoid method. Second, we can extend our framework easily to online settings that are more pertinent to online markets where decision-making usually occurs sequentially.

3 Our Algorithm and Theoretical Guarantees

In this section, we present our algorithm 1 and provide its theoretical guarantees. Recall that we focus on functions that are difficult to optimize, where even the individual problems $\max_{x\in\mathcal{C}} f_t^i(x)$ for each $i \in [L], t \in [T]$ are NP-hard. To solve the online variants of such problems with group fairness constraints, we need access to the online algorithm used to solve the problem with a single group. In particular, we assume the existence of the following no-regret online oracle $\mathcal{Z}$.

Definition 1 (γ-no-regret Online Oracle)**.** *A γ-no-regret online oracle, denoted as $\mathcal{Z}$, is an algorithm that produces a sequence of probability distributions $\{P_t \in \Delta(\mathcal{C})\}_{t\in[T]}$ such that the following γ-average regret is diminishing over time,*

$$\max_{x\in\mathcal{C}} \frac{\gamma}{T}\sum_{t\in[T]} f_t(x) - \frac{1}{T}\sum_{t\in[T]} \mathbb{E}_{x_t\sim P_t}\left[f_t(x_t)\right] = \frac{o(T)}{T}. \quad (7)$$

In the full-information setting, at time t, the oracle knows its past actions $x_1, \dots, x_{t-1}$ and the past functions $f_1, \dots, f_{t-1} \in \mathcal{F}$, while in the bandit setting, it is only aware of its past actions and the past function values $f_1(x_1), \dots, f_{t-1}(x_{t-1})$.

We can obtain a γ-no-regret online oracle for several combinatorial problems, including our two running examples.

Remark 2 (Oracle for Online Shortest Path Problem)**.** *Online shortest path problem can be solved exactly, with $\gamma = 1$ to get an $\mathcal{O}\left(1/\sqrt{T}\right)$ averaged regret. One standard approach (see [Dekel, 2012] for details) first obtains a convex relaxation of the problem at each time-step to select an action y_t from the convex hull of all the paths in $\mathcal{C}$. This reduces the problem into an online linear optimization problem with full-information feedback, which can be optimized with an algorithm such as projected online gradient descent to get $\mathcal{O}(\sqrt{T})$ regret. This is followed by a randomization procedure that samples a path x_t at each time step using y_t.*

Remark 3 (Oracle for Online Sub-modular Maximization)**.** *There are several online approximation oracles known for different variants of SMM. [Niazadeh* et al.*, 2021] provide one instance of such an algorithm that achieves $O(1/\sqrt{T})$ γ-average regret in the full-information and $O(1/T^{1/3})$ γ-average regret in the bandit-information setting for various classes of sub-modular functions, including monotone sub-modular functions with cardinality constraints ($\gamma = 1 - 1/e$) and non-monotone discrete sub-modular functions ($\gamma = 1/2$). Other works have also presented no-regret oracles for specific sub-modular problems. [Streeter and Golovin, 2008] proposed an oracle that achieves $O(1/\sqrt{T})$ $(1-1/e)$-average regret for monotone sub-modular functions with cardinality constraints, based on the randomized weighted majority algorithm, [Chen* et al.*, 2018] introduced a Frank-Wolfe based online oracle that achieves $O(1/\sqrt{T})$ $(1-1/e)$-average regret for monotone continuous sub-modular maximization in the downward closed convex set, and [Roughgarden and Wang, 2018] presented a potential function based oracle for non-monotone set sub-modular maximization. For a survey, see the paper by [Niazadeh* et al.*, 2021].*

Algorithm 1 Robust Online Algorithm

Input: A γ-no-regret online approximation oracle $\mathcal{Z}$ with action space $\Delta(\mathcal{C})$ (such as Blackwell algorithms in [Niazadeh *et al.*, 2021]), a no-regret online linear optimization algorithm $\mathcal{A}$ with action space $[0, L/\delta]^L$ (such as projected online gradient descent with optimal step sizes), the fairness violation threshold δ, and the total period T.

Output: An (random) action for each period.

1: Initialize $\boldsymbol{\alpha}_1 \leftarrow (L/\delta, L/\delta, \dots, L/\delta)$.

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

3: Max player calls the oracle $\mathcal{Z}$ to get $P_t \in \Delta(\mathcal{C})$.

4: Max player plays $x_t \sim P_t$.

5: Max player receives one of the following feedbacks:

Full-information: functions $\{f_t^i\}_{i\in[L]}$.

Bandit: function values $\{f_t^i(x_t)\}_{i\in[L]}$.

6: The min player gets a loss vector

$$\ell_t \leftarrow \begin{bmatrix} f_t^1(x_t) - \gamma\tau^1 \\ \vdots \\ f_t^i(x_t) - \gamma\tau^i \\ \vdots \\ f_t^L(x_t) - \gamma\tau^L \end{bmatrix}$$

.

7: The min player updates $\boldsymbol{\alpha}_{t+1} \leftarrow \mathcal{A}\left(\boldsymbol{\alpha}_t, \{\ell_k\}_{1\le k\le t}\right)$.

8: The oracle $\mathcal{Z}$ receives one of the following feedbacks:

Full-information: function

$$g_t := \sum_{i\in[L]} (1 + \alpha_t^i) f_t^i.$$

Bandit: function value

$$g_t(x_t) = (1 + \alpha_t^i) f_t^i(x_t).$$

9: **end for**

We will also need an online linear optimization oracle.

Definition 2 (No-regret OLO Oracle). *Let $\mathcal{X}, \mathcal{L} \subset \mathbb{R}^d$ be two compact spaces. Assume at each time step $j \in [N]$ we see an arbitrary loss function $l_j \in \mathcal{L} \subset \mathbb{R}^d$, and an algorithm $\mathcal{A} : (\mathcal{X} \times \mathcal{L})^\star \to \mathcal{X}$ outputs an action $x_j \in \mathcal{X}$ based on the history at time j. Then $\mathcal{A}$ is a no-regret online linear optimization (OLO) if it has a zero average regret $R^{\mathcal{A}}(N)$ in the limit $N \to \infty$, where*

$$R^{\mathcal{A}}(N) := \frac{1}{N} \sum_{j \in [N]} \langle \ell_j, x_j \rangle - \inf_{x \in \mathcal{X}} \frac{1}{N} \sum_{j \in [N]} \langle \ell_j, x \rangle . \quad (8)$$

Remark 4. *A few examples of no-regret online linear optimization oracles are projected Online Gradient Descent (OGD) [Zinkevich, 2003], Follow the Regularized Leader (FTRL), and Follow the Perturbed Leader (FTPL) [Kalai and Vempala, 2005] algorithms. For a survey of these algorithms and their relative comparisons, see [Hazan, 2016].*

Our complete online algorithm using the above two oracles is presented in Algorithm 1. Our algorithm is primal-dual-based and simulates the problem as an online max-min game: the min player $\mathcal{A}$ runs a no-regret online minimization algorithm, and the max player $\mathcal{Z}$ runs an approximate online algorithm on a conic combination of the functions at any given time $t \in [T]$. For the min player, who is ignorant of the combinatorial structure of the problem, there is no difference between full and bandit information feedback because it only requires the values of functions for each group at each round and each group's fairness threshold. Algorithm 1 satisfies the bounds in the following Theorem (c.f. full paper for proof).

Theorem 1. *Suppose that the online maximization problem with ex-ante fairness constraints, defined by $(\{f_t^i\}_{i\in[L],t\in[T]}, \{\tau_i\}_{i\in[L]})$ as shown in Equation 1 is feasible. Let $\delta > 0$ be some fairness violation threshold, and let γ be the approximation constant. Suppose that we have a γ-no-regret online approximation oracle $\mathcal{Z}$ with a γ-average regret of $R^{\mathcal{Z}}(\cdot)$ and an OLO algorithm $\mathcal{A}$ with an average regret of $R^{\mathcal{A}}(\cdot)$. Then, Algorithm 1 outputs a (random) sequence of actions $\{x_t\}_{t\in[T]}$ that satisfies the following:*

- *For each $i \in [L]$,*
$$\frac{1}{T} \sum_{t\in[T]} \mathbb{E}_{x_t \sim P_t, f_t^i} \left[f^i(x_t) | \mathcal{H}_t \right] \geq \gamma \cdot \tau^i - \delta f_{\max} - \frac{\delta}{L} R^{\mathcal{A}}(T) - \frac{\delta}{L} \cdot R^{\mathcal{Z}}(T),$$
meaning that the sequence is γ-approximately fair, and the violation converges to δf_{max} as $T \to \infty$.
- *The total reward function evaluated on the sequence is a γ-approximation to the optimal value in hindsight, $\gamma \cdot$ ONLINE-OPT, with the following regret*
$$\frac{\gamma}{T} \cdot \text{ONLINE-OPT} - \frac{1}{T} \sum_{i\in[L],t\in[T]} \mathbb{E}_{x_t \sim P_t, f_t^i} \left[f^i(x_t) | \mathcal{H}_t \right] \leq R^{\mathcal{A}}(T) + R^{\mathcal{Z}}(T).$$

Here, the functions $\{f_t^i\}_{i\in[L],t\in[T]} \in \mathcal{F}$ satisfy $\mathbb{E}_{f_t^i}[f_t^i(\cdot)|\mathcal{H}_t] = f^i(\cdot)$ where $\mathcal{H}_t = \sigma\Big(\{f_k^i\}_{i\in[L],k\in[t-1]}, \{x_k\}_{k\in[t-1]}\Big)$.

Proof sketch. We rely on the expected average Lagrangian dual quantity, defined as

$$\frac{1}{T} \sum_{i\in[L],t\in[T]} \Big((1+\alpha_t^i) \mathbb{E}_{x_t \sim P_t, f_t^i} \left[f_t^i(x_t) | \mathcal{H}_t \right] - \gamma \alpha_t^i \tau^i \Big),$$

to establish an upper bound on $\gamma \cdot$ ONLINE-OPT and a lower bound on the algorithm's total reward, $\sum_{i\in[L],t\in[T]} \mathbb{E}_{x_t \sim P_t, f_t^i} \left[f_t^i(x_t) | \mathcal{H}_t \right]$. Specifically, we first upper-bound the expected average Lagrangian dual quantity by taking the minimum overall expected average Lagrangian duals evaluated at $\{x_t\}_{t\in[T]}$, where the minimum is taken over $\boldsymbol{\alpha} \in [0, L/\delta]^L$, along with using the definition of the no-regret OLO oracle $\mathcal{A}$. This step gives the following bound,

$$\frac{1}{T} \sum_{i\in[L],t\in[T]} \Big((1+\alpha_t^i) \mathbb{E}_{x_t \sim P_t, f_t^i} \left[f_t^i(x_t) | \mathcal{H}_t \right] - \gamma \alpha_t^i \tau^i \Big) \leq \inf_{\boldsymbol{\alpha} \in [0,L/\delta]^L} \frac{1}{T} \sum_{i\in[L]} \sum_{t\in[T]} \left((1+\alpha^i) \mathbb{E}_{x_t \sim P_t} \left[f^i(x_t) \right] - \gamma \alpha^i \tau^i \right) + R^{\mathcal{A}}(T). \quad (9)$$

Next, we establish a lower bound on the expected average dual by leveraging the γ-approximation of ONLINE-OPT by $\mathcal{Z}$ and the definition of the functions g_t.

$$\frac{1}{T} \sum_{i\in[L],t\in[T]} \Big((1+\alpha_t^i) \mathbb{E}_{x_t \sim P_t, f_t^i} \left[f_t^i(x_t) | \mathcal{H}_t \right] - \gamma \alpha_t^i \tau^i \Big) \geq \gamma \cdot \text{ONLINE-OPT} - R^{\mathcal{Z}}(T) \quad (10)$$

Combining the upper bound (9) and the lower bound (10), followed by some algebraic manipulations, directly gives the fairness bound. Additionally, using the realization that the first quantity on the right of (9), i.e., the minimum overall expected average Lagrangian duals evaluated on $\{x_t\}_{t\in[T]}$ is upper bounded by the algorithm's average reward finishes the proof for the γ-averaged regret bound. □

From the statement of Theorem 1, it might seem like δ can be set to 0 to avoid the slack in the fairness constraint completely. However, there is a hiding delta in $\mathcal{R}^{\mathcal{A}}(T)$ because the action space of the oracle $\mathcal{A}$ is restricted to $[0, L/\delta]^L$ in Algorithm 1. For instance, if we pick $\mathcal{A}$ to be the projected online gradient descent algorithm, then with optimal step size, $R^{\mathcal{A}}(T) = \mathcal{O}\left(\frac{f_{max} L^2}{\delta\sqrt{T}}\right)$ (c.f., [Hazan, 2016] and the discussion in our full version). This implies that decreasing δ improves the slackness in the fairness constraint but worsens the upper bound for γ-averaged regret. Fixing $\mathcal{A}$ as projected online gradient descent with optimal step sizes (c.f., discussion in our full version), we recover the final guarantees for our two running examples.

Remark 5 (Online Shortest Paths). *For our first running example using $\mathcal{Z}$ as the oracle described in remark 2, the above theorem implies an average regret guarantee of $\mathcal{O}(1/\sqrt{T})$ as both $R^{\mathcal{A}}(T)$ and $R^{\mathcal{Z}}(T)$ are $\mathcal{O}(1/\sqrt{T})$.*

Remark 6 (Sub-modular Maximization). *For our second running example by picking $\mathcal{Z}$ as the Blackwell*

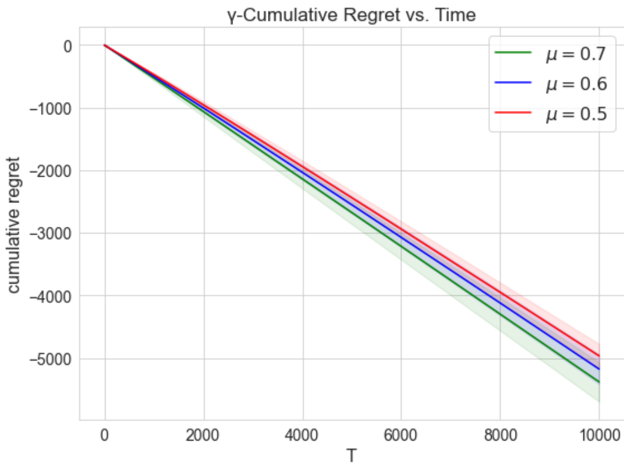

Figure 1: Cumulative $(1 - 1/e)$-regret over time of Algorithm 1 for various fairness thresholds μ (same as τ in the main text).

approachability-based algorithm by [Niazadeh et al.*, 2021], we can get $R^{\mathcal{Z}}(T) = O(1/\sqrt{T})$, implying an overall γ averaged regret of $\mathcal{O}(1/\sqrt{T})$ (γ factor depends on the additional structure in SMM) in the full-information setting. Similarly, in the bandit information setting the algorithm by [Niazadeh* et al.*, 2021] implies $R^{\mathcal{Z}}(T) = O(1/T^{1/3})$, implying an overall regret of $\mathcal{O}(1/T^{1/3})$.*

4 Experiments

In this section, we present a case study using the MovieLens 1M dataset due to [Harper and Konstan, 2015] to evaluate the effectiveness of our offline (c.f., our full version for more details) and online algorithms for movie recommendation. Specifically, we simulate a movie recommendation platform that aims to offer an assortment of movies tailored to users' diverse preferences. We partition the users into demographic groups based on gender and seek to maximize the overall market share while ensuring that the market share of each group meets a predetermined threshold to satisfy the group fairness constraint. Here, the market share of an assortment for a group represents the probability that a user in the respective group selects at least one movie from the assortment. Our experiments demonstrate that our online algorithm produces a sequence of sets with diminishing γ regret while meeting the fairness constraint. Furthermore, we examine the impact of various thresholds on the algorithm's performance.

Diminishing regret of the online algorithm. Figure 1 displays the cumulative regret of Algorithm 1 over time. To compute the regret of our algorithms, we use $(1 - 1/e)$ of the optimal market share as the regret benchmark, where $(1 - 1/e)$ is the γ factor for the offline greedy algorithm. The optimal market share is the highest possible among all assortments of size 5, assuming that we know all the parameters (weights and probability of each mixture) beforehand, which we obtain by enumerating all possible assortments. In the experiment, we consider three sets of thresholds for the average market share for each group: $(0.5, 0.5)$, $(0.6, 0.6)$, and $(0.7, 0.7)$, where all of these thresholds are feasible. We set the maximum allowed violation δ to be 0.01. For each threshold, we take the average performance over 50 runs. The results in Figure 1 show that the cumulative regrets are negative for all thresholds. The cumulative regret is also smaller (better) for a bigger feasible threshold. This may be because our online algorithm learns faster by updating the dual variables more aggressively when the threshold is bigger.

Demographic Groups. For the demographic groups, we first divide the users into two groups based on gender and further divide each group into three subgroups based on age (1-24, 25-44, 45+), which we use when constructing the market share function. To ensure diverse preferences for each group, we select the top 20 movies with the largest absolute difference in average ratings between the two groups, which also have been rated by at least 1000 users. We use a mixed multinomial logit (MMNL) as our choice model for each group, where the mixture comes from the different age groups. We set the probability of each mixture as the empirical probability of each age group within the specific gender in the user data. Moreover, we set the weight for each movie proportional to users' average rating of the corresponding movie in the particular subgroup (based on gender and age). We provide an example in the full version of our paper and highlight that our market share function is monotone submodular, and the class of monotone submodular functions is conic-closed.

In the offline setting, each group's mixture probability and market share parameters are known beforehand. However, the decision time (choosing the set) does not know these parameters in the online setting. To introduce stochasticity into the market share function in the online setting, we make the probability of the mixtures of each group stochastic. This is achieved by drawing the probabilities from a Dirichlet distribution, parameterized by the means obtained from the offline setting. Specifically, for group i at time t, the probability of each mixture is sampled from a Dirichlet distribution with mean $(p_{i,1}, p_{i,2}, p_{i,3})$ obtained from the offline setting. Moreover, we set the number of periods to be $T = 10,000$ and the allowed violation to be $\delta = 0.01$ for both settings. The allowed violation controls the maximum violation allowed to satisfy the group fairness.

Implementation. For our offline algorithm, we use projected OGD as our online linear optimization method (c.f., discussion in our full version). We then use the greedy algorithm for maximizing monotone submodular functions with a cardinality constraint, as proposed by [Nemhauser *et al.*, 1978], as the $(1 - 1/e)$-approximate oracle. Meanwhile, we also use projected OGD as our OLO for the online algorithm. Additionally, we apply the framework proposed by [Niazadeh *et al.*, 2021] for the offline greedy algorithm. We use the Multiplicative Weight/Hedge algorithm to update the weights of each item for each position in the assortment with a learning rate of $\epsilon_t = \sqrt{1/t}$. The Hedge algorithm is a no-regret adversarial learning algorithm that keeps weight over the different arms at each time step and updates those weights based on the observed feedback. In our problem, subproblem i corresponds to determining the item for the i^{th} position in the assortment. (See [Niazadeh *et al.*, 2021] for more details.)

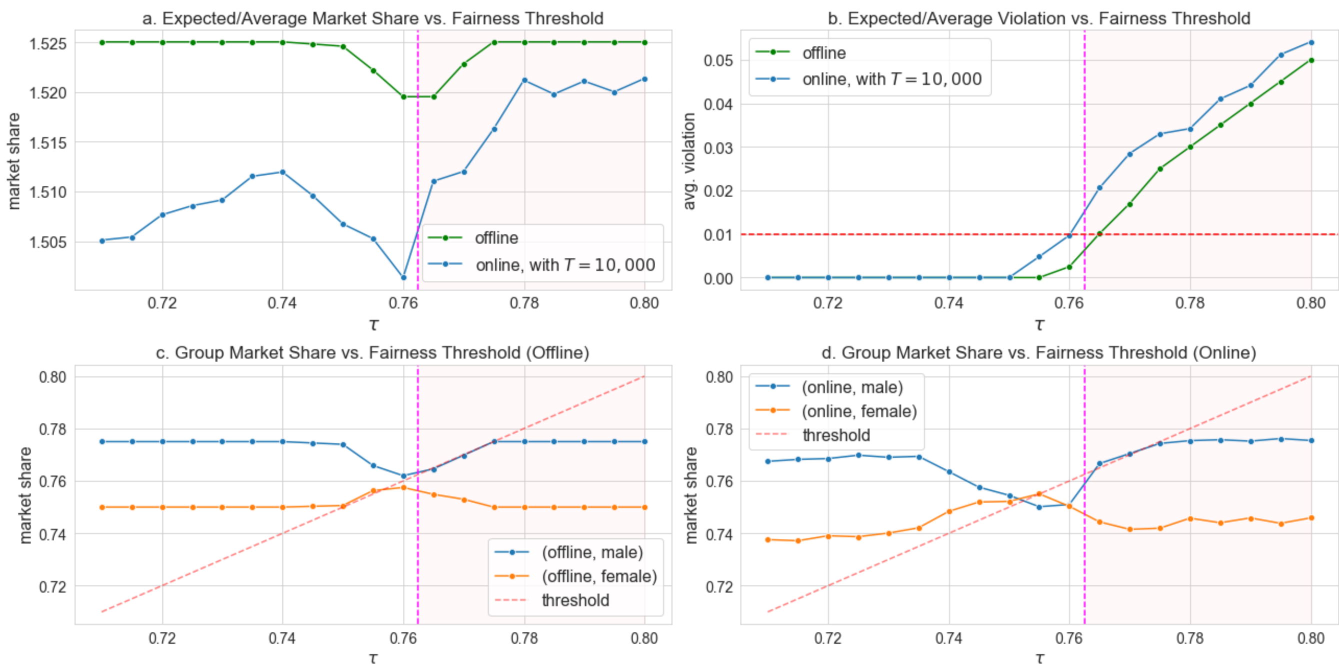

Figure 2: The effect of fairness constraints on market share. Under different levels of fairness thresholds, we show a) the expected/average market share obtained by our offline and online algorithms (the shaded pink region shows the infeasible space), b) the expected/average total violation of the fairness constraints obtained by our algorithms (the dashed red line is our fairness violation threshold $\delta = 0.01$), c) the expected/average market share for each group obtained by our offline algorithm, and d) the expected/average market share for each group obtained by our online algorithm.

The effect of fairness threshold. The results in Figure 2 provide insights into the effect of varying the fairness threshold on the performance of the offline and online algorithms. The presented values are the averages of 50 runs. In the offline setting, the total expected market share remains similar (almost constant) when the thresholds are feasible, as depicted in Figure 2a (green line). However, as the threshold approaches the boundaries of feasibility (in the range of $[0.74, 0.76]$), the performance decreases since the algorithm has to adjust the assortment to increase the market share of one group (female users in this case), while in turn decreasing the market share of the other group (male users), as seen in Figure 2c. Moreover, when the threshold becomes infeasible, the algorithm no longer tries to balance the two groups, resulting in a higher total expected market share. This is expected, as we do not have any fairness guarantees when the optimization problem is infeasible. Additionally, we limit the domain of the dual variable $\boldsymbol{\alpha}$, causing the algorithm to no longer balance the two groups.

In contrast, in the online setting, the total average market share increases as the threshold increases (averaged over $T = 10,000$ rounds), as indicated by the blue line in Figure 2a. This may be because when the thresholds are higher, the online algorithm updates the dual variable faster, thus converging to the optimal solution more quickly. However, as the threshold approaches the boundaries of feasibility (in $[0.74, 0.76]$), the performance decreases and fluctuates. The algorithm adjusts the market shares of both groups to either satisfy or violate the fairness constraints by at most $\delta = 0.01$. Notably, in the boundary region, the average market share of females increases, while the market share of males decreases while still satisfying the constraint. When the threshold becomes infeasible, the algorithm maximizes the total average market share with no fairness guarantee.

5 Discussion

This paper provides a general framework for incorporating fairness constraints into online combinatorial problems, as our methods use approximation algorithms for the base problem and OLO algorithms as black boxes. Our framework is easy to state and implement, making it especially relevant for applications in online markets. There are two important future directions for our work. In the online setting, we consider a stochastic adversary, and the stronger adaptive adversary, which can provide arbitrary functions at each time step, is much more challenging and probably requires new techniques. Note that there is a slack variable in our fairness guarantees, appearing from technical requirements for online gradient descent. We believe this is avoidable in some settings (where a barrier function can be constructed) using an interior point method to solve our lagrangian max-min game.

Acknowledgements

R.N. was supported by the Asness junior faculty fellowship at Chicago Booth. F.S. and N.G. were supported by an award from the Office of Naval Research (ONR) (Award Number: N00014-23-1-2584). K.K.P. was supported through the

NSF TRIPOD Institute on Data, Economics, Algorithms and Learning (IDEAL) and other awards from DARPA and NSF.

Contribution Statement

Author names appear alphabetically, and all authors contributed equally to this submission.

References

[Asadpour *et al.*, 2022] Arash Asadpour, Rad Niazadeh, Amin Saberi, and Ali Shameli. Sequential submodular maximization and applications to ranking an assortment of products. *Operations Research*, 2022.

[Balseiro *et al.*, 2021] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. Regularized online allocation problems: Fairness and beyond. In *International Conference on Machine Learning*, pages 630–639. PMLR, 2021.

[Ben-Tal *et al.*, 2013] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.

[Bertsimas *et al.*, 2011] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. The price of fairness. *Operations research*, 59(1):17–31, 2011.

[Bertsimas *et al.*, 2012] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012.

[Besbes *et al.*, 2014] Omar Besbes, Yonatan Gur, and Assaf Zeevi. Stochastic multi-armed-bandit problem with non-stationary rewards. *Advances in neural information processing systems*, 27, 2014.

[Biega *et al.*, 2018] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*, pages 405–414, 2018.

[California Legislative Service, 2020] California Legislative Service. The california privacy rights act of 2020, 2020. Available at: https://cppa.ca.gov/regulations/.

[Celis *et al.*, 2018] Elisa L. Celis, Damian Straszak, and Nisheeth K. Vishnoi. Ranking with fairness constraints. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[Chen *et al.*, 2018] Lin Chen, Hamed Hassani, and Amin Karbasi. Online continuous submodular maximization. In *International Conference on Artificial Intelligence and Statistics*, pages 1896–1905. PMLR, 2018.

[Chen *et al.*, 2022] Qinyi Chen, Negin Golrezaei, Fransisca Susan, and Edy Baskoro. Fair assortment planning. *arXiv preprint arXiv:2208.07341*, 2022.

[Chierichetti *et al.*, 2019] Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvtiskii. Matroids, matchings, and fairness. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2212–2220. PMLR, 2019.

[Chouldechova and Roth, 2018] Alexandra Chouldechova and Aaron Roth. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810*, 2018.

[Dekel, 2012] Ofer Dekel. Combinatorial optimization. Lecture Notes, 2012. [Online; accessed 2024-01-17]. Available: https://courses.cs.washington.edu/courses/cse599s/12sp/scribes/lecture16.pdf.

[Deng *et al.*, 2022] Yuan Deng, Negin Golrezaei, Patrick Jaillet, Jason Cheuk Nam Liang, and Vahab Mirrokni. Fairness in the autobidding world with machine-learned advice. *arXiv preprint arXiv:2209.04748*, 2022.

[Dobzinski and Schapira, 2006] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1064–1073, 2006.

[Duchi *et al.*, 2021] John C. Duchi, Peter W. Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 46(3):946–969, 2021.

[Dudik *et al.*, 2020] Miroslav Dudik, Nika Haghtalab, Haipeng Luo, Robert E. Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. Oracle-efficient online learning and auction design. *Journal of the ACM (JACM)*, 67(5):1–57, 2020.

[Dwork *et al.*, 2012] 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, 2012.

[Garcia-Soriano and Bonchi, 2021] David Garcia-Soriano and Francesco Bonchi. Maxmin-fair ranking: individual fairness under group-fairness constraints. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 436–446, 2021.

[Goyal *et al.*, 2016] Vineet Goyal, Retsef Levi, and Danny Segev. Near-optimal algorithms for the assortment planning problem under dynamic substitution and stochastic demand. *Operations Research*, 64(1):219–235, 2016.

[Harper and Konstan, 2015] Maxwell Harper and Joseph Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.

[Hazan, 2016] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.

[Hitsch *et al.*, 2010] Gunter J. Hitsch, Ali Hortacsu, and Dan Ariely. What makes you click?—mate preferences in online dating. *Quantitative marketing and Economics*, 8:393–427, 2010.

[Jafarnia-Jahromi *et al.*, 2021] Mehdi Jafarnia-Jahromi, Liyu Chen, Rahul Jain, and Haipeng Luo. Online learning for stochastic shortest path model via posterior sampling. *arXiv preprint arXiv:2106.05335*, 2021.

[Joseph *et al.*, 2016] Matthew Joseph, Michael Kearns, Jamie H. Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. *Advances in neural information processing systems*, 29, 2016.

[Kakade *et al.*, 2007] Sham M. Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 546–555, 2007.

[Kalai and Vempala, 2005] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.

[Kearns *et al.*, 2018] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International conference on machine learning*, pages 2564–2572. PMLR, 2018.

[Kearns *et al.*, 2019] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. An empirical study of rich subgroup fairness for machine learning. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 100–109, 2019.

[Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.

[Kleinberg *et al.*, 2018] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. Algorithmic fairness. In *Aea papers and proceedings*, volume 108, pages 22–27, 2018.

[Liu *et al.*, 2022] Qingsong Liu, Weihang Xu, Siwei Wang, and Zhixuan Fang. Combinatorial bandits with linear constraints: Beyond knapsacks and fairness. *Advances in Neural Information Processing Systems*, 35:2997–3010, 2022.

[Loi *et al.*, 2019] Michele Loi, Anders Herlitz, and Hoda Heidari. A philosophical theory of fairness for prediction-based decisions. *Available at SSRN 3450300*, 2019.

[Manshadi *et al.*, 2021] Vahideh Manshadi, Rad Niazadeh, and Scott Rodilitz. Fair dynamic rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 694–695, 2021.

[Miettinen, 1999] Kaisa Miettinen. *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business Media, 1999.

[Nemhauser *et al.*, 1978] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978.

[Niazadeh *et al.*, 2021] Rad Niazadeh, Negin Golrezaei, Joshua R. Wang, Fransisca Susan, and Ashwinkumar Badanidiyuru. Online learning via offline greedy algorithms: Applications in market design and optimization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 737–738, 2021.

[Park and Kwan, 2020] Yoo Min Park and Mei-Po Kwan. Understanding racial disparities in exposure to traffic-related air pollution: Considering the spatiotemporal dynamics of population distribution. *International journal of environmental research and public health*, 17(3):908, 2020.

[Rahmattalabi *et al.*, 2019] Aida Rahmattalabi, Phebe Vayanos, Anthony Fulginiti, Eric Rice, Bryan Wilder, Amulya Yadav, and Milind Tambe. Exploring algorithmic fairness in robust graph covering problems. *Advances in neural information processing systems*, 32, 2019.

[Ribeiro, 2002] Carlos H. C. Ribeiro. Reinforcement learning agents. *Artificial intelligence review*, 17:223–250, 2002.

[Roughgarden and Wang, 2018] Tim Roughgarden and Joshua R. Wang. An optimal learning algorithm for online unconstrained submodular maximization. In *Conference On Learning Theory*, pages 1307–1325. PMLR, 2018.

[Sagawa *et al.*, 2019] Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*, 2019.

[Shalev-Shwartz and Singer, 2007] Shai Shalev-Shwartz and Yoram Singer. A primal-dual perspective of online learning algorithms. *Machine Learning*, 69:115–142, 2007.

[Singh and Joachims, 2018] Ashudeep Singh and Thorsten Joachims. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2219–2228, 2018.

[Soma *et al.*, 2022] Tasuku Soma, Khashayar Gatmiry, and Stefanie Jegelka. Optimal algorithms for group distributionally robust optimization and beyond. *arXiv preprint arXiv:2212.13669*, 2022.

[Streeter and Golovin, 2008] Matthew Streeter and Daniel Golovin. An online algorithm for maximizing submodular functions. *Advances in Neural Information Processing Systems*, 21, 2008.

[Tang and Yuan, 2023] Shaojie Tang and Jing Yuan. Beyond submodularity: a unified framework of randomized set selection with group fairness constraints. *Journal of Combinatorial Optimization*, 45(4):102, 2023.

[Zinkevich, 2003] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 928–936, 2003.