
Fiduciary Bandits

Gal Bahar¹ Omer Ben-Porat¹ Kevin Leyton-Brown² Moshe Tennenholtz¹

Abstract

Recommendation systems often face exploration-exploitation tradeoffs: the system can only learn about the desirability of new options by recommending them to some user. Such systems can thus be modeled as multi-armed bandit settings; however, users are self-interested and cannot be made to follow recommendations. We ask whether exploration can nevertheless be performed in a way that scrupulously respects agents' interests—i.e., by a system that acts as a *fiduciary*. More formally, we introduce a model in which a recommendation system faces an exploration-exploitation tradeoff under the constraint that it can never recommend any action that it knows yields lower reward in expectation than an agent would achieve if it acted alone. Our main contribution is a positive result: an asymptotically optimal, incentive compatible, and *ex-ante* individually rational recommendation algorithm.

1. Introduction

Multi-armed bandits (henceforth MABs) (Bubeck et al., 2012; Cesa-Bianchi & Lugosi, 2006) is a well-studied problem domain in online learning. In that setting, several arms (i.e., actions) are available to a planner; each arm is associated with an unknown reward distribution, from which rewards are sampled independently each time the arm is pulled. The planner selects arms sequentially, aiming to maximize her sum of rewards. This often involves a tradeoff between *exploiting* arms that have been observed to yield good rewards and *exploring* arms that could yield even higher rewards. Many variations of this model exist, including stochastic (Abbasi-Yadkori et al., 2011; Karmin et al., 2013), Bayesian (Chapelle & Li, 2011; Agrawal & Goyal, 2012), contextual (Chu et al., 2011; Slivkins, 2014), adver-

sarial (Auer et al., 1995) and non-stationary (Besbes et al., 2014; Levine et al., 2017) bandits.

This paper considers a setting motivated by recommender systems. Such systems suggest actions to agents based on a set of current beliefs and assess agents' experiences to update these beliefs. For instance, in navigation applications (e.g., Waze; Google maps) the system recommends routes to drivers based on beliefs about current traffic congestion. The planner's objective is to minimize users' average travel time. The system cannot be sure of the congestion on a road segment that no agents have recently traversed; thus, navigation systems offer the best known route most of the time and explore occasionally. Of course, users are not eager to perform such exploration; they are self-interested in the sense that they care more about minimizing their own travel times than they do about conducting surveillance about traffic conditions for the system.

A recent line of work (Kremer et al., 2014; Mansour et al., 2015), inspired by the viewpoint of algorithmic mechanism design (Nisan & Ronen, 1999; Nisan et al., 2007), deals with that challenge by *incentivizing exploration*—that is, setting up the system in such a way that no user would ever rationally choose to decline an action that was recommended to him. The key reason that it is possible to achieve this property while still performing a sufficient amount of exploration is that the planner has more information than the agents. At each point in time, each agent holds beliefs about the arms' reward distributions; the planner has the same information, but also knows about all of the arms previously pulled and the rewards obtained in each case. More specifically, Kremer et al. (2014) consider a restricted setting and devise an MAB algorithm that is *incentive compatible* (IC), meaning that whenever the algorithm recommends arm i to an agent, the best response of the agent is to select arm i .

Although this approach explicitly reasons about agents' incentives, it does not treat agents *fairly*: agents who are asked to explore receive lower expected rewards. More precisely, in their attempt to reach optimality (in the static setting) or minimize regret (in the online setting), these IC MAB algorithms are intentionally providing (*a priori*) sub-optimal recommendations to some of the agents. In particular, some of the agents could be better off by not using the system and follow their *default arm*—the *a priori* superior arm, which

¹Technion – Israel Institute of Technology ²University of British Columbia, Canada. Correspondence to: Omer Ben-Porat <omerbp@campus.technion.ac.il>.

would be every agent’s rational choice in the absence both of knowledge of other agents’ experiences and of a trusted recommendation. Thus, it would be natural for agents to see the recommendations of such IC MAB algorithms as a betrayal of trust; they might ask “why should I trust a recommender that occasionally gives out recommendations it has every reason to believe could make me worse off?”

In this work, we explicitly suggest that a social welfare maximization standpoint might raise societal issues, harming the trust agents put in recommender systems. The central premise of this paper is that explore-and-exploit AI systems should satisfy *individual guarantees*—guarantees that the system should fulfill for each agent *independently from the other agents and their recommendations*. At the one end of the spectrum are current MAB algorithms—successful in maximizing welfare, but do not offer the slightest individual guarantee. At the other end is the *fiduciary duty*: borrowed from law applications, it requires that the mechanism acts in the interest of its clients with all its knowledge. This is the strictest, and strongest, individual guarantee the system could provide. However, if we applied this standard, we would be left only with the mechanism that greedily picks the apparently best arm in each iteration. In some settings, perhaps this is the best that can be achieved; however, note that this mechanism is rarely able to learn anything. It is therefore natural to ask for an approach that enjoys both worlds—maximizing welfare while satisfying individual guarantees.

Our contribution We explore a novel compromise between these two extreme points, which we call *ex-ante* individual rationality (EAIR). To motivate it, we consider the benchmark reward of each agent to be that of the default arm: the reward agents would get if the recommender system is unavailable. A mechanism is EAIR if the reward of every recommendation it makes beats that benchmark in expectation, per the mechanism’s knowledge. More technically, a mechanism is EAIR if any probability distribution over arms that it selects has expected reward that is always at least as great as the reward of the default arm, both calculated based on the mechanism’s knowledge (which is more extensive than that of agents). While it is possible for the mechanism to sample a recommendation from a distribution that is *a priori* inferior to the (realization of the) default arm, the agent receiving the recommendation is nevertheless guaranteed to realize expected reward weakly greater than that offered by the default arm. Satisfying this requirement makes a MAB algorithm more appealing to agents; we foresee that in some domains, such a requirement might be imposed as fairness constraints by authorities.

Algorithmically, we focus on constructing optimal EAIR mechanisms. Our model is a bandit model with $K \geq 2$ arms and n agents (rounds). Similarly to Kremer et al. (2014),

we assume that rewards are fixed but initially unknown.

We consider two agent schemes. In the first part of the paper, we assume that agents follow recommendations, as in the classical MAB literature. This is the case if, e.g., agents are oblivious to some of the actions’ desirability, unaware of the entire set of alternatives, or if the cognitive overload of computing expectations is high. The main technical contribution of this paper is an EAIR mechanism, which obtains the highest possible social welfare by any EAIR mechanism up to an additive factor of $o(\frac{1}{n})$. Due to our static setting (rewards are realized only once), following the wrong exploration policy for even one agent has detrimental effect on social welfare. The optimality of our mechanism, which we term Fiduciary Explore & Exploit (FEE) and outline as Algorithm 1, follows from a careful construction of the exploration phase. Our analysis uses an intrinsic property of the setting, which is further elaborated in Theorem 1.

Later on, in Section 4, we adopt a different agent scheme, which is fully aligned with the incentivizing exploration literature. We assume that agents are strategic and have (the same) Bayesian prior over the rewards of the arms. In this context, a mechanism is *incentive compatible* (IC) if each agent’s expected reward is maximized by the recommended action. We provide a positive result in this challenging case as well. Our second technical contribution is Incentive Compatible Fiduciary Explore & Exploit (IC-FEE), which uses FEE as a black box, and is IC, EAIR and asymptotically optimal.

To complement this analysis, we also propose the more demanding concept of *ex-post* individual rationality (EPIR). The EPIR condition requires that a recommended arm must never be *a priori* inferior to the default arm given the planner’s knowledge. The EAIR and EPIR requirements differ in the guarantees that they provide to agents and correspondingly allow the system different degrees of freedom in performing exploration. We design an asymptotically optimal IC and EPIR mechanism. Finally, we analyze the social welfare cost of adopting either EAIR or EPIR mechanisms.

Related work Background on MABs can be found in Cesa-Bianchi & Lugosi (2006) and a recent survey (Bubeck et al., 2012). Despite that many works address MAB rounds as interacting agents, Kremer et al. (2014) is the first work of which we are aware that suggests that vanilla algorithms should be modified to deal with agents due to human nature and incentives. The authors considered two deterministic arms, a prior known both to the agents and the planner, and an arrival order that is common knowledge among all agents, and presented an optimal IC mechanism. Cohen & Mansour (2019) extended this optimality result to several arms under further assumptions. This setting has also been extended to regret minimization (Mansour et al., 2015), so-

cial networks (Bahar et al., 2016; 2019), and heterogeneous agents (Chen et al., 2018; Immorlica et al., 2019). All of this literature disallows paying agents; monetary incentives for exploration are discussed in e.g., (Chen et al., 2018; Frazier et al., 2014). None of this work considers the orthogonal, societal consideration of individual rationality constraint as we do here.

Our work also contributes to the growing body of work on fairness in Machine Learning (Ben-Porat & Tennenholtz, 2018; Dwork et al., 2012; Hardt et al., 2016; Liu et al., 2018). In the context of MABs, some recent work focuses on fairness in the sense of treating *arms* fairly. In particular, Liu et al. (2017) aim at treating similar arms similarly and Joseph et al. (2016) demand that a worse arm is never favored over a better one despite a learning algorithm’s uncertainty over the true payoffs. Finally, we note that the EAIR requirement we impose—that agents be guaranteed an expected reward at least as high as that offered by a default arm—is also related to the burgeoning field of safe reinforcement learning (Garcia & Fernández, 2015).

2. Model

Let $A = \{a_1, \dots, a_K\}$ be a set of K arms (actions). Rewards are deterministic but initially unknown: the reward of arm a_i is a random variable X_i , and $(X_i)_{i=1}^K$ are mutually independent. We denote by R_i the observed value of X_i . To clarify, rewards are realized only once; hence, once R_i is observed, $X_i = R_i$ for the rest of the execution. Further, we denote by μ_i the expected value of X_i , and assume for notational convenience that $\mu_1 \geq \mu_2 \geq \dots \geq \mu_K$. We also make the simplifying assumption that the rewards $(X_i)_{i=1}^K$ are fully supported on the set $[H]^+ \stackrel{\text{def}}{=} \{0, 1, \dots, H\}$, and refer to the continuous case in Section 6.

There are n agents, who arrive sequentially. We denote by a^l the action of the agent arriving at stage l . The reward of the agent arriving at stage l is denoted by R^l , and is a function of the arm she chooses. For instance, by selecting arm a_r the agent obtains $R^l(a_r) = X_r$. Agents are fully aware of the distribution of $(X_i)_{i=1}^K$. Each and every agent cares about her own reward, which she wants to maximize.

A mechanism is a recommendation engine that interacts with agents. The input for the mechanism at stage l is the sequence of arms pulled and rewards received by the previous $l - 1$ agents. The output of the mechanism is a recommended arm for agent l . Formally, a mechanism is a function $M : \bigcup_{l=1}^n (A \times \mathbb{R}_+)^{l-1} \rightarrow \Delta(A)$; of course, we can also define a deterministic notion that maps simply to A . The mechanism has a global objective, which is to maximize agents’ social welfare $\sum_{l=1}^n R^l(a^l)$.

We consider two agent schemes. The first is *non-strategic agents*, i.e., agents always follow the recommendation. An

underlying assumption of classical MAB algorithms, such behavior could be explicit in case the mechanism makes decisions for the agents; or implicit, e.g., agents are unaware of the entire set of alternatives or their desirability, or high cognitive overload is required to compute it. The second agent scheme is *strategic agents*: the mechanism makes action recommendations, but cannot compel agents to follow these recommendations. In this scheme, we say that a mechanism is incentive compatible (IC) when following its recommendations is a dominant strategy: that is, when given a recommendation, an agent’s best response is to follow her own recommendation. Formally,

Definition 1 (Incentive Compatibility). *A mechanism M is incentive compatible (IC) if $\forall l \in \{1, \dots, n\}$, for every history $h \in (A \times \mathbb{R}_+)^{l-1}$ and for all actions $a_r, a_i \in A$,*

$$\mathbb{E}(R^l(a_r) - R^l(a_i) \mid M(h) = a_r) \geq 0. \quad (1)$$

Unless stated otherwise, we address the non-strategic agents scheme. We handle the other agent scheme in Section 4.

When agents follow the mechanism, we can represent the mechanism’s (expected) social welfare by

$$SW(M) = \mathbb{E} \left[\frac{1}{n} \sum_{l=1}^n X_{M(h_l)} \right], \quad (2)$$

where $X_{M(h_l)} = \sum_{r=1}^K \Pr_{M(h_l)}(a_r) \mathbb{E}(X_r \mid h_l)$ is the reward agent l receives. Notice that $X_{M(h_l)}$ depends on the randomness of the rewards *and*, possibly, the randomness of $M(h_l)$.

Denote the highest possible social welfare under non-strategic agents by OPT. A mechanism M^* is said to be *optimal* if $SW(M^*) = \text{OPT}$. A mechanism M^* is *asymptotically optimal* if, for every “large enough” number of agents n greater than some number n' , it holds that $SW(M^*) \geq \text{OPT} - o(\frac{1}{n})$. This definition of approximation is equivalent to sub-linear regret in the MAB literature.

2.1. Individual Guarantees

An individual guarantee is a guarantee that a mechanism can provide to the agents it interacts with, independently of the other agents. In this subsection, we present our main conceptual contribution: a meaningful individual guarantee that allows exploration.

To put our guarantee in the right context, we first present the strictest and the strongest guarantee that could be provided. A mechanism is a *delegate* if it acts as the agent would have acted had it revealed the information it has with her. Formally, A mechanism M is a delegate if for every agent $l \in \{1, \dots, n\}$, every history $h \in (A \times \mathbb{R}_+)^{l-1}$ and every distribution \mathbf{p} over A , it holds that $\mathbb{E}(X_{M(h)} \mid h) \geq$

$\sum_{r=1}^K \mathbf{p}(r) \mathbb{E}(X_r | h)$. Indeed, this definition provides the strongest individual guarantee. It characterizes the greedy mechanism, GREEDY, which exploits in every round (according to the information it has). Noticeably, GREEDY performs little exploration, and probably leads to low social welfare. While sometimes relaxing this strong guarantee is impossible (e.g., banking or health-care), in many situations the planner is willing to relax individual guarantees to favor better social welfare.

The other extreme is to adopt a policy that we term FULL-EXPLORATION. FULL-EXPLORATION is the mechanism that first explores all arms sequentially, and then exploit the best arm. Clearly, at least for the non-strategic agent scheme, FULL-EXPLORATION is optimal when the number of agents is large enough. Nevertheless, with very high probability, it picks sub-optimal arms for the first K agents, which can be a highly undesired property.

Our guarantee builds on the popular economic concept of individual rationality. To introduce it, we propose the following thought experiment. Assume that agents have to make decisions without the mechanism. The agents know that $\mu_1 = \max_i \mu_i$; hence, we shall assume that every agent's *default action* is a_1 .¹ The default action is the action each agent would have selected if she did not use the mechanism. We compare the two options: picking the default arm or following the mechanism's action. If a mechanism guarantees that the latter is higher in expectation according to its knowledge, agents are better off using the mechanism. As a result, an individually rational mechanism should guarantee each agent at least the reward obtained by her default action. The next definition relies on this reasoning.

Definition 2 (*Ex-Ante Individual Rationality*). *A mechanism M is ex-ante individually rational (EAIR) if for every agent $l \in \{1, \dots, n\}$, and every history $h \in (A \times \mathbb{R}_+)^{l-1}$,*

$$\sum_{r=1}^K \Pr_{M(h)}(a_r) \mathbb{E}(X_r | h) \geq \mathbb{E}(X_1 | h). \quad (3)$$

The EAIR definition is conditioned on histories, i.e., the mechanism's knowledge. The right hand side is what an agent would get, given the knowledge of the mechanism, if she follows the default arm (which is optimal according to her knowledge). The left hand side is the expected value (over lotteries selected by the mechanism and reward distribution) guaranteed by the mechanism. Due to the mutual independence assumption, we must have $E(X_r | h) = R_r$ if arm a_r was observed under the history h and $E(X_r | h) = \mu_r$ otherwise. An EAIR mechanism must select a *portfolio* of arms with expected reward never inferior to the reward of the default arm a_1 .

¹As it will become apparent later, if agents have different default arms the social welfare can only increase since more arms could be explored.

Example. We now give an example to illustrate our setting and to familiarize the reader with our notation. Consider $K = 3$ arms, $H = 30$ and $X_1 \sim \text{Uni}\{0, \dots, 30\}$, $X_2 \sim \text{Uni}\{0, \dots, 20\}$, $X_3 \sim \text{Uni}\{0, \dots, 10\}$; thus $\mu_1 = 15$, $\mu_2 = 10$, and $\mu_3 = 5$. As always, a_1 is the default arm. To satisfy EAIR, a mechanism should recommend a_1 to the first agent, since EAIR requires that the expected value of any recommendation should weakly exceed R_1 . Let $h_1 = (a_1, R_1)$ be the history after the first agent. Now, we have three different cases. First, if $R_1 > \mu_2 = 10$, we know that $\mathbb{E}(X_2 | h_1) < R_1$ and $\mathbb{E}(X_3 | h_1) < R_1$; therefore, an EAIR mechanism can never explore any other arm, since any distribution over $\{a_2, a_3\}$ would violate Inequality (3). Second, if $R_1 \leq \mu_3 = 5$, then $\mathbb{E}(X_2 | h_1) \geq R_1$ and $\mathbb{E}(X_3 | h_1) \geq R_1$, and hence an EAIR mechanism can explore both a_2 and a_3 .

The third and most interesting case is where $\mu_3 < R_1 \leq \mu_2$, as when $R_1 = 8$. In this case, arm a_3 could only be recommended through a portfolio. An EAIR mechanism could select any distribution over $\{a_2, a_3\}$ that satisfies Inequality (3): any $p \in [0, 1]$ such that $p \cdot \mu_2 + (1 - p) \cdot \mu_3 \geq R_1$. This means that an EAIR mechanism can potentially explore arm a_3 , yielding higher expected social welfare overall than simply recommending a non-inferior arm deterministically.

3. Asymptotically Optimal EAIR Mechanism

In this section, we consider the case of non-strategic agents. We present the main technical contribution of this paper: a mechanism that asymptotically optimally balances the explore-exploit tradeoff while satisfying the EAIR property. The mechanism, which we term Fiduciary Explore & Exploit (FEE), is described as Algorithm 1. FEE is an event-based protocol that triggers every time an agent arrives. We now give an overview of FEE, focusing on the case where all agents adopt the recommendation of the mediator (we treat the other case in Section 4). We explain the algorithm's exploration phase in Subsection 3.1, describe the overall algorithm in Subsection 3.2, and prove the algorithm's formal guarantees in Subsection 3.3. We provide a comprehensive example of the way FEE operates in the appendix.

FEE is composed of three phases: primary exploration (Lines 1–6), secondary exploration (Lines 7–18), and exploitation (Lines 19). During the primary exploration phase, the mechanism compares the default arm a_1 to whichever other arms are permitted by the individual rationality constraint. This turns out to be challenging for two reasons. First, the order in which arms are explored matters; tackling them in the wrong order can reduce the set of arms that can be considered overall. Second, it is nontrivial to search in the continuous space of probability distributions over arms. To address this latter issue, we present a key lemma that allows us to use dynamic programming and find the optimal

exploration policy in time $O(2^K K^2 H^2)$. Because we expect K either to be fixed or to be significantly smaller than n, H , this policy is computationally efficient. Moreover, we note that the optimal exploration policy can be computed offline prior to the agents' arrival.

The primary exploration phase terminates in one of two scenarios: either the reward R_1 of arm a_1 is the best that was observed and thus no other arm could be explored (as in our example when $R_1 > 10$, or when $R_1 = 8$ and exploring a_2 yielded $R_2 \leq R_1$ and thus a_3 could not be explored), or another arm a_i was found to be superior to a_1 : i.e., an arm a_i was observed for which $R_i > R_1$. In the latter case, the mechanism gains the option of conducting a secondary exploration, using arm a_i to investigate all the arms that were not explored in the primary exploration phase. The third and final phase—to which we must proceed directly after the primary exploration phase if that phase does not identify an arm superior to the default arm—is to exploit the most rewarding arm observed.

Remark. In this section we assume that agents are non-strategic and follow the mechanism's recommendation.

3.1. Primary Exploration Phase

Performing primary exploration optimally requires solving a planning problem; it is a challenging one, because it involves a continuous action space and a number of states exponential in K and H . We approach this task as a Goal Markov Decision Process (GMDDP) (see, e.g., (Barto et al., 1995)) that abstracts everything but pure exploration. In our GMDDP encoding, all terminal states fall into one of two categories. The first category is histories that lead to pure exploitation of a_1 , which can arise either because EAIR permits no arm to be explored or because all explored arms yield rewards inferior to the observed R_1 ; the second is histories in which an arm superior to a_1 was found. Non-terminal states thus represent histories in which it is still permissible for some arms to be explored. The set of actions in each non-terminal state is the set of distributions over the non-observed arms (i.e., portfolios) corresponding to the history represented in that state, which satisfy the EAIR condition. The transition probabilities encode the probability of choosing each candidate arm from a portfolio; observe that the rewards of each arm are fixed, so this is not a source of additional randomness in our model. GMDDP rewards are given in terminal nodes only: either the observed R_1 if no superior arm was found or the expected value of the maximum between the superior reward discovered and the maximal reward of all unobserved arms (since in this case, as we show later on, the mechanism is able to explore all arms w.h.p. during the secondary exploration phase).

Formally, the GMDDP is a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R} \rangle$, where

- \mathcal{S} is a finite set of states. Each state s is a pair (O, U) , where $O \subseteq \{(a, c) \mid a \in A, c \in H\}$ is the set of arm-reward pairs that have been observed so far, with each a appearing at most once in O (since rewards from the arms are deterministic): for every (O, U) and every $a \in A$, $|\{c \mid (a, c) \in O\}| \leq 1$. $U \subseteq A$ is the set of arms not yet explored. The initial state is thus $s_0 = (\emptyset, A)$. For every non-empty² set of pairs O we define $\alpha(O)$ to be the reward observed for arm a_1 , and $\beta(O) = \max_{c: \exists a, (a, c) \in O} c$ to be the maximal reward observed.

- $\mathcal{A} = \bigcup_{s \in \mathcal{S}} \mathcal{A}_s$ is an infinite set of actions. For each $s = (O, U) \in \mathcal{S}$, \mathcal{A}_s is defined as follows:

1. If $s = s_0$, then $\mathcal{A}_{s_0} = \Delta(\{a_1\})$: i.e., a deterministic selection of a_1 .
2. Else, if $\alpha(O) < \beta(O)$, then $\mathcal{A}_s = \emptyset$. This condition implies that we can move to secondary exploration.
3. Otherwise, \mathcal{A}_s is a subset of $\Delta(U)$, such that $\mathbf{p} \in \mathcal{A}_s$ if and only if $\sum_{a_i \in U} \mathbf{p}(a_i) \mu_{a_i} \geq \alpha(O)$. Notice that this resembles the EAIR condition given in Inequality (3). Moreover, the case where none of the remaining arms have strong enough priors to allow exploration falls here as a vacuous case of the above inequality.

We denote by \mathcal{S}_T the set of *terminal* states, namely $\mathcal{S}_T = \{s \in \mathcal{S} \mid \mathcal{A}_s = \emptyset\}$.

- \mathcal{P} is the transition probability function. Let $s = (O, U) \in \mathcal{S}$, and let $s' = (O', U')$ such that $O' = O \cup \{(a_i, c)\}$ and $U' = U \setminus \{a_i\}$ for some $a_i \in U, c \in [H]^+$. Then, the transition probability from s to s' given an action \mathbf{p} is defined by $\mathcal{P}(s' | s, \mathbf{p}) = \mathbf{p}(a_i) \Pr(X_i = c)$. If s' is some other state that does not meet the conditions above, then let $\mathcal{P}(s' | s, \mathbf{p}) = 0$ for every $\mathbf{p} \in \mathcal{A}_s$.

- $\mathcal{R} : \mathcal{S}_T \rightarrow \mathbb{R}$ is the reward function, defined on terminal states only. For each terminal state $s = (O, U) \in \mathcal{S}_T$,

$$\mathcal{R}(s) = \begin{cases} \alpha(O) & \alpha(O) = \beta(O) \\ \mathbb{E}[\max\{\beta(O), \max_{a_{i'} \in U} X_{i'}\}] & \alpha(O) < \beta(O) \end{cases}.$$

That is, when a_1 was the highest-reward arm observed, the reward of a terminal state is $\alpha(O)$; otherwise, it is the expectation of the maximum between $\beta(O)$ and the highest reward of all unobserved arms. The reward depends on unobserved arms since the secondary exploration phase allows us to explore all these arms; hence, their values are also taken into account.

A policy $\pi : (\mathcal{S} \times \mathcal{A})^* \times \mathcal{S} \rightarrow \mathcal{A}$ is a function from all GMDDP histories (sequences of states and actions) and a current state to an action. A policy π is *valid* if for every

²Due to the construction, every non-empty O must contain (a_1, c) for some $c \in [H]^+$.

history h and every non-terminal state s , $\pi(h, s) \in \mathcal{A}_s$. A policy π is *stationary* if for every two histories h, h' and a state s , $\pi(h, s) = \pi(h', s)$. When discussing a stationary policy, we thus neglect its dependency on h , writing $\pi(s)$.

Given a policy π and a state s , we denote by $W(\pi, s)$ the expected reward of π when initialized from s , which is defined recursively from the terminal states:

$$W(\pi, s) = \begin{cases} \mathcal{R}(s) & \text{if } s \in S_T \\ \sum_{s' \in \mathcal{S}} \mathcal{P}(s'|s, \pi(s)) W(\pi, s') & \text{otherwise.} \end{cases}$$

We now turn to our technical results. The following lemma shows that we can safely focus on stationary policies that effectively operate on a significantly reduced state space.

Lemma 1. *For every policy π there exists a stationary policy π' such that (1) $\pi'(s) = \pi'(s')$ for every pair of states $s = (O, U)$ and $s' = (O', U)$ with $\alpha(O) = \alpha(O')$ and $\beta(O) = \beta(O')$; and (2) for every state s , $W(\pi', s) \geq W(\pi, s)$.*

Lemma 1 tells us that there exists an optimal, stationary policy that selects the same action in every pair of states that share the same unobserved set U and values $\alpha(O)$ and $\beta(O)$, but are distinguished in the O component. Thus, we do not need a set of states whose size depends on the number of possible arm-reward observation histories: all we need to record is U and a real value for either $\alpha(O)$ and $\beta(O)$, reducing the number of states to $O(2^K H)$.

We still have one more challenge to overcome: the set of actions \mathcal{A}_s available in each state s is infinite. Despite that \mathcal{A}_s is a convex polytope and thus we can apply Linear Programming, our approach is much more computationally efficient and interpretable. We prove that there exists an optimal “simple” policy, which we denote π^* . Given two indices $i, r \in \{2, \dots, K\}$, we denote by \mathbf{p}_{ir}^α (for $i \neq r$) and by \mathbf{p}_{ii}^α (for $i = r$) the distributions over $\{a_1, \dots, a_K\}$ such that

$$\mathbf{p}_{ir}^\alpha(a) = \begin{cases} \frac{|\alpha - \mu_r|}{|\alpha - \mu_i| + |\alpha - \mu_r|} & \text{if } a = a_i \\ \frac{|\alpha - \mu_i|}{|\alpha - \mu_i| + |\alpha - \mu_r|} & \text{if } a = a_r \\ 0 & \text{otherwise} \end{cases},$$

and $\mathbf{p}_{ii}^\alpha(a) = 1$ if and only if $a = a_i$. When $\alpha = \alpha(O)$ is clear from context, we omit it from the superscript.

We are now ready to describe the policy π^* , which we later prove to be optimal. For the initial state s_0 , $\pi^*(s_0) = \mathbf{p}_{11}$. For every non-terminal state $s = (O, U) \in \mathcal{S}$ with $s \neq s_0$, $\pi^*(s) = \mathbf{p}_{i^*r^*}$ such that $(i^*, r^*) \in \mathcal{A}_s$ maximize

$$\left(1 - \frac{\mathbb{1}_{i=r}}{2} \right) \left[\mathbf{p}_{ir}(i) \sum_{s' \in \mathcal{S}} \mathcal{P}(s'|s, \mathbf{p}_{ii}) W(\pi^*, s') + \mathbf{p}_{ir}(r) \sum_{s' \in \mathcal{S}} \mathcal{P}(s'|s, \mathbf{p}_{rr}) W(\pi^*, s') \right].$$

The optimality of π^* follows from a property that is formally proven in Theorem 1: any policy π that satisfies the conditions of Lemma 1 can be presented as a mixture of policies that solely take actions of the form $(\mathbf{p}_{ir})_{i,r}$. As a result, we can improve π by taking the best such policy from that mixture. We derive π^* via dynamic programming, where the base cases are the set of terminal states. For any other state, $\pi^*(s)$ is the best action of the form \mathbf{p}_{ir} as defined above, considering all states that are reachable from s . While any policy π' can be encoded as a weighted sum over such “simple” policies, π^* is the best one, and hence is optimal.

Theorem 1. *For every valid policy π and every state s , it holds that $W(\pi^*, s) \geq W(\pi, s)$.*

Since our compressed state representation consists of $O(2^K H)$ states, the computation of π^* in each stage requires us to consider $O(K^2)$ candidate actions, each of which involves summation of at most $H + 1$ summands; thus, π^* can be computed in $O(2^K K^2 H^2)$ time.

3.2. Intuitive Description of FEE

We now present the FEE algorithm, stated formally as Algorithm 1. The primary exploration phase (Lines 1–6) is based on the GMDP from the previous subsection. It is composed of computing π^* and then producing recommendations according to its actions, each of which defines a distribution over (at most) two actions. Let (U, O) denote the terminal state reached by π^* (the primary exploration selects a fresh arm in each stage; hence such a state is reached after at most K agents).

We then enter the secondary exploration phase. If $\beta(O) = R_1$ then this phase is vacuous: no distribution over the unobserved arms can satisfy the EAIR condition and/or all the observed arms are inferior to arm a_1 . On the other hand, if $\beta(O) > R_1$ (Line 7), we found an arm $a_{\bar{r}}$ with a reward superior to R_1 , and can use it to explore all the remaining arms. For every $a_i \in U$, the mechanism operates as follows. If the probability of a_i yielding a reward greater than $a_{\bar{r}}$ is zero, we neglect it (Lines 11–13). Else, if $\mu_i \geq R_1$, we recommend a_i . This is manifested in the second condition in Line 15. Otherwise, $\mu_i < R_1$. In this case, we select a distribution over $\{a_{\bar{r}}, a_i\}$ that satisfies the EAIR condition and explore a_i with the maximal possible probability, which is $\mathbf{p}_{\bar{r}i}(i)$. As we show formally in the proof of Lemma 2, the probability of exploring a_i in this case is at least $\frac{1}{H}$, implying that after H tries in expectation the algorithm would succeed to explore a_i .

Ultimately (Line 19), FEE recommends the best observed arm to all the remaining agents.

3.3. Algorithmic Guarantees

Algorithm 1 Fiduciary Explore & Exploit (FEE)

- 1: Initialize a GMDP instance $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R} \rangle$, and compute π^* .
 - 2: Set $s = (O, U) = (\emptyset, A)$.
 - 3: **while** s is not terminal **do**
 - 4: Draw arm $a_i \sim \pi^*(s)$, recommend a_i and observe R_i .
 - 5: $O \leftarrow O \cup \{(a_i, R_i)\}, U \leftarrow U \setminus \{a_i\}$.
 - 6: $s \leftarrow (O, U)$.
 - 7: **if** $\beta(O) > R_1$ **then**
 - 8: **while** U is not empty **do**
 - 9: Let $a_{\bar{r}}$ s.t. $a_{\bar{r}} \in \arg \max_{a_r \in A \setminus U} R_r$.
 - 10: Select an arbitrary arm $a_i \in U$.
 - 11: **if** $\Pr(X_i > R_{\bar{r}}) = 0$ **then**
 - 12: $U \leftarrow U \setminus \{a_i\}$.
 - 13: **continue**.
 - 14: Draw $Y \sim \text{Uni}[0, 1]$.
 - 15: **if** $Y \leq \frac{R_{\bar{r}} - R_1}{R_{\bar{r}} - \mu_i}$ or $\mu_i \geq R_1$ **then**
 - 16: Recommend a_i and observe R_i .
 - 17: $U \leftarrow U \setminus \{a_i\}$.
 - 18: **else**, recommend $a_{\bar{r}}$.
 - 19: Recommend $a_{i^*} \in \arg \max_{a_i \in A \setminus U} R_i$ to all agents.
-

We begin by arguing that FEE is indeed EAIR.

Proposition 1. FEE satisfies the EAIR condition.

The proof of Proposition 1 is highly intuitive: the reward of every recommendation FEE makes always exceed R_1 in expectation. We now move on to consider the social welfare of FEE. Let OPT_{EAIR} denote the highest welfare attained by any EAIR mechanism. First, we show that the expected value of π^* at s_0 , denoted by $W(\pi^*, s_0)$, upper bounds the social welfare of any EAIR mechanism.

Theorem 2. It holds that $\text{OPT}_{\text{EAIR}} \leq W(\pi^*, s_0)$.

The proof proceeds by contradiction: given an EAIR mechanism M , we construct a series of progressively-easier-to-analyze EAIR mechanisms with non-decreasing social welfare; we modify the final mechanism by granting it oracular capabilities, making it violate the EAIR property and yet preserving reducibility to a policy for the GMDP of Subsection 3.1. We then argue via the optimality of π^* that the oracle mechanism cannot obtain a social welfare greater than $W(\pi^*, s_0)$. Next, we lower bound the social welfare of FEE.

Lemma 2. $SW_n(\text{FEE}) \geq \text{OPT}_{\text{EAIR}} - O\left(\frac{KH^2}{n}\right)$.

The proof relies mainly on an argument that the primary and secondary explorations will not be too long on average: after $(K+1)H$ agents the mechanism is likely to begin exploiting. Noting that the lower bound of Lemma 2 asymptotically approaches the upper bound of Theorem 2, we conclude that FEE is asymptotically optimal.

4. Incentive Compatibility

In this section, we consider the second and more challenging agent scheme: strategic agents. Our main goal is to show that FEE, which we developed in Section 3.3, can be modified to satisfy IC as well.³ We remark that there are cases that an IC mechanism cannot explore all arms, regardless of individual rationality constraints. To illustrate, assume that $\Pr(X_1 \geq \mu_2) = 1$, i.e., the reward of arm a_1 is always greater or equal to the expected reward of arm a_2 . In this case, no agent will ever follow a recommendation for arm a_2 . Consequently, we shall make the following standard assumption (see, e.g., (Mansour et al., 2015))

Assumption 1. For every i, j such that $1 \leq i < j \leq K$, it holds that $\Pr(X_i < \mu_j) > 0$.

If Assumption 1 does not hold for some pair (i, j) , arm a_j would never be explored; hence, we can remove such arms from A . We shall also make the simplifying assumption that $\mu_1 > \mu_2 \geq \dots \geq \mu_K$, as otherwise the problem becomes easier to solve.

Among other factors, the expectation in Inequality (1) is taken over agents' information on the arrival order. On the one extreme, the arrival order could be uniform, i.e., each agent l is entirely oblivious about her "place in line." In this case, as we show in the appendix, FEE satisfies IC as is assuming that there are sufficiently many agents. On the other extreme, which is the more popular in prior work (Kremer et al., 2014; Mansour et al., 2015), agents have complete information about their rounds. Namely, the agent arriving at time l knows that she is the l 'th agent. The complete information case is the more demanding one, and an IC mechanism for this case will also be IC under any distributional assumption on the arrival order. Nevertheless, as we demonstrate shortly, it requires more technical work.

We build on the techniques of Mansour et al. (2015) and use *phases*: each phase contains one round of exploration (that is, following FEE) and the other rounds are either exploitation via GREEDY (defined in Subsection 2.1) or recommendation of arm a_1 . An IC version of FEE, which we term IC-FEE, is outlined as Algorithm 2.

IC-FEE works as follows. It initializes an instance of FEE, and uses it seldom in the earlier rounds, and regularly afterward (every time IC-FEE makes a recommendation, it updates FEE). In Line 2, it recommends a_1 to the first agent. Recall that π^* employed by FEE is only allowed to pick a_1 w.p. 1 in the first round; hence, FEE and IC-FEE coincide with the first recommendation. Then, depending on the

³For simplicity, we formulated IC-FEE to satisfy IC in the best response sense: given that all other agents follow their recommendations, it is an agent's best response to adopt the recommendation as well. However, IC-FEE can be easily modified to offer dominant strategies to agents.

value of R_1 , it recommends agents $2, \dots, K$ either greedily (maximizing the reward in each round, Line 3) or arm a_1 (Line 4). Later, in Line 5, it splits the remaining rounds to *phases* of size B (B will be determined later on). In each such phase k , we first ask whether FEE is exploring or exploiting (Line 7). If FEE exploits (Line 19 in Algorithm 1), every agent of every phase from here on will be recommended by FEE. If that is not the case (see the else block starting at Line 9), IC-FEE picks one agent from the B agents of this phase uniformly at random, denoted $l(k)$. Then, agent $l(k)$ gets the recommendation from FEE. The recommendation policy for the rest of the agents in this phase depends on the observed arms. If IC-FEE already discovered an arm a_i with $R_i > R_1$ (Line 11), we let agent l exploit using GREEDY. Otherwise (Line 12), IC-FEE recommends a_1 .

Lines 11 and 12 are also where our mechanism departs from the principles of prior work. For example, in the work of Mansour et al. (2015), each phase contains one round of exploration and the rest are exploitation rounds. In our setting, agents that are not exploring might still not exploit. The distinction between Lines 11 and 12 is crucial: exploiting unobserved arms might lead to sub-optimal welfare, since they are the chance to explore arms with expected reward below R_1 . We elaborate more in the proof of Theorem 3.

To determine the phase length B , we introduce the following quantities ξ and γ . Due to Assumption 1, there exist $\xi > 0$ and $\gamma > 0$ such that for all $i \in [K]$, it holds that $\Pr(\forall i' \in [K] \setminus \{i\} : \mu_i - X_{i'} > \xi) > \gamma$. In words, it says that the reward of every arm i is greater than all other arms by at least ξ , w.p. of at least γ . The following Theorem 3 summarizes the properties of IC-FEE.

Theorem 3. *Let the phase length be $B = \left\lceil \frac{H}{\xi\gamma} \right\rceil + 1$. Under Assumption 1, IC-FEE satisfies EAIR and IC. In addition, $SW_n(\text{IC-FEE}) \geq \text{OPT}_{\text{EAIR}} - O\left(\frac{KH^3}{n\xi\gamma}\right)$.*

5. Further Analysis

Notice that EAIR mechanisms guarantee each agent the value of the default arm, but only in expectation. We now propose a more strict form of individual rationality, *ex-post* individual rationality (EPIR).

Definition 3 (*Ex-Post Individual Rationality*). *A mechanism M is ex-post individually rational (EPIR) if for every agent $l \in \{1, \dots, n\}$, every history $h \in (A \times \mathbb{R}_+)^{l-1}$, and every arm a_r such that $\Pr_{M(h)}(a_r) > 0$, it holds that $\mathbb{E}(X_r - X_1 | h) \geq 0$.*

Satisfying EPIR means that the mechanism never recommends an arm that is *a priori* inferior to arm a_1 given the mechanism's knowledge. It is immediate to see that every EPIR mechanism is also EAIR. EPIR mechanisms are quite

Algorithm 2 IC Fiduciary Explore & Exploit (IC-FEE)

- 1: Initialize an instance of FEE and update it after every recommendation.
 - 2: Recommend a_1 to the first agent, observe R_1 .
 - 3: **if** $R_1 < \mu_K$ **then** recommend as GREEDY to agents $2, \dots, K$.
 - 4: **else**, recommend a_1 to agents $2, \dots, K$.
 - 5: Split the remaining rounds into consecutive phases of B rounds each.
 - 6: **for** phase $k = 1, \dots$ **do**
 - 7: **if** FEE exploits (Line 19 in FEE) **then**
 - 8: follow FEE .
 - 9: **else**
 - 10: Pick one agent $l(k)$ from the B agents in this phase uniformly at random, and recommend her according to FEE.
 - 11: **if** an arm a_i with $R_i > R_1$ was revealed **then** recommend as GREEDY.
 - 12: **else**, recommend a_1 .
-

conservative, since they can only explore arms that yield expected rewards of at least the value R_1 obtained for a_1 . We develop an optimal IC/EPIR mechanism in the appendix.

5.1. Social Welfare Analysis

We now analyze the loss in social welfare due to individual rationality constraints. For simplicity, we consider the case of non-strategic agents. Recall that OPT is the highest possible social welfare, and OPT_{EAIR} is its counterpart after imposing EAIR. In addition, let OPT_{EPIR} and OPT_{DEL} denote the best asymptotic social welfare (w.r.t. some instance $\langle K, A, (X_i) \rangle$ and infinitely many agents) achievable by an EPIR and a delegate mechanisms, respectively. Noticeably, for every instance $\langle K, A, (X_i) \rangle$, it holds that $\text{OPT} \geq \text{OPT}_{\text{EAIR}} \geq \text{OPT}_{\text{EPIR}} \geq \text{OPT}_{\text{DEL}}$. In the rest of this subsection, we analyze the ratio of two subsequent optimal welfares. We begin by showing that individual guarantees can deteriorate welfare even for the most flexible notion, EAIR.

Proposition 2. *For every $K, H \in \mathbb{N}$, there exists an instance $\langle K, A, (X_i) \rangle$ with $\frac{\text{OPT}}{\text{OPT}_{\text{EAIR}}} \geq H \left(1 - e^{-\frac{K}{H}}\right)$.*

Proposition 2 shows that when K and H have the same magnitude, the ratio is on the order of H , meaning that EAIR mechanisms perform poorly when a large number of different reward values are possible. However, this result describes the worst case; it turns out that optimal EAIR mechanisms have constant ratio under some reward distributions. For example, as we show in the appendix this ratio is at most $\frac{8}{7}$ if $X_i \sim \text{Uni}\{0, 1, \dots, H\}$ for every $i \in \{2, \dots, K\}$ and X_1 is only slightly better a-priori.

Next, we consider the cost of adopting the stricter EPIR condition rather than EAIR. As Proposition 3 shows, by providing a more strict fiduciary guarantee the social welfare may be harmed by a factor of H .

Proposition 3. *For every $K, H \in \mathbb{N}$, there exists an instance $\langle K, A, (X_i) \rangle$ with $\frac{\text{OPT}_{\text{EAIR}}}{\text{OPT}_{\text{EPIR}}} \geq \frac{H+2}{3} \left(1 - e^{-\frac{K-2}{H}}\right)$.*

Finally, we show that the EPIR guarantee still allows us to significantly improve upon OPT_{DEL} .

Proposition 4. *For every $K, H \in \mathbb{N}$, there exists an instance $\langle K, A, (X_i) \rangle$ with $\frac{\text{OPT}_{\text{EPIR}}}{\text{OPT}_{\text{DEL}}} \geq \frac{H}{3} \left(1 - e^{-\frac{K-2}{H}}\right)$.*

6. Conclusions and Discussion

This paper introduces a model in which a recommender system must manage an exploration-exploitation tradeoff under the constraint that it may never knowingly make a recommendation that will yield lower reward than any individual agent would achieve if he/she acted without relying on the system.

We see considerable scope for follow-up work. First, from a technical point of view, our algorithmic results are limited to discrete reward distributions. One possible future direction would be to present an algorithm for the continuous case. More conceptually, we see natural extensions of EPIR and EAIR to stochastic settings, either by assuming a prior and requiring the conditions w.r.t. the posterior distribution or by requiring the conditions to hold with high probability. Moreover, we are intrigued by non-stationary settings—where e.g., rewards follow a Markov process—since the planner would be able to sample *a priori* inferior arms with high probability assuming the rewards change fast enough, thereby reducing regret.

Acknowledgements

We thank the participants of the Computational Data Science seminar at Technion – Israel Institute of Technology and the participants of Young Researcher Workshop on Economics and Computation for their comments and suggestions. Additionally, we thank ICML 2020 anonymous reviewers who provided comments that improved the manuscript. The work of G. Bahar, O. Ben-Porat and M. Tennenholtz is funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement n° 740435). The work of K. Leyton-Brown is funded by the NSERC Discovery Grants program, DND/NSERC Discovery Grant Supplement, Facebook Research and Canada CIFAR AI Chair Amii. Part of this work was done while K. Leyton-Brown was a visiting researcher at Technion – Israel Institute of Science and was partially funded by the European Union’s Horizon 2020 research and innovation programme (grant

agreement n° 740435).

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pp. 2312–2320, 2011.
- Agrawal, S. and Goyal, N. Analysis of thompson sampling for the multi-armed bandit problem. In *Proceedings of the 25th Annual Conference on Learning Theory (COLT)*, pp. 1–39, 2012.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pp. 322–331. IEEE, 1995.
- Bahar, G., Smorodinsky, R., and Tennenholtz, M. Economic recommendation systems: One page abstract. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC ’16, pp. 757–757, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3936-0. doi: 10.1145/2940716.2940719. URL <http://doi.acm.org/10.1145/2940716.2940719>.
- Bahar, G., Smorodinsky, R., and Tennenholtz, M. Social learning and the innkeeper challenge. In *ACM Conf. on Economics and Computation (EC)*, 2019.
- Barto, A. G., Bradtke, S. J., and Singh, S. P. Learning to act using real-time dynamic programming. *Artificial Intelligence*, 72(1-2):81–138, 1995.
- Ben-Porat, O. and Tennenholtz, M. A game-theoretic approach to recommendation systems with strategic content providers. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pp. 1118–1128, 2018.
- Besbes, O., Gur, Y., and Zeevi, A. Stochastic multi-armed-bandit problem with non-stationary rewards. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 199–207, 2014.
- Bubeck, S., Cesa-Bianchi, N., et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, learning, and games*. Cambridge Univ Press, 2006.
- Chapelle, O. and Li, L. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pp. 2249–2257, 2011.

- Chen, B., Frazier, P., and Kempe, D. Incentivizing exploration by heterogeneous users. In Bubeck, S., Perchet, V., and Rigollet, P. (eds.), *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 798–818. PMLR, 06–09 Jul 2018. URL <http://proceedings.mlr.press/v75/chen18a.html>.
- Chu, W., Li, L., Reyzin, L., and Schapire, R. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 208–214, 2011.
- Cohen, L. and Mansour, Y. Optimal algorithm for bayesian incentive-compatible. In *ACM Conf. on Economics and Computation (EC)*, 2019.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science conference (ITCS)*, pp. 214–226. ACM, 2012.
- Frazier, P., Kempe, D., Kleinberg, J., and Kleinberg, R. Incentivizing exploration. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pp. 5–22, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2565-3. doi: 10.1145/2600057.2602897. URL <http://doi.acm.org/10.1145/2600057.2602897>.
- Garcia, J. and Fernández, F. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- Hardt, M., Price, E., Srebro, N., et al. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 3315–3323, 2016.
- Immorlica, N., Mao, J., Slivkins, A., and Wu, Z. S. Bayesian exploration with heterogeneous agents, 2019.
- Joseph, M., Kearns, M., Morgenstern, J. H., and Roth, A. Fairness in learning: Classic and contextual bandits. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 325–333. Curran Associates, Inc., 2016.
- Karnin, Z., Koren, T., and Somekh, O. Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, pp. 1238–1246, 2013.
- Kremer, I., Mansour, Y., and Perry, M. Implementing the wisdom of the crowd. *Journal of Political Economy*, 122: 988–1012, 2014.
- Levine, N., Crammer, K., and Mannor, S. Rotting bandits. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 3074–3083. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6900-rotting-bandits.pdf>.
- Liu, L., Dean, S., Rolf, E., Simchowitz, M., and Hardt, M. Delayed impact of fair machine learning. In *International Conference on Machine Learning*, pp. 3156–3164, 2018.
- Liu, Y., Radanovic, G., Dimitrakakis, C., Mandal, D., and Parkes, D. C. Calibrated fairness in bandits, 2017.
- Mansour, Y., Slivkins, A., and Syrgkanis, V. Bayesian incentive-compatible bandit exploration. In *ACM Conf. on Economics and Computation (EC)*, 2015.
- Nisan, N. and Ronen, A. Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM Symposium on Theory of Computing (STOC)*, pp. 129–140. ACM, 1999.
- Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge, 2007.
- Slivkins, A. Contextual bandits with similarity information. *The Journal of Machine Learning Research*, 15(1):2533–2568, 2014.