

Bandit Learning with Biased Human Feedback

Wei Tang

Washington University in St. Louis
w.tang@wustl.edu

Chien-Ju Ho

Washington University in St. Louis
chienju.ho@wustl.edu

ABSTRACT

We study a multi-armed bandit problem with biased human feedback. In our setting, each arm is associated with an unknown reward distribution. When an arm is played, a user receives a realized reward drawn from the distribution of the arm. She then provides feedback, a biased report of the realized reward, that depends on both the realized reward and the feedback history of the arm. The learner can observe only the biased feedback but not the realized rewards. The goal is to design a strategy to sequentially choose arms to maximize the total rewards users receive while only having access to the biased user feedback. We explore two natural feedback models. When user feedback is biased only by the average feedback of the arm (i.e., the ratio of positive feedback), we demonstrate that the evolution of the average feedback over time is mathematically equivalent to users performing online gradient descent for some latent function with a decreasing step size. With this mathematical connection, we show that under some mild conditions, it is possible to design a bandit algorithm achieving *regret* (i.e., the difference between the algorithm performance and the optimal performance of always choosing the best arm) sublinear in the number of rounds. However, in another model when user feedback is biased by both the average feedback and the number of feedback instances, we show that there exist no bandit algorithms that could achieve sublinear regret. Our results demonstrate the importance of understanding human behavior when applying bandit approaches in systems with humans in the loop.

CCS CONCEPTS

• **Theory of computation** → **Sequential decision making**; • **Human-centered computing** → **User models**;

KEYWORDS

multi-armed bandits; herding effects; social learning

ACM Reference Format:

Wei Tang and Chien-Ju Ho. 2019. Bandit Learning with Biased Human Feedback. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

In a multi-armed bandit problem, a learner sequentially selects from a set of arms. Each arm is associated with some unknown reward distribution. After selecting an arm, the learner observes the realized reward for the selected arm. The goal of the learner is to maximize the total rewards obtained from selected arms over

time. The performance of the bandit algorithm is often measured in terms of *regret*, defined as the difference between the algorithm performance and the performance of an oracle which can select the best arm in hindsight. The multi-armed bandit formulation provides a theoretical framework for resolving the classical exploration-exploitation tradeoffs in online decision problems under uncertainty. Therefore, multi-armed bandits have been studied in a wide range of applications in various domains, such as medical trials, online auctions, or web advertisements.

We explore the applications of bandits settings to human-in-the-loop systems. For example, consider user-generated content platforms, such as Youtube, Quora, or Stack Exchange. On these platforms, content qualities vary across a wide spectrum. Ideally, the platform would like to select the best content to display to users to optimize users' experience. However, the content qualities are often not known in advance, and the platform needs to learn the content qualities through user feedback (e.g., number of likes, upvotes, etc). This naturally leads to a bandit problem, in which the platform needs to balance exploration (display content with fewer feedback instances to users to acquire more information) and exploitation (display content with higher empirical ratings to optimize users' happiness), as studied in the literature [13, 19].

Many challenges arise when humans are involved in the bandit learning process. In recent years, researchers have addressed various strategic issues brought up by humans involved in bandit learning [13, 19, 20, 24]. However, in these works, it is assumed that users' feedback is *unbiased* in representing the reward of selecting an arm (e.g., in user-generated content platforms, users' average ratings are used as the estimates for content qualities). On the other hand, as the empirical studies suggest [21, 27, 28], user feedback is often biased by other users' feedback. For example, users have the tendency to provide feedback that agrees with the majority opinion even if their experience disagrees (i.e., the herding effect). These empirical evidences suggest a different stochastic model in that each observed feedback instance might be biased by the feedback history. Moreover, this biased user feedback introduces additional challenges. Since user feedback only represents biased reports of the realized rewards, suppose the goal of the platform is to maximize the total rewards over time (which may be interpreted as the overall user experience), can a platform achieve sublinear regret from only observing biased feedback?

In this paper, we study a variant of the multi-armed bandit problem with human biased feedback. In our setting, the learner/platform only observes human-generated feedback instead of the realized reward when selecting an arm. The human feedback depends on both the realized reward and other users' feedback for the selected arm. The goal of the learner is to maximize the total realized rewards for the selected arms while only having access to biased human-generated feedback.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada. © 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

To address the issues of user biased feedback, we explore two natural user feedback models and study their impacts to the design of bandit algorithms. The first model, avg-herding feedback model, assumes that user feedback for an arm depends on the realized reward and the *average feedback* (i.e., the ratio of positive feedback) of the arm so far. We show that, under this model, the dynamics of user feedback over time is mathematically connected to asymptotic approximation [26]. In particular, the average feedback changes over time as if users are performing online gradient descent on a latent function with a decreasing step size. With this mathematical connection, we characterize the convergence and convergence rates for the average feedback of an arm under some mild conditions. These convergence results enable us to design a bandit algorithm based on UCB (Upper Confidence Bound) algorithm and achieve sublinear regret.

While the results on the first model are promising, our results on another natural model, beta-herding feedback model, paint a very different picture. In this model, user feedback is biased by not only the average feedback in the past, but also the number of feedback instances the arm has received so far. This model captures a natural scenario that users might be biased more heavily if there exists more feedback instances in the history. We show that, under this model, the average feedback of an arm converges to a random variable with non-zero variance. This implies that, even with an infinitely number of feedback instances for the arm, the learner is not able to infer the expected reward of the arm through observing the average user feedback. We further show that, using arguments from information theory, there exist no bandit algorithms that can achieve sublinear regret when user feedback follows beta-herding feedback model.

We next present a toy example to demonstrate that it is possible to get around the above impossible result by modifying the information structure to break the assumption that users follow beta-herding feedback model. In particular, if the learner is allowed to hide the historical information from a small portion of the users, under some styled user models on how users respond to information structures, it is possible to design an algorithm achieving sublinear regret. This result opens up a potentially interesting line of future research: Can the learner adaptively *design* the information structures to improve the overall utility?

Our results demonstrate the importance of understanding human behavior when learning from human generated feedback. A small deviation on the user behavior model and/or the design of the information structure could have significantly impacts on the overall system outcome. Therefore, platforms and decision makers should carefully take these into account when designing learning algorithms in systems with humans in the loop.

2 RELATED WORK

In this section, we review the relevant literature in multi-armed bandit problems, recent studies on human-in-the-loop bandit learning, and the literature on social influences and social learning that share similar motivations of this work.

Multi-armed bandit problems. Our work is a variant of the well-studied multi-armed bandit problem [18]. Bandit problems traditionally assume the rewards generated by each arm at each round

are directly observable, and the research focus has been divided into settings in which rewards are either independent and identically distributed (i.i.d.) [3] or adversarial [2, 4]. There exist other works that assume rewards are neither i.i.d. drawn nor adversarial. For example, bandits with Markovian rewards [22, 23] assume the state of each arm evolves according to a Markov process. Other non-stationary bandit problems [6, 11] consider the setting in which the rewards distribution might change over time, independent of previous actions. More recently, researchers have addressed the setting in which the rewards are strategic choices of humans and could be influenced by how the bandit algorithm is designed [13, 19]. Our work differs from the above works in that, in our setting, the “state” (history information) of each arm is correlated with learner’s actions and there might be infinitely many states. Moreover, in our setting, the algorithm cannot observe realized rewards but only has access to biased feedback while previous work assume the realized rewards are observable.

Human-in-the-loop bandit learning. Recently, there have been works exploring bandit learning with humans in the loop [9, 17, 20, 24]. In the setting of these works, the learner cannot directly choose which arms to play. Instead, at each time step, a myopic agent, who only aims to maximize her own reward at the single time step she is involved in, chooses which arm to play. Since the agent only cares about her instant payoff, she does not have incentives to explore and tends to always exploit, and this collective arm playing will lead to the convergence to the suboptimal arm. Researchers have been attempting to address this problem by considering different ways of *persuading* agents to perform exploration, including offering agents payments to perform exploration [9] or utilizing information asymmetry to lead agents to explore by designing what information to show to each agent [17, 20, 24]. The idea of utilizing information asymmetry to persuade agents is similar to Bayesian Persuasion [16] in economics. The above line of work has focused on settings in which humans are involved in *arm selection*, i.e., which arm is played in each round. In this work, we focus on a parallel aspect of human involvements, in which humans are involved in *feedback generation*.

Social influences and social learning. Our feedback models are motivated by the empirical evidences that users’ decisions are influenced by not only their own experience but also other users’ decisions [21, 27, 28]. For example, Muchnik et al. [21] empirically show that, a post on a forum is more likely to receive positive feedback (i.e., *upvotes*) if the platform insert an upvote right after the post is made. Similar discussion also appears in the social learning literature in economics [5, 7, 29]. They discuss the setting in which users’ decisions might be influenced by other users’ decisions. Therefore, under certain conditions, users might collectively make the bad decision since they might follow what other users do regardless of what they privately know. In prior work, there is not much discussion on either the convergence rate of users’ aggregate behavior or the impacts on the system designer’s perspective. In this work, we focus on deriving the dynamics of user feedback over time and explore the impacts on the design of bandit algorithms.

3 PROBLEM SETTING

Let K be the number of arms. Each arm $k \in [K] = \{1, \dots, K\}$ is associated with an unknown quality $\theta_k \in [0, 1]$. Let $I^* = \operatorname{argmax}_k \theta_k$ and $\theta^* = \theta_{I^*}$ be the index of the best arm and the associated highest expected quality. At each round t , a user randomly drawn from the population arrives, the learner selects an arm $I_t \in \{1, \dots, K\}$ for the arriving user. The user then gets a binary reward Z_t (positive or negative experience) with mean θ_{I_t} .

$$Z_t \sim \text{Bernoulli}[\theta_{I_t}]$$

The reward is not observable to the learner. However, after receiving the reward, each user provides a binary feedback $X_t \in \{0, 1\}$ about this arm. The goal of the learner is to maximize the total rewards users receive while observing only the (potentially biased) feedback. Note that when the feedback is the same as the realized reward, i.e., $X_t = Z_t$ for all t , this problem reduces to standard bandit setting. Below we describe the user feedback models, i.e., how X_t is generated.

User feedback models

Users' feedback depends on both the realized rewards and the feedback history of the arms. The feedback history of arm k up to time t can be summarized by $n_{k,t}$ and $\rho_{k,t}$, which represent the number of feedback instances and the ratio of positive feedback for arm k up to round t . We assume $n_{k,0} = \rho_{k,0} = 0$ to simplify the presentation, however, our results can be easily extended to settings with non-zero $n_{k,0}$ and $\rho_{k,0}$, which can be used to represent the users' *prior* of the arm quality. Again, if users provide unbiased feedback, we should have $X_t = Z_t$ for all t .

In this paper, we model the feedback generation as a stochastic process. We define a feedback function to model the probability of obtaining positive feedback for an arm from a user randomly drawn from the population. Note that a feedback function describes the characteristics of the *user population* the platform is interacting with instead of a single specific user. In particular, we introduce $\text{Feedback}(\theta, \rho, n)$ to model the probability of obtaining positive feedback from a user given that the arm quality is θ and the history information of the arm is summarized by its average feedback ρ and the number of feedback instances n .

As a special case, when $\text{Feedback}(\theta, \rho, n) = \theta$, user feedback represents unbiased samples of the arm quality.

In this paper, we explore two natural feedback models.

- Avg-herding feedback model:
In this feedback model, user feedback is biased by the average feedback of the arm. In particular, the feedback function has the form

$$\text{Feedback}(\theta, \rho, n) = F(\theta, \rho).$$

In Section 4, we study the stochastic process of user feedback specified by a general set of feedback functions F . We then discuss the impacts of this stochastic feedback generation on the design of bandit algorithms.

- Beta-herding feedback model:
In this feedback model, user feedback is biased by the average feedback and the number of feedback instances. In particular, we consider a natural setting and assume users

update their beliefs about the arm quality in a Bayesian manner. Users treat the historical ratings as the prior signals of the arm quality and update the posterior based on their own experience. They then provide feedback according to their posterior.

We introduce a factor $m \geq 0$, which can be interpreted as the weights users put on their own experience. When the arm quality is θ and the arm history is (n, ρ) , the expected number of *positive signals* the user will obtain is $m\theta + n\rho$, where the first term is the expected positive signals the users receive from their own experience (i.e., arm quality multiplied by the weight) and the second term is the number of positive signals from other users. The total number of signals is $m+n$.

Therefore, the probability of obtaining positive feedback for arm k at round t can be written as

$$\text{Feedback}(\theta, \rho, n) = \frac{m\theta + n\rho}{m+n}. \quad (1)$$

Note that when $m \rightarrow \infty$, user feedback provides unbiased samples of the arm quality.

Regret notions

The goal of the learner is to maximize the sum of rewards users receive over time. Let \mathcal{A} be the algorithm the learner deploys and $\{I_t\}$ are the arms selected by \mathcal{A} . We define the *regret* as $R_{\mathcal{A}}(T)$.

$$\mathbb{E}[R_{\mathcal{A}}(T)] = T\theta^* - \mathbb{E}_{\mathcal{A}} \left[\sum_{t=1}^T \theta_{I_t} \right],$$

where the expectation is taken over the randomness of the reward realization and the algorithm. In particular, we are interested in the region of $T \rightarrow \infty$ and aim to understand under what conditions we can achieve asymptotic sublinear regret, i.e., $\mathbb{E}[R(T)] = o(T)$, when user feedback is biased by historical feedback.

4 BANDITS WITH AVG-HERDING FEEDBACK MODEL

In this section, we explore the bandit learning problem when user feedback follows avg-herding feedback model. We first derive the stochastic process of the feedback generation for a single arm and characterize the convergence and convergence rate of users' average feedback over time. We then discuss how this user feedback model impacts the design and analysis of bandit algorithms.

4.1 Stochastic process of feedback generation

In the following discussion, we explore the feedback dynamics of a single arm, i.e., the stochastic process of feedback generation. We omit the arm's index k in the subscript when it is clear from the context. Also, since user feedback is biased by the history of only the selected arm, to simplify the presentation, we consider the case that the same arm is repeatedly selected and therefore $n_t \equiv t$ when studying the stochastic process for a single arm.

4.1.1 Connection to stochastic approximation. Recall that in avg-herding feedback model, when the quality of the arm is θ and the average feedback of the arm is ρ , the probability for a user to provide a positive feedback is $F(\theta, \rho)$. The stochastic process of the feedback dynamics can be expressed as follows: at the $(t +$

1)-th round, the feedback X_{t+1} provided by the user is drawn randomly from a Bernoulli distribution: $\text{Bernoulli}[F(\theta, \rho_t)]$. The history information of the arm (n_{t+1}, ρ_{t+1}) are updated based on the realized feedback.

As mentioned, we simplify the presentation by setting $n_t \equiv t$. Therefore, we focus on how ρ_t evolves over time. By simple weighted averaging, we have

$$\rho_{t+1} = \frac{t}{t+1}\rho_t + \frac{1}{t+1}X_{t+1} = \rho_t - \frac{1}{t+1}(\rho_t - X_{t+1}).$$

Define the noise term $\xi_t = \mathbb{E}[X_t | \mathcal{F}_{t-1}] - X_t = F(\theta, \rho_{t-1}) - X_t$, where $\mathcal{F}_t = \sigma(\{X_t\}_{t \geq 1})$ is the filtration of the stochastic process. It is easy to see that $\mathbb{E}[\xi_t | \mathcal{F}_{t-1}] = 0$. Also let $\eta_t = 1/t$ be the step size (learning rate). We can rewrite the above recursive definition as an update rule in stochastic approximation [10, 26].

$$\rho_{t+1} = \rho_t - \eta_{t+1}(\rho_t - F(\theta, \rho_t) + \xi_{t+1}) \quad (2)$$

In particular, suppose there exists a latent function $G(\theta, \rho)$, such that $\partial G / \partial \rho = \rho - F(\theta, \rho)$, then Equation (2) is equivalent to the update rule for stochastic gradient descent with step size η_{t+1} :

$$\rho_{t+1} = \rho_t - \eta_{t+1}(\nabla_\rho G(\theta, \rho_t) + \xi_{t+1})$$

With this observation, the stochastic process of the average feedback updates is equivalent to users collectively performing stochastic gradient descent for a latent function G with a decreasing step size. Below we utilize this mathematical connection and discuss conditions on the convergence and convergence rates of the average feedback ρ . We then discuss the impacts of this stochastic process on the design and analysis of bandit algorithms.

4.1.2 On the convergence and convergence rate of $\lim_{t \rightarrow \infty} \rho_t$.

We first specify the assumptions needed to establish the asymptotic behavior of the limit of average feedback.

- A1. $F(\theta, \rho)$ is strictly increasing in θ and non-decreasing in ρ ;
- A2. $F(\theta, \rho)$ is differentiable and L_F^ρ -Lipschitz continuous with respect to ρ .

A1 implies that, conditional on the same quality (average feedback), an arm with better average feedback (quality) receives more positive feedback in expectation. A2 assumes the improvement is smooth with respect to ρ . While the differentiable property of F can be satisfied if the population is large and smooth, we note that the differentiable property is only for analytical convenience. Our results still hold even if F is only continuously differentiable in some local neighbourhood of equilibrium points.

We would also like to note that these two assumptions are relatively mild. As an example, below we give a general set of feedback functions F that satisfy the above assumptions.

Example 4.1. Consider the following set of feedback functions: $F(\theta, \rho) = w_1\theta + w_2\rho$, for any $w_1, w_2 \geq 0$ and $w_1 + w_2 = 1$. This set of feedback functions satisfies both of the assumptions. It also has very natural interpretations. In particular, it specifies that, the probability of receiving a positive feedback from a random user (drawn from the population) $F(\theta, \rho)$ is the weighted average of the arm quality θ and other users' average feedback ρ .

Armed with the above assumptions, we can formally characterize the convergence of ρ_t .

LEMMA 4.2. *Let $\mathcal{S}_\theta := \{\rho : \rho - F(\theta, \rho) = 0\}$. We have $\mathbb{P}(\lim_{t \rightarrow \infty} \rho_t \in \mathcal{S}_\theta) = 1$.*

The above lemma demonstrates that ρ_t converges to one of the points in a set \mathcal{S}_θ and characterizes the points in \mathcal{S}_θ . Recall that the latent function $G(\theta, \rho)$ satisfies $\partial G / \partial \rho = \rho - F(\theta, \rho)$. Therefore, the lemma illustrates that the average feedback will converge to one of the points in \mathcal{S}_θ , the set of the local optimal points for the latent function G . This intuition suggests that, when the latent function G is strongly convex, since there exists only one local optimal point (which is the global optimal), we should be able to show that ρ_t will almost surely converge to the global optimal.

Moreover, the convexity of G is correlated with the value of the Lipschitz constant L_F^ρ . In particular, when $L_F^\rho < 1$, by definition, we have $\nabla_\rho F(\theta, \rho) < 1$ for all θ and ρ . Since $\nabla_\rho^2 G(\theta, \rho) = 1 - \nabla_\rho F(\theta, \rho)$, when $L_F^\rho < 1$, $\nabla_\rho^2 G(\theta, \rho) > 0$ for all θ and ρ . Therefore, G is strongly convex when $L_F^\rho < 1$. Below we formally characterize the convergence of ρ_t when G is strongly convex.

COROLLARY 4.3. *Given $L_F^\rho < 1$, i.e., G is strongly convex, there exists a unique ρ^* that satisfies $\rho^* - F(\theta, \rho^*) = 0$, such that $\mathbb{P}(\lim_{t \rightarrow \infty} \rho_t = \rho^*) = 1$.*

Next we provide the results on the convergence rate of ρ_t and focus on the case when G is strongly convex. In particular, we introduce $\bar{\lambda} > 0$, such that $\nabla_\rho^2 G \geq \bar{\lambda} > 0$.

THEOREM 4.4. *Given $L_F^\rho < 1$, i.e., G is strongly convex. $\forall \epsilon > 0$, we have,*

$$\mathbb{P}(|\rho_t - \rho^*| \geq \epsilon) \leq \exp\left(-\frac{(\epsilon - \epsilon_t)^2}{2 \sum_{i=1}^t L_i}\right),$$

where $L_i = \eta_i^2 (\prod_{j=i}^{t-1} (\eta_{j+1}^2 (L_F^\rho - 1)^2) - 2\bar{\lambda}\eta_{j+1} + 1)$,

$\epsilon_t = \exp(-\bar{\lambda}S_t)|\rho_0 - \rho^*| + \sqrt{\sum_{i=0}^{t-1} \eta_{i+1}^2 \exp(-2\bar{\lambda}(S_t - S_{i+1}))}$,

and $S_t = \sum_{i=1}^t \eta_i$.

PROOF SKETCH. In the proof, we decompose $|\rho_t - \rho^*|$ into two terms, with the first term corresponding to the *empirical iterate error*, i.e., the difference between the absolute value of error at a given time and its mean, and the second term corresponding to the *expectation error*:

$$|\rho_t - \rho^*| = (|\rho_t - \rho^*| - \mathbb{E}[|\rho_t - \rho^*|]) + \mathbb{E}[|\rho_t - \rho^*|] \quad (3)$$

To derive the probabilistic tail bound for the $|\rho_t - \rho^*|$, we bound the empirical iterate error $|\rho_t - \rho^*| - \mathbb{E}[|\rho_t - \rho^*|]$ and the expectation error $\mathbb{E}[|\rho_t - \rho^*|]$ separately using martingale concentration bounds. \square

REMARK 1. *We would like to offer a few observations to help interpret the convergence bound¹. In particular,*

- when $t \rightarrow \infty$, $\epsilon_t \rightarrow 0$,
- when $\bar{\lambda} \in (0, 1/2)$, $\sum_{i=1}^t L_i = O(t^{-2\bar{\lambda}})$, and
- when $\bar{\lambda} \in [1/2, \infty)$, $\sum_{i=1}^t L_i = O(1/t)$.

So we can characterize the bound in two regions based on whether $\bar{\lambda} \geq 1/2$. As a special case, when user feedback is unbiased, i.e., $F(\theta, \rho) = \theta$, we have $\bar{\lambda} = 1$, and the bound reduces to $\mathbb{P}(|\rho_t - \rho^*| \geq$

¹The detailed derivations are included in the appendix of the full paper.

$\epsilon) \leq O(e^{-\epsilon^2 t})$, the same as the standard Chernoff bound. Moreover, in our setting, since $F(\theta, \rho)$ is non-decreasing in ρ , i.e., $\nabla_\rho F \geq 0$. We have $\nabla_\rho^2 G = 1 - \nabla_\rho F \leq 1$. Therefore, while our bound holds for the region $\bar{\lambda} \in (0, \infty)$, in our setting, we focus on the region $\bar{\lambda} \in (0, 1]$.

Note that in this theorem, the convergence rate is a function of $\bar{\lambda}$, which is the property of the function G (hence the property of the feedback model F). As an intuitive interpretation, recall that $\nabla_\rho^2 G \geq \bar{\lambda}$ and $\nabla_\rho G = \rho - F(\theta, \rho)$. Therefore, small $\bar{\lambda}$ implies large $\partial F / \partial \rho$, which means user's feedback is influenced more by the other users' feedback and relatively less by the arm quality. When users' feedback depends less on the arm quality, it requires more feedback to infer the arm quality, and therefore the convergence is slower. This intuition aligns with the theorem, in which smaller $\bar{\lambda}$ leads to a slower convergence rate.

4.2 Designing bandit algorithms

Given the convergence bound in Theorem 4.4, we can design a UCB-like algorithm that achieves sublinear regret. We assume the learner has knowledge of the feedback model F . Note that since F models the behavior of feedback generation for the *user population* the platform is interacting with, this assumption only requires the platform to have knowledge of the population instead of any particular users².

In each round of our algorithm, the learner maintains an estimator $\hat{\theta}_{k,t}$ of arm k 's quality from the observation of average feedback $\rho_{k,t}$. From Lemma 4.2, an asymptotically unbiased and consistent estimator of arm's quality $\hat{\theta}_{k,t}$ can be obtained by solving the following equation.

$$\hat{\theta}_{k,t} = \max\{\min\{\hat{\theta}_{k,t} : F(\hat{\theta}_{k,t}, \rho_{k,t}) = \rho_{k,t}\}, 1, 0\} \quad (4)$$

Intuitively, the solutions of the above equation represent the set of local optimal points of G . Moreover, we can show that the estimator $\hat{\theta}_{k,t}$ is unique for every $\rho_{k,t}$ if $A1$ is satisfied.

LEMMA 4.5. *Suppose $A1$ is satisfied, for any $\rho_{k,t}$, there exists a unique $\hat{\theta}_{k,t}$ that satisfies Equation (4).*

Given the convergence bounds and the estimator $\hat{\theta}_{k,t}$, we are ready to describe our proposed UCB-like algorithm Avg-UCB, as specified in Algorithm 1. The key differences to the standard UCB algorithms are that: First, we maintain a quality estimate $\hat{\theta}_{k,t}$ for each arm k at each time t by solving Equation (4) instead of using empirical average feedback. Second, the confidence interval in the UCB index is derived from the convergence rates as specified in Theorem 4.4. Our algorithm takes as input parameters β and $\bar{\lambda}$. β plays a similar role as the constant in UCB confidence radius to balance exploration and exploitation. $\bar{\lambda}$ is the parameter of the problem instance. Note that our algorithm only requires to find some $\bar{\lambda}$ such that $\nabla_\rho^2 G \geq \bar{\lambda}$.

The following theorem gives the regret bound for the algorithm Avg-UCB.

²In practice, this assumption can be approximately satisfied through market research or behavioral experiments, which study the connection between users' real experience (i.e., Z_t) and reported feedback (i.e., X_t). Moreover, our results are robust to small estimation noises of F .

Algorithm 1 Avg-UCB for Avg-Herding Feedback Model

```

1: Input:  $\beta, \bar{\lambda}, K$ .
2: Initializations: first  $K$  rounds, play each arm once
3: for  $t = K + 1, \dots, T$  do
4:   for each  $k \in \{1, \dots, K\}$  do
5:     Compute  $\hat{\theta}_{k,t-1}$  from (4).
6:      $\text{UCB}_{k,t} = \hat{\theta}_{k,t-1} + \sqrt{\frac{\beta \ln(t-1)}{n_{k,t-1}^{\min\{1, 2\bar{\lambda}\}}}}$ .
7:   Choose arm  $I_t \in \operatorname{argmax}_{k=1, \dots, K} \text{UCB}_{k,t}$ .
8:   (Ties are broken in some consistent way)
9:   Receive feedback  $X_t$ .
10:   $\rho_{I_t,t} \leftarrow (\rho_{I_t,t-1} \times n_{I_t,t-1} + X_t) / (n_{I_t,t-1} + 1)$ 
11:   $\rho_{k,t} = \rho_{k,t-1}, \forall k \neq I_t$ .
12:   $n_{I_t,t} \leftarrow n_{I_t,t-1} + 1$ 
13:   $n_{k,t} \leftarrow n_{k,t-1}, \forall k \neq I_t$ .

```

THEOREM 4.6. *Suppose $A1$ and $A2$ are satisfied and $L_F^\rho < 1$. Let $\bar{\lambda}' = \max\{1, 1/(2\bar{\lambda})\}$, $\Delta_k = \theta^* - \theta_k$. With appropriately chosen β^3 the expected regret for Avg-UCB is bounded by:*

$$\mathbb{E}[R(T)] \leq \sum_{k \neq I^*} \Delta_k (4 \ln T / (C \Delta_k^2))^{\bar{\lambda}'} + K \pi^2 / 6,$$

where C is a constant that is dependent on the properties of feedback function F .

We introduce an additional notion $\bar{\lambda}' = \max\{1, 1/(2\bar{\lambda})\}$ to simplify the presentation due to the different convergence rates on whether $\bar{\lambda} < 1/2$ as discussed in Remark 1. Similar to the discussion on the convergence rate, the dependency of the above upper regret bound on $\bar{\lambda}'$ implies that it is harder to learn the quality of an arm if users are biased more by the historical information rather than the arm quality.

The above regret bound is a gap-dependent bound. In particular, let $\Delta_{\min} = \min_{k:k \neq I^*} \Delta_k$. The regret bound can be written as: $\mathbb{E}[R(T)] = O\left(\frac{(\ln T)^{\bar{\lambda}'}}{\Delta_{\min}^{2\bar{\lambda}'-1}}\right)$. Observe that $\lim_{T \rightarrow \infty} \mathbb{E}[R(T)]/T \rightarrow 0$ for any $\bar{\lambda}' > 0$. Therefore, the algorithm achieves sublinear regret as long as G is strongly convex (i.e., $\bar{\lambda}' > 0$).

Moreover, we can derive gap-independent bounds from the above bound. When $\bar{\lambda} \geq 1/2$ (which includes the unbiased feedback setting with $\bar{\lambda} = 1$), we can show that $\mathbb{E}[R(T)] = O(\sqrt{T \ln T})$, which matches the standard regret bound without biased feedback.

What if G is not convex. Our algorithm relies on the assumption that the latent function G is convex, i.e., $L_F^\rho < 1$. This assumption implies that users' feedback is not influenced too heavily by the change of feedback history. While this assumption seems mild, it is natural to wonder whether we can obtain similar results when G is not convex.

We would like to note that even in settings when G is non-convex, the statements of Lemma 4.2 and 4.5 still hold. This means the average user feedback for each arm will still converge to some point, and we can infer the arm quality from the converged average

³The choice of β depends on the parameters of $F(\theta, \rho)$. The detailed derivation is tied with the proof and is included in the appendix of the full paper.

feedback. The main obstacle to overcome is to derive the convergence rate as in Theorem 4.4. This problem is challenging as it is equivalent to deriving the convergence rate of optimization for non-convex functions. Recently, There have been recent works focusing on deriving the convergence rates for non-convex optimization in different settings [1, 12]. As long as one could characterize the convergence rate of ρ_t even non-convex G , our bandit strategy can be adapted to generate a sublinear regret strategy (by changing the “confidence interval” in the UCB index based on the derived convergence rate).

5 BANDITS WITH BETA-HERDING FEEDBACK MODEL

In the previous section, we explore avg-herding feedback model, in which user feedback is biased only by the average feedback of the selected arm. We show that, under some mild conditions, the average feedback for an arm almost surely converges to some value, and we can infer the arm quality from the average feedback, and therefore we can design a UCB-like algorithm for achieving sublinear regret.

However, in some scenarios, user feedback may be biased by not only the average feedback but also the number of feedback instances of the arm. In this section, we explore another natural feedback model, beta-herding feedback model, and prove impossibility results. In particular, we assume users give feedback in a Bayesian manner. They treat the feedback history as the prior, i.e., for an arm with history (n, ρ) , there are $n\rho$ positive signals and $n(1 - \rho)$ negative signals for the arm. After they experience the binary reward (drawn according to the arm’s quality distribution), they update their posterior by treating their experience as m signals and then provide feedback according to the posterior. Therefore, in expectation, the probability for them to provide positive feedback for an arm with quality θ and history (n, ρ) is $\text{Feedback}(\theta, \rho, n) = (m\theta + n\rho)/(m + n)$.

5.1 Stochastic process of feedback generation

The first natural attempt is to replace $F(\theta, \rho_t)$ with $\text{Feedback}(\theta, \rho, n)$ in Equation (2) and apply similar analysis using stochastic approximation. However, when $\text{Feedback}(\theta, \rho, n)$ follows beta-herding feedback model, one can not directly apply this approach. Briefly speaking, the update rule in Equation (2) aims to find the equilibrium points of the feedback function. However, in beta-herding feedback model, the feedback function is changing over time, and it is not trivial whether the converged points satisfy the set of properties as derived with avg-herding feedback model.

Instead, we make the observation that the stochastic process of beta-herding feedback model is similar to the urn process [14]. We utilize the property of *exchangeability* for the feedback history to give the characterization of ρ_t process. Below we formally characterize the stochastic process of ρ_t with beta-herding feedback model.

LEMMA 5.1. *Consider the stochastic process in Equation (2) with the feedback model described in Equation (1), $\lim_{t \rightarrow \infty} \rho_t$ converges almost surely to a random variable specified by a beta distribution.*

In particular,

$$\lim_{t \rightarrow \infty} \rho_t \sim \text{Beta}(m\theta, m(1 - \theta)).$$

PROOF SKETCH. Let $S_t = \sum_{i=1}^t x_i$, where x_i is the realization of the feedback random variable X_i . It is easy to show that the sequence of random variables X_i satisfies the exchangeable property. Therefore, the probability that $S_t = l$, i.e., there are l positive feedback among t feedback, can be written as

$$\frac{\prod_{i=0}^{l-1} (m\theta + i) \cdot \prod_{j=0}^{t-1-l} (m(1 - \theta) + j)}{\prod_{i=0}^{t-1} (m + i)}.$$

By Stirling’s approximation, when $t \rightarrow \infty$, we have

$$\mathbb{P}(S_t = l) = \frac{l^{m\theta-1}}{B(m\theta, m(1-\theta))} \cdot (t-l)^{m(1-\theta)-1} \cdot t^{1-m},$$

where $B(\cdot)$ is the Beta function. Denote $l = \rho t$ for some $0 < \rho < 1$,

$$\mathbb{P}\left(\frac{S_t}{t} \leq \rho\right) = \sum_{i=0}^{\lfloor t\rho \rfloor} \mathbb{P}\left(\frac{S_t}{t} = \frac{i}{t}\right)$$

When $t \rightarrow \infty$, the summation can be written as an integral

$$\mathbb{P}\left(\frac{S_t}{t} \leq \rho\right) = t \int_0^\rho \mathbb{P}\left(\frac{S_t}{t} = u\right) du.$$

Plug in the above $\mathbb{P}(S_t = l)$ expression and replace l with ρt ,

$$\mathbb{P}\left(\frac{S_t}{t} \leq \rho\right) = \frac{1}{B(m\theta, m(1-\theta))} \int_0^\rho u^{m\theta-1} (1-u)^{m(1-\theta)-1} du,$$

which is the CDF of the beta distribution. This completes the proof. \square

Note that when the feedback is unbiased, i.e., when $m \rightarrow \infty$, the beta distribution will shrink to a Dirac delta function which has the point mass exactly in θ .

5.2 The impossibility result

In this section, we show that there exist no bandit algorithms that achieve sublinear regret if user feedback follows beta-herding feedback model.

Lemma 5.1 implies that, even if we obtain an infinite number of feedback instances for an arm, we cannot accurately infer the arm quality with high probability from the empirical average feedback ρ_∞ . A natural next question to ask is, if we take into account all the feedback generated in the process, whether it is possible to infer the true arm quality. Below we use the notion of Fisher information to answer the question. In short, Fisher information provides a way to quantify the amount of information about the latent parameter θ we can obtain for observing each sample of a random variable X_i . Since Fisher information is additive, we can show that,

LEMMA 5.2. *Consider the stochastic process in Equation (2) with the feedback model described in Equation (1). Let $\bar{I}_t(\theta)$ denote the Fisher information of θ for observing t -th sample. We have*

$$\lim_{t \rightarrow \infty} \sum_{i=1}^t \bar{I}_i(\theta) = O(1).$$

PROOF SKETCH. Let $f(x|\theta)$ be the probability mass function of random variable X and x_t be the realization of X_t . Considering the stochastic process specified in Equation (1),

$$f(x_t|\theta) = \left(\frac{m\theta + S_{t-1}}{m+t-1}\right)^{x_t} \cdot \left(1 - \frac{m\theta + S_{t-1}}{m+t-1}\right)^{1-x_t},$$

where $x_t = 1$ or $x_t = 0$, $S_t = \sum_{i=1}^t X_i$. Define $l(x_t|\theta) = \log f(x_t|\theta)$. By definition of Fisher information for a single observation, and the chain rule for multiple observations, we have:

$$\begin{aligned} \sum_{i=1}^t I_i(\theta) &= \sum_{i=1}^t -\mathbb{E}[l''(x_i|\theta)] \\ &= \sum_{i=1}^t \frac{m^2}{m+i-1} \left(\mathbb{E}\left[\frac{1}{m\theta+S_{i-1}}\right] + \mathbb{E}\left[\frac{1}{m(1-\theta)+i-1-S_{i-1}}\right] \right) \end{aligned}$$

Since we know that $\{X_t\}_{t \geq 1}$ are exchangeable random variables,

$$\begin{aligned} \mathbb{E}\left[\frac{1}{m\theta+S_t}\right] &= \sum_{l=0}^t \binom{t}{l} \frac{\prod_{i=0}^{l-1} (m\theta+i) \cdot \prod_{j=0}^{t-l-1} (m(1-\theta)+j)}{\prod_{i=0}^{t-1} (m+i)} \cdot \frac{1}{m\theta+l} \\ &= O(t^{-1}) \end{aligned}$$

Similarly, we also have:

$$\mathbb{E}\left[\frac{1}{m(1-\theta)+t-1-S_{t-1}}\right] = O(t^{-1})$$

Thus:

$$\begin{aligned} \lim_{t \rightarrow \infty} \sum_{i=1}^t I_i(\theta) &= \lim_{t \rightarrow \infty} \sum_{i=1}^t \frac{m^2}{m+i-1} O(i^{-1}) \\ &= O(1) \end{aligned}$$

where $O(1)$ is a constant. This completes the proof. \square

Using this fact, by the general Cramér-Rao bound, we know that, for any estimator $\hat{\theta}_t$, the variance of $\hat{\theta}_t$ must follow:

$$\text{Var}(\hat{\theta}_t) \geq \Theta \left(\frac{1}{\sum_{i=1}^t I_i(\theta)} \right)$$

Since $\lim_{t \rightarrow \infty} \sum_{i=1}^t I_i(\theta)$ is bounded, the variance of any estimator will not shrink to zero even with infinitely many observations. Therefore, the learner cannot accurately infer the arm quality with high probability in the beta-herding feedback model and therefore cannot guarantee to identify the best arms even with infinitely many feedback instances. Since the learner only observes the feedback, we can conclude the following.

THEOREM 5.3. *If users' feedback follows beta-herding feedback model, there exists no bandit algorithm that can achieve sublinear regrets in our setting.*

PROOF SKETCH. We prove this by contradiction. Consider the case with two arms. Without loss of generality, assume arm 1 is optimal and arm 2 is suboptimal, i.e., $\theta_1 > \theta_2$, and suppose there exists an algorithm \mathcal{A} which can achieve sublinear regret, i.e., $\mathbb{E}(R_{\mathcal{A}}(T)) = o(T)$. Let k_t denote the arm chosen by algorithm \mathcal{A} at time t . One must have $\lim_{t \rightarrow \infty} \mathbb{P}(k_t = 1) = 1$. Let $\hat{\theta}_1^t, \hat{\theta}_2^t$ be the algorithm's estimators on θ_1, θ_2 given the history information accumulated till time round t . The ability to almost surely choose

arm 1 by algorithm \mathcal{A} when $t \rightarrow \infty$ indicates that we are able to differentiate the two arms, i.e.,

$$\lim_{t \rightarrow \infty} \mathbb{P}(\hat{\theta}_1^t > \hat{\theta}_2^t) = 1$$

However, as shown in Lemma 5.2, since the fisher information on the estimator are always bounded even when given infinitely many observations. It implies the estimators are not consistent, and that $\lim_{t \rightarrow \infty} \mathbb{P}(\hat{\theta}_1^t < \hat{\theta}_2^t) > 0$. This leads to the contradiction and completes the proof. \square

We note that the technique used in the proof can be extended to a more general feedback model for impossibility results. The intuition is to use Fisher information to quantify how informative a given data is with respect to a set of parameters and the influence of the data itself on the estimate. For different models, if the amount information for each feedback can be quantified, the same techniques can be applied.

5.3 An alternative approach: Designing information structures

Theorem 5.3 presents a strong impossibility result: if all feedback instances are generated according to beta-herding feedback model, we cannot design any bandit algorithms to achieve sublinear regret. A natural approach to get over this impossibility results is to break the assumption by taking interventions. Inspired by Bayesian persuasion [16], which designs the information structure to *persuade* agents to take certain actions, we explore whether we could design information structures to induce certain types of "feedback". For example, in the extreme case, if we do not show any historical information to users, and assume users provide unbiased feedback when no information is presented, then the problem reduces to standard bandit settings. However, in practice, we might not want to dramatically change the whole platform and might want to take as few interventions as possible. This leads to an interesting research question on whether we can minimally intervene the existing design of information structure, such that it is possible to design bandit algorithms with sublinear regrets.

In this section, we present a simple algorithm as a *toy example* to demonstrate the idea. A full study along this direction requires a careful and thorough modeling and is out of the scope of this paper. We consider the constrained setting in which the platform can only choose among two information design in each round, either showing all history information to users (and assuming users' feedback follow beta-herding feedback model) or showing no history information (and assuming users provide unbiased feedback). Our goal is to minimize the number of rounds that show no information to users while achieving sublinear regret. In particular, we propose a *two-stage policy*, as described in Algorithm 2, which shows no historical information for the first $\lceil T^\alpha \rceil$ rounds and resumes to standard design afterwards.

The regret bound of Algorithm 2 is given as follows.

THEOREM 5.4. *Let $\Theta = \{\theta_1, \dots, \theta_K\}$ be a bandit instance, and $\alpha \geq \ln(K+2)/\ln T$, then the expected regret of two-stage policy,*

Algorithm 2 *two-stage policy*

-
- 1: **Input:** learning rounds parameter $\alpha \in (0, 1)$, exploration parameter $\beta > 0$, number of arms K .
 - 2: **Initializations:** first K rounds, play each arm once
 - 3: **for** $t = K + 1, \dots, \lfloor T^\alpha \rfloor$ **do**
 - 4: **for** each $k \in \{1, \dots, K\}$ **do**
 - 5: $\text{UCB}_{k,t} = \hat{\theta}'_{k,t-1} + \sqrt{\frac{\beta \ln(t-1)}{n_{k,t-1}}}$, where $\hat{\theta}'_{k,t-1} = \frac{\sum_{s=1}^{n_{k,t-1}} \mathbb{1}\{I_s=k\} X_{t-1}}{n_{k,t-1}}$.
 - 6: Choose arm $I_t \in \text{argmax}_{k=1, \dots, K} \text{UCB}_{k,t}$.
 - 7: Present arm I_t without showing its history information to the user, and get feedback X_t .
 - 8: $\rho_{I_t,t} \leftarrow (\rho_{I_t,t-1} \times n_{I_t,t-1} + X_t) / (n_{I_t,t-1} + 1)$.
 - 9: $\rho_{k,t} = \rho_{k,t-1}$ for $k \neq I_t$.
 - 10: $n_{I_t,t} \leftarrow n_{I_t,t-1} + 1$.
 - 11: $n_{k,t} \leftarrow n_{k,t-1}$ for $k \neq I_t$.
 - 12: Let $I_\tau \in \text{argmax}_{k=1, \dots, K} n_{k, \lfloor T^\alpha \rfloor}$.
 - 13: Present arm I_τ with associated history information to the user in the remaining rounds.
 - 14: (all ties broken in some consistent way)
-

where $\beta > 1$, is bounded from above by:

$$\mathbb{E}[R(T)] \leq \sum_{k \neq I^*} \left(\frac{4\alpha\beta \ln T}{\Delta_k} + 8\beta\Delta_k \right) + (T - T^\alpha) \left(\sqrt{\frac{4K\alpha\beta \ln T}{T^\alpha - K}} + \frac{K}{\beta - 1} \left(\frac{T^\alpha - K}{K} \right)^{2-2\beta} \right)$$

where the second term is in an order of $\mathcal{O}\left((T - T^\alpha)\sqrt{\frac{K\beta\alpha \ln T}{T^\alpha}}\right)$.

To interpret the bound, when $\alpha \geq 1/2$, the above regret bound is in the order of $\mathcal{O}(\sqrt{\alpha T^\alpha \ln T})$, while when $\alpha < 1/2$, the above regret bound is in the order of $\mathcal{O}(\sqrt{\alpha T^{1-\alpha} \ln T})$.

Algorithm 2 presents an example that we can achieve sublinear regrets by modifying the information structures presented to users. In particular, we only need to hide the historical information from T^α users, with $\alpha < 1$, out of T users to achieve sublinear regrets. Note that we only consider a naive approach in a styled model, i.e., showing no information at all in some rounds, and assume simple user feedback models. We hope our results will encourage research that considers more fine-tuned information design and more thorough models of user feedback and platform utility.

6 DISCUSSION ON THE APPLICATIONS

In this section, we provide discussion on the applications of our setting. As the motivating example of this paper, we consider user-generated content platforms that need to learn content qualities through user feedback. Our analysis and results naturally extend to platforms that rely on user reviews to provide recommendations (such as Yelp or Amazon). However, to formulate the recommendation problem as a bandit learning problem, we need to make a simplifying assumption, as made in prior work [13, 19], that users are going to *follow* the recommendations. While this assumption

seems strong, in practice, it approximates users' behavior to a certain degree. In particular, empirical studies demonstrate that the probability for a users to view an item drops significantly when the position of the item decreases [8, 15, 25]. These empirical observations suggest that a significant amount of users are indeed following recommendations (since recommended items are ranked higher). Moreover, there have been recent studies on incentivizing exploration using information asymmetry [17, 20, 24] which demonstrate it is possible to make recommendations that users will *choose* to follow. The techniques in this paper can be applied in that line of work to explore the dynamics of feedback generation.

In addition to the above example, our setting applies to scenarios when the platform cannot observe the true objective but can only use (potentially biased) estimates as the proxy for the objective. Consider the following illustrating scenario: the police station needs to decide which area to send police officers to patrol at each time step. Each area i has an intrinsic, unknown crime rate p_i . When sending police officers to an area i , the police station obtains an unobserved reward $u(p_i)$, representing the value of increased safety for the area. Assume $u(p_i)$ is increasing in p_i . After the patrol, police officers need to report the amount of criminal activities during their patrol. However, these reports might be biased by the history of *reported* crime rate of the area. For example, if there are more reports of illegal activities in the area in the history, they might stop more people for inspection. This creates biases in the reports. If the goal is to maximize the sum of $u(p_i)$, this problem can be formulated using our setting, since the objective is a function of *true* crime rates, while the decision maker only has access to *reported* crime rates. Now assume the feedback model follows beta-herding feedback model. According to our results, without additional interventions, the police station might make *unfair* decisions in where to patrol using only the biased feedback, since it is impossible for them to infer the true crime rate from the reports. This example further emphasizes the importance of understanding human behavior in learning problems, especially when the corresponding actions have significant impacts on humans.

7 CONCLUSION AND FUTURE WORK

We explore bandit problems with biased human feedback under two different feedback models. In avg-herding feedback model, where users' feedback is biased only by the average feedback of the arm, we show that the updates of average feedback over time is mathematically equivalent to users collectively performing stochastic gradient descent. With this connection, we design a UCB-like algorithm that achieves sublinear regret under some mild conditions. However, in beta-herding feedback model, where users' feedback is biased both by the average feedback and the number of feedback instances of the arm, using arguments from information theory, we show that there exist no bandit algorithms that can achieve sublinear regret.

We hope our work will open more discussion on better understanding human behavior when designing algorithms for systems with humans in the loop. Our results also point to potentially future research directions on designing interfaces (e.g., in terms of how information is exchanged) between humans and machine learning algorithms to leverage the power of both ends.

REFERENCES

- [1] Zeyuan Allen-Zhu and Elad Hazan. 2016. Variance reduction for faster non-convex optimization. In *International Conference on Machine Learning (ICML)*. 699–707.
- [2] Jean-Yves Audibert and Sébastien Bubeck. 2009. Minimax policies for adversarial and stochastic bandits. In *The 22nd Conference on Learning Theory (COLT)*. 217–226.
- [3] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47, 2-3 (2002), 235–256.
- [4] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. 1995. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on*. IEEE, 322–331.
- [5] Abhijit V Banerjee. 1992. A simple model of herd behavior. *The quarterly journal of economics* 107, 3 (1992), 797–817.
- [6] Omar Besbes, Yonatan Gur, and Assaf Zeevi. 2014. Stochastic multi-armed-bandit problem with non-stationary rewards. In *Advances in Neural Information Processing Systems (NIPS)*. 199–207.
- [7] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy* 100, 5 (1992), 992–1026.
- [8] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 international conference on Web Search and Data Mining (WSDM)*. ACM, 87–94.
- [9] Peter Frazier, David Kempe, Jon Kleinberg, and Robert Kleinberg. 2014. Incentivizing Exploration. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC)*.
- [10] Noufel Frikha, Stéphane Menozzi, et al. 2012. Concentration bounds for stochastic approximations. *Electronic Communications in Probability* 17 (2012).
- [11] Aurélien Garivier and Eric Moulines. 2011. On Upper-Confidence Bound Policies for Switching Bandit Problems. In *Algorithmic Learning Theory*, Jyrki Kivinen, Csaba Szepesvári, Esko Ukkonen, and Thomas Zeugmann (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 174–188.
- [12] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. 2015. Escaping from saddle points—Online stochastic gradient for tensor decomposition. In *Conference on Learning Theory (COLT)*. 797–842.
- [13] Arpita Ghosh and Patrick Hummel. 2013. Learning and Incentives in User-generated Content: Multi-armed Bandits with Endogenous Arms. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science (ITCS)*.
- [14] Bruce M Hill, David Lane, and William Sudderth. 1980. A strong law for some generalized urn processes. *The Annals of Probability* (1980), 214–226.
- [15] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2017. Accurately interpreting clickthrough data as implicit feedback. In *ACM SIGIR Forum*, Vol. 51. ACM, 4–11.
- [16] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.
- [17] Ilan Kremer, Yishay Mansour, and Motty Perry. 2013. Implementing the "Wisdom of the Crowd". In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC)*. 605–606.
- [18] Tze Leung Lai and Herbert Robbins. 1985. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* 6, 1 (1985), 4–22.
- [19] Yang Liu and Chien-Ju Ho. 2018. Incentivizing High Quality User Contributions: New Arm Generation in Bandit Learning. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [20] Yishay Mansour, Aleksandrs Slivkins, and Vasilis Syrgkanis. 2015. Bayesian incentive-compatible bandit exploration. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation (EC)*. ACM, 565–582.
- [21] Lev Muchnik, Sinan Aral, and Sean J Taylor. 2013. Social influence bias: A randomized experiment. *Science* 341, 6146 (2013), 647–651.
- [22] Gergely Neu, Andras Antos, András György, and Csaba Szepesvári. 2010. Online Markov decision processes under bandit feedback. In *Advances in Neural Information Processing Systems (NIPS)*. 1804–1812.
- [23] Ronald Ortner, Daniil Ryabko, Peter Auer, and Rémi Munos. 2012. Regret bounds for restless Markov bandits. In *International Conference on Algorithmic Learning Theory (ALT)*. Springer, 214–228.
- [24] Yiangos Papanastasiou, Kostas Bimpikis, and Nicos Savva. 2017. Crowdsourcing exploration. *Management Science* (2017).
- [25] Matthew Richardson, Ewa Dominowska, and Robert Ragno. 2007. Predicting clicks: estimating the click-through rate for new ads. In *Proceedings of the 16th international conference on World Wide Web (WWW)*. ACM, 521–530.
- [26] Herbert Robbins and Sutton Monro. 1951. A stochastic approximation method. *The annals of mathematical statistics* (1951), 400–407.
- [27] Matthew J Salganik, Peter Sheridan Dodds, and Duncan J Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 5762 (2006), 854–856.
- [28] Ruben Sipos, Arpita Ghosh, and Thorsten Joachims. 2014. Was This Review Helpful to You?: It Depends! Context and Voting Patterns in Online Content. In *Proceedings of the 23rd International Conference on World Wide Web (WWW)*. 337–348.
- [29] Lones Smith and Peter SÅyrensen. 2000. Pathological Outcomes of Observational Learning. *Econometrica* 68, 2 (2000), 371–398.