

Incentivized Exploration for Multi-Armed Bandits under Reward Drift

Zhiyuan Liu*

Department of Computer Science
University of Colorado, Boulder
zhiyuan.liu@colorado.edu

Huazheng Wang*[†]

Department of Computer Science
University of Virginia
hw7ww@virginia.edu

Fan Shen

Technology, Cybersecurity and Policy
University of Colorado, Boulder
fan.shen@colorado.edu

Kai Liu

Computer Science Division
Clemson University
kail@clemson.edu

Lijun Chen

Department of Computer Science
University of Colorado, Boulder
lijun.chen@colorado.edu

Abstract

We study incentivized exploration for the multi-armed bandit (MAB) problem where the players receive compensation for exploring arms other than the greedy choice and may provide biased feedback on reward. We seek to understand the impact of this drifted reward feedback by analyzing the performance of three instantiations of the incentivized MAB algorithm: UCB, ϵ -Greedy, and Thompson Sampling. Our results show that they all achieve $\mathcal{O}(\log T)$ regret and compensation under the drifted reward, and are therefore effective in incentivizing exploration. Numerical examples are provided to complement the theoretical analysis.

Introduction

Multi-armed bandit (MAB) problem is a classical model for sequential decision making under uncertainty, and finds applications in many real world systems such as recommender systems (Li et al. 2010; Bouneffouf, Bouzeghoub, and Gañarski 2012), search engine systems (Radlinski, Kleinberg, and Joachims 2008; Yue and Joachims 2009) and cognitive radio networks (Gai, Krishnamachari, and Jain 2010), to just name a few. In the traditional MAB model, the decision maker (the principal) who selects the arm to pull and the action performers (the players) who actually pull the arm are assumed to be the same entity. This is, however, not true for several important real world applications where the principal and players are different entities with different interests. Take the Amazon product rating as an example: Amazon (the principal) would like the customers (the players) to buy and try different products (arms) of certain type in order to identify the best product (i.e., exploration), while the customers are heavily influenced by the current ratings on the products and behave myopically, i.e., select the product that currently has the highest rating (i.e., exploitation). It is well known that such exploitation-only behavior can be far from the optimal (Bubeck and Cesa-Bianchi 2012; Sutton and Barto 2018). In the traditional MAB setting, the

principal, who is also the player, strives to find the optimal tradeoff between exploration and exploitation and execute it accordingly. When the principal and players are different entities, misaligned interests between them need to be reconciled in order to balance exploration and exploitation in an optimal manner.

Incentivized learning has been proposed for the MAB problem to reconcile different interests between the principal and the players (Frazier et al. 2014; Mansour, Slivkins, and Syrgkanis 2015; Wang and Huang 2018; Immorlica et al. 2019). In order to incentivize exploration, the principal provides certain compensation to the player so that s/he will pull the arm other than the greedy choice that currently has the best empirical reward. The goal of the principal is to maximize the cumulative reward while minimizing the total compensation to the players.

However, existing incentivized MAB models (Han, Kempe, and Qiang 2015; Wang and Huang 2018; Immorlica et al. 2018; Liu and Ho 2018; Hirschschall et al. 2018) assume that the players provide unbiased stochastic feedback on reward¹ even after they receive certain incentive from the principal. This assumption does not always hold in the real world scenarios: work based on industrial level experiments in (Martensen, Gronholdt, and Kristensen 2000; Razak, Nirwanto, and Triatmanto 2016; Ehsani and Ehsani 2015) shows that the customers are inclined to give higher evaluation (i.e., increased reward) with incentive such as discount and coupon. The compensation could even be the primary driver of customer satisfaction (Martensen, Gronholdt, and Kristensen 2000; Lee and Lin 2005). This drift in reward feedback may cause negative impact on the exploration and exploitation tradeoff, e.g., a suboptimal arm is mistaken as the optimal one because of the incentives and the players will keep pulling it even after the compensation is removed. This has been ignored in previous research.

In this paper, we aim to investigate the impact of drifted reward feedback in the incentivized MAB problem. Specifically, we consider a general incentivized exploration algorithm where the player receives a compensation that is the difference in reward between the principal's choice and the greedy choice, and provides biased feedback that is the

*Equal Contribution

[†]Huazheng Wang is supported by Bloomberg Data Science Ph.D. Fellowship.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We will use reward and feedback interchangeably in this paper

sum of the true reward of an arm and a drift term that is a non-decreasing function of the compensation received for pulling this arm. We seek to answer the important question if the compensation scheme is effective in incentivizing exploration under reward drift from two intertwining aspects: (1) if the algorithm is robust to drifted reward so that the sequential decisions based on biased feedback still enjoy a small regret, and (2) if the proposed incentive mechanism is cost-efficient to the principal. We analyze the regret and compensation for three instantiations of the algorithm where the principal employs UCB, ϵ -Greedy, and Thompson Sampling, respectively. Our analytical results, complemented by numerical experiments, show that with drifted reward the proposed compensation scheme achieves both $\mathcal{O}(\log T)$ regret and $\mathcal{O}(\log T)$ compensation, and is thus effective in incentivizing exploration.

Related Work

Incentivized learning has attracted a lot of attention since the work (Frazier et al. 2014). In (Frazier et al. 2014), the authors proposed a Bayesian incentivized model with discounted regret and compensation, and characterized the relationship between the reward, compensation, and discount factor. In (Mansour, Slivkins, and Syrgkanis 2015), the authors studied the non-discount case and proposed an algorithm that has $\mathcal{O}(\sqrt{T})$ regret. In (Wang and Huang 2018), the authors analyzed the non-Bayesian and non-discount reward case and showed $\mathcal{O}(\log T)$ regret and compensation for incentivized exploration based on simplified MAB algorithms. But all the models and analysis are under the assumption that the players' feedbacks are unbiased under compensation. In contrast, we consider biased feedback under compensation, and show that the incentivized exploration with reward drift can still achieve $\mathcal{O}(\log T)$ regret and compensation.

Related work also includes those in robustness of MAB under adversarial attack. In (Lykouris, Mirrokni, and Paes Leme 2018), the authors proposed a multi-layer active arm elimination race algorithm for stochastic bandits with adversarial corruptions whose performance degrades linearly to the amount of corruptions. In (Feng, Parkes, and Xu 2019), the authors studied strategic behavior of rational arms and show that UCB, ϵ -Greedy, and Thompson sampling achieve $\mathcal{O}(\max\{B, \log T\})$ regret bound under any strategy of the strategic arms, where B is the total budget across arms. On the other hand, in (Jun et al. 2018), the authors constructed attacks by decreasing the reward of non-target arms, and showed that their algorithm can trick UCB and ϵ -Greedy to pull the optimal arm only $o(T)$ times under an $\mathcal{O}(\log T)$ attack budget. All the modeled attacks are from exogenous sources, e.g., malicious users, while in our paper, the reward drift can have an interpretation as arising from attacks but generated endogenously by the incentivized exploration algorithm itself.

Model, Notation, and Algorithm

Consider a variant of the multi-armed bandit problem where a principal has K arms, denoted by the set $[K] =$

$\{1, \dots, K\}$. The reward of each arm $i \in [K]$ follows a distribution with support $[0, 1]$ and mean μ_i that is unknown. Without loss of generality, we assume that arm 1 is the unique optimum with the maximum mean. Denote by $\Delta_i = \mu_1 - \mu_i$ the reward gap between arm 1 and arm i ($i \neq 1$), and let $\Delta = \min_i \Delta_i$. At each time $t = 1, \dots, T$, a new player will pull one arm $I_t \in [K]$ and receive a reward r_t that will be fed back to the principal and other players. Let $n_i(t) = \sum_{\tau=1}^{t-1} \mathbb{I}(I_\tau = i)$ denote the number of times that arm i has been pulled up to time t and $\hat{\mu}_i(t) = \frac{1}{n_i(t)} \sum_{\tau=1}^{t-1} r_\tau \mathbb{I}(I_\tau = i)$ the corresponding empirical average reward, where the indicator function $\mathbb{I}(A) = 1$ if A is true and $\mathbb{I}(A) = 0$ otherwise.

In real world applications, the principal and players may exhibit different behaviors. The principal would like to see the players select the best arm and maximize the cumulative reward. On the other hand, the players may be heavily influenced by other players' feedback, e.g., the reward history of the arms, and behave myopically, i.e., pull the arm that currently achieves the highest empirical reward (exploitation). It is well known that such a myopic exploitation-only behavior can be far from the optimum due to the lack of exploration (Sutton and Barto 2018). The principal cannot pull the arm directly, but can provide certain compensation to incentivize the players to pull arms with suboptimal empirical reward (exploration). However, this compensation may affect the players' feedback (Martensen, Gronholdt, and Kristensen 2000), which results in a biased reward history and disturbs both the principal and players' decisions. Specifically, we assume that at time t there is a drift b_t in feedback that is caused by compensation x , captured by an unknown function $b_t = f_t(x)$ with the following properties.

Assumption 1. *The reward drift function $f_t(x)$ is non-decreasing with $f_t(0) = 0$, and is Lipschitz continuous, i.e., for any x and y , there exists a constant l_t such that*

$$|f_t(x) - f_t(y)| \leq l_t |x - y|. \quad (1)$$

The biased feedback $r_t + b_t$ is then collected, and the principal and players know only the sum and cannot distinguish each part.

Let $l = \max_t l_t$ for later use. Denote by $E_t^i = 1$ the event that arm i is pulled with compensation at time t and $E_t^i = 0$ otherwise. Denote $B_i(t) = \sum_{\tau=1}^{t-1} b_\tau \mathbb{I}(E_\tau^i = 1)$ be the cumulative drift of arm i up to time t and $\bar{\mu}_i(t) = \hat{\mu}_i(t) + \frac{B_i(t)}{n_i(t)}$ be the corresponding average drifted reward. The general incentive mechanism and algorithm are described in Algorithm 1.

We characterize the performance of the incentivized exploration algorithm in terms of two metrics – the expected cumulative regret that quantifies the total loss because of not pulling the best arm, and the cumulative compensation that the principal pays for incentivizing exploration:

$$\mathbb{E}(R(T)) = \mathbb{E}\left(\sum_{t=1}^T (\mu_1 - \mu_{I_t})\right) = \sum_{i=2}^K \Delta_i \mathbb{E}(n_i(T+1)),$$

$$\mathbb{E}(C(T)) = \mathbb{E}\left(\sum_{t=1}^T (\bar{\mu}_{I_t} - \bar{\mu}_{G_t})\right).$$

Algorithm 1: Incentivized MAB under Reward Drift

```
1 for  $t = 1, 2, 3, \dots, T$  do
2   The principal selects arm  $I_t$  according to certain
   algorithm;
3   The player will choose  $G_t = \arg \max \bar{\mu}_i(t)$  w/o
   compensation;
4   if  $G_t = I_t$  then
5     The principal does not provide compensation;
6     The player pulls arm  $I_t$  and receives reward  $r_t$ ;
7   else
8     The principal provides compensation
      $\bar{\mu}_{G_t} - \bar{\mu}_{I_t}$ ;
9     The player pulls arm  $I_t$  and receives reward
      $r_t + b_t$ , with  $b_t = f_t(\bar{\mu}_{G_t} - \bar{\mu}_{I_t})$ .
10  end
11  Update average reward  $\bar{\mu}_{I_t}$ 
12 end
```

Notice that in Algorithm 1 the compensation and pulled arm are decided based on biased feedback which may not be an accurate reflection of an arm's reward, while the regret is in terms of the "true" reward that is unknown. We seek to answer the important question if the proposed compensation scheme is effective in incentivizing exploration from two intertwining aspects: (1) if the algorithm is robust to drifted reward so that the sequential decisions based on biased feedback still enjoy a sublinear regret, and (2) if the proposed incentive mechanism is cost-efficient to the principal. There are different arm selection strategies that the principal can employ, i.e., Step 2 of Algorithm 1. In the next section, in order to answer the above question, we analyze the cumulative regret and compensation under several typical multi-armed bandit algorithms such as UCB, ϵ -Greedy, and Thompson Sampling.

Regret and Compensation Analysis

In this section, we consider three instantiations of Algorithm 1 when the principal employs UCB, ϵ -Greedy, and Thompson Sampling at Step 2, respectively. As will be seen later, our analysis shows that the proposed compensation scheme is effective in incentivizing exploration under reward drift.

UCB policy

Consider first the case where the principal applies the UCB policy, i.e., uses the sum of average biased reward and upper confidence bound $\bar{\mu}_i(t) + \sqrt{\frac{2 \log t}{n_i(t)}}$ as the criterion to choose the arm to explore, as shown in Algorithm 2. The main result is summarized in Theorem 1.

Theorem 1. *For the incentivized UCB algorithm, the expected regret $R(T)$ and compensation $C(T)$ are bounded as follows:*

$$\mathbb{E}(R(T)) \leq \sum_{i=2}^N \frac{8(l+1)^2 \log T}{\Delta_i} + \frac{\Delta_i(K-1)\pi^2}{3}, \quad (2)$$

Algorithm 2: Incentivized UCB under Reward Drift

```
1 for  $t = 1, 2, 3, \dots, T$  do
2   The principal selects arm
    $I_t = \arg \max_i \bar{\mu}_i(t) + \sqrt{\frac{2 \log t}{n_i(t)}}$ ;
3   Steps 3-11 of Algorithm 1.
4 end
```

$$\mathbb{E}(C(T)) \leq \sum_{i=2}^N \frac{16(l+1) \log T}{\Delta_i} + \frac{16(l+1) \log T}{\Delta} + 2\pi K \sqrt{\frac{2 \log T}{3}}. \quad (3)$$

Proof. Notice that compensation is incurred under the conditions:

$$\begin{aligned} \bar{\mu}_{I_t}(t) &\leq \bar{\mu}_{G_t}(t), \\ \bar{\mu}_{I_t}(t) + \sqrt{\frac{2 \log t}{n_{I_t}(t)}} &\geq \bar{\mu}_{G_t}(t) + \sqrt{\frac{2 \log t}{n_{G_t}(t)}}. \end{aligned}$$

By the second condition, the compensation

$$\bar{\mu}_{G_t}(t) - \bar{\mu}_{I_t}(t) \leq \sqrt{\frac{2 \log t}{n_{I_t}(t)}}, \quad (4)$$

and further by Assumption 1, the drift $b_t \leq l_t \sqrt{\frac{2 \log t}{n_{I_t}(t)}}$. The total drift $B_i(t)$ of arm i can be bounded as follows (due to space limit, the details of inequality (5) are provided in supplementary material):

$$B_i(t) = \sum_{\tau=1}^t b_\tau \mathbb{I}(E_\tau^i = 1) \leq 2l \sqrt{2n_i(t) \log t}. \quad (5)$$

For each sub-optimal arm $i \neq 1$, if this arm is pulled by the player at t (with or without compensation), it must hold that

$$\hat{\mu}_i(t) + \frac{B_i(t)}{n_i(t)} + \sqrt{\frac{2 \log t}{n_i(t)}} \geq \hat{\mu}_1(t) + \frac{B_1(t)}{n_1(t)} + \sqrt{\frac{2 \log t}{n_1(t)}}.$$

So, the probability that arm i is pulled by the player at time t can be bounded by the following:

$$\begin{aligned} &\Pr(I_t = i) \\ &\leq \Pr \left(\hat{\mu}_i(t) + \frac{B_i(t)}{n_i(t)} + \sqrt{\frac{2 \log t}{n_i(t)}} \geq \hat{\mu}_1(t) + \frac{B_1(t)}{n_1(t)} + \sqrt{\frac{2 \log t}{n_1(t)}} \right) \\ &\leq \Pr \left(\hat{\mu}_i(t) + (2l+1) \sqrt{\frac{2 \log t}{n_i(t)}} \geq \hat{\mu}_1(t) + \frac{B_1(t)}{n_1(t)} + \sqrt{\frac{2 \log t}{n_1(t)}} \right) \\ &\leq \Pr \left(\hat{\mu}_i(t) + (2l+1) \sqrt{\frac{2 \log t}{n_i(t)}} \geq \hat{\mu}_1(t) + \sqrt{\frac{2 \log t}{n_1(t)}} \right), \end{aligned}$$

where the second inequality is due to the bound (5) on cumulative drift. Similar to the analysis in (Auer, Cesa-Bianchi, and Fischer 2002), notice that if the event $\hat{\mu}_i(t) + (2l+1)$

1) $\sqrt{\frac{2 \log t}{n_i(t)}} \geq \hat{\mu}_1(t) + \sqrt{\frac{2 \log t}{n_1(t)}}$ happens, one of the following three events must happen:

$$X_i(t) : \hat{\mu}_i(t) \geq \mu_i + \sqrt{\frac{2 \log t}{n_i(t)}},$$

$$Y_1(t) : \hat{\mu}_1(t) \leq \mu_1 - \sqrt{\frac{2 \log t}{n_1(t)}},$$

$$Z_i(t) : 2(l+1) \sqrt{\frac{2 \log t}{n_i(t)}} \geq \Delta_i.$$

Therefore, $\Pr(I_t = i) \leq \Pr(X_i(t)) + \Pr(Y_1(t)) + \Pr(Z_i(t))$. By the Chernoff-Hoeffding's inequality (Hoeffding 1994),

$$\Pr(X_i(t)) \leq \frac{1}{t^2}, \quad \Pr(Y_1(t)) \leq \frac{1}{t^2},$$

and their sum from $t = 1$ to T is bounded by $\frac{\pi^2}{3}$. If $n_i(t) \geq \frac{8(l+1)^2 \log T}{\Delta_i^2}$, the event $Z_i(t)$ will not happen, and thus $\sum_{t=1}^T \Pr(Z_i(t)) \leq \frac{8(l+1)^2 \log T}{\Delta_i^2}$. We can bound $\mathbb{E}[n_i(T)]$ as follows:

$$\begin{aligned} \mathbb{E}[n_i(T)] &= \sum_{t=1}^T \Pr(I_t = i) \\ &\leq \sum_{t=1}^T (\Pr(X_i(t)) + \Pr(Y_1(t)) + \Pr(Z_i(t))) \\ &\leq \frac{8(l+1)^2 \log T}{\Delta_i^2} + \frac{\pi^2}{3}. \end{aligned}$$

So, the expected regret

$$\mathbb{E}(R(T)) \leq \sum_{i=2}^K \frac{8(l+1)^2 \log T}{\Delta_i} + \frac{\Delta_i(K-1)\pi^2}{3}.$$

The calculation of compensation is a bit different from that of regret since compensation can be incurred even if the best arm is pulled. The player will be compensated to pull arm 1 only when

$$\begin{aligned} \bar{\mu}_1(t) &\leq \bar{\mu}_i(t), \\ \bar{\mu}_1(t) + \sqrt{\frac{2 \log t}{n_1(t)}} &\geq \bar{\mu}_i(t) + \sqrt{\frac{2 \log t}{n_i(t)}}, \end{aligned}$$

which requires $n_1(t) \leq n_i(t)$. So, the average number of times when the players are compensated to pull arm 1 is smaller than $\max_{i \neq 1} \mathbb{E}(n_i(T))$. Denote by $C_i(t)$ the total compensation the players have received to pull arm i up to time t . Recall the bound (4), we can bound the total compensation as follows:

$$\begin{aligned} \mathbb{E}(C(T)) &= \mathbb{E} \left(C_1(T) + \sum_{i=2}^K C_i(T) \right) \\ &\leq \sum_{m=1}^{\max_{i \neq 1} \mathbb{E}(n_i(T))} \sqrt{\frac{2 \log T}{m}} + \sum_{i=2}^K \sum_{m=1}^{\mathbb{E}(n_i(T))} \sqrt{\frac{2 \log T}{m}} \end{aligned}$$

$$\begin{aligned} &\leq \frac{16(l+1) \log T}{\Delta} + 2K\pi \sqrt{\frac{2 \log T}{3}} \\ &\quad + \sum_{i=2}^K \frac{16(l+1) \log T}{\Delta_i}. \end{aligned}$$

□

ε -Greedy policy

We now consider the case where the principal uses the ε -Greedy policy as shown in Algorithm 3, with the choice of exploration probability ε_t from that shows diminishing ε_t achieves better performance. Algorithm 3 involves a random exploration phase (Step 3), and its analysis is more involved. Recall that the ‘‘true’’ reward has a normalized support of $[0, 1]$, we therefore assume that the drifted reward $r_t + b_t$ is projected onto $[0, 1]$. This assumption is also consistent with real world applications such as Amazon and Yelp as their rating systems usually have lower and upper bounds.

Algorithm 3: Incentivized ε -Greedy under Reward Drift

```

1 for  $t = 1, 2, 3, \dots, T$  do
2   Let  $\varepsilon_t = \min(1, \frac{cK}{t})$ ;
3   With probability  $1 - \varepsilon_t$ , the principal chooses arm
    $I_t = \arg \max_i \bar{\mu}_i(t)$ ; with probability  $\varepsilon_t$ , the
   principal uniformly selects an arm  $I_t \in [K]$ ;
4   Steps 3-11 of Algorithm 1 with  $[r_t + b_t]_0^1$  where  $[\cdot]_0^1$ 
   denotes the projection onto  $[0, 1]$ 
5 end

```

Theorem 2. For the incentivized ε -Greedy algorithm with $\varepsilon_t = \min\{1, \frac{cK}{t}\}$ and $c \geq \frac{36}{\Delta}$, with a high probability the expected regret $R(T)$ and compensation $C(T)$ are bounded as follows:

$$\mathbb{E}(R(T)) \leq \sum_{i=2}^K cS_i(l)(\log T + 1) + c(K-1)(K + \frac{\pi^2}{6}), \quad (6)$$

$$\mathbb{E}(C(T)) \leq \max(l, 1)(c + \sqrt{3c})K \log T, \quad (7)$$

where $S_i(l) = 1.5 + 3(1 + \sqrt{3/c})l + 18c/\Delta_i^2$.

Proof. Since the biased feedback lies in the interval $[0, 1]$, the drift $b_t \leq l_t(\bar{\mu}_{G_t} - \bar{\mu}_{I_t}) \leq l_t$. A compensation for pulling arm i will be incurred only when the arm is chosen by the principal to explore. By Lemma 2 (Agarwal et al. 2014) in supplementary material, the number of explorations that arm i can receive up to time t is bounded by

$$m_i(t) \leq c(\log t + 1) + \sqrt{3c \log \frac{K}{\delta} (\log t + 1)} \quad (8)$$

with a probability of at least $1 - \delta$. When t is large enough such that $\log t \geq \log \frac{K}{\delta} - 1$, the right hand side of (8) is upper bounded by

$$\bar{m}_i(t) = (c + \sqrt{3c})(\log t + 1),$$

and the total drift $B_i(t)$ on arm i up to time t is upper bounded by $l\bar{m}_i(t)$ with a probability of at least $1 - \delta$.

Let $L = \frac{3l\bar{m}_i(T)}{\Delta_i}$ that is chosen to facilitate the analysis. We can bound $\mathbb{E}(n_i(T))$ as follows:

$$\begin{aligned}
& \mathbb{E}(n_i(T)) \\
& \leq \sum_{t=1}^T \frac{\varepsilon_t}{K} + \mathbb{E} \left(\sum_{t=1}^T (1 - \varepsilon_t) \mathbb{I}(I_t = i, n_i(t) \leq L) \right) \\
& \quad + \mathbb{E} \left(\sum_{t=1}^T (1 - \varepsilon_t) \mathbb{I}(I_t = i, n_i(t) \geq L) \right) \\
& \leq \underbrace{\sum_{t=1}^T \frac{\varepsilon_t}{K}}_A + L + \mathbb{E} \left(\sum_{t=1}^T \mathbb{I}(I_t = i, n_i(t) \geq L) \right) \\
& \leq A + \sum_{t=1}^T \Pr \left(\hat{\mu}_i(t) + \frac{B_i(t)}{n_i(t)} \geq \hat{\mu}_1(t) + \frac{B_1(t)}{n_1(t)}, n_i(t) \geq L \right) \\
& \leq A + \sum_{t=1}^T \Pr \left(\hat{\mu}_i(t) + \frac{\Delta_i}{3} \geq \hat{\mu}_1(t) \right) \\
& \leq A + \sum_{t=1}^T \Pr \left(\hat{\mu}_i(t) \geq \mu_i + \frac{\Delta_i}{3} \right) + \sum_{t=1}^T \Pr \left(\hat{\mu}_1(t) \leq \mu_1 - \frac{\Delta_i}{3} \right),
\end{aligned}$$

where the second last inequality is due to

$$\frac{B_i(t)}{n_i(t)} \leq \frac{g\bar{m}_i(t)}{3g\bar{m}_i(T)/\Delta_i} \leq \frac{\Delta_i}{3},$$

and the last inequality uses the fact that $\mu_i = \mu_1 - \Delta_i$. By Lemma 3 in supplementary material, when $c \geq \frac{36}{\Delta_i}$, we have

$$\begin{aligned}
& \sum_{t=1}^T \Pr \left(\hat{\mu}_i(t) \geq \mu_i + \frac{\Delta_i}{3} \right) + \sum_{t=1}^T \Pr \left(\hat{\mu}_1(t) \leq \mu_1 - \frac{\Delta_i}{3} \right) \\
& \leq \left(\frac{c}{2} + \frac{18}{\Delta_i^2} \right) \log T + c(K + \frac{\pi^2}{\Delta_i^2}) + \frac{18}{\Delta_i^2}.
\end{aligned}$$

We can also show that $A \leq c(1+3g(1+\sqrt{3/c}))(\log T+1)$, and further obtain the bound (6) on expected regret after some straightforward mathematical manipulations.

For the compensation analysis, notice again that the drifted reward is in $[0, 1]$, so the compensation at each time is less than 1 and the total compensation the players receive to pull arm i is bounded by the bound \bar{m}_i on the number of explorations it receives. To be consistent with the case with no drift, we write the bound on expected compensation as

$$\mathbb{E}(C(T)) \leq \max(l, 1)(c + \sqrt{3c})K(\log T + 1). \quad \square$$

Thompson Sampling

Consider now the case where the principal uses Thompson Sampling as shown in Algorithm 4. Thompson Sampling starts with a (prior) distribution on each arm's reward, and updates the distribution after the arm being pulled. At each

time, the principal samples the reward of each arm according to its posterior distribution, and then selects the arm with the highest sample reward. In this paper, we consider Gaussian prior adopted from (Agrawal and Goyal 2013) since the often used Beta priors are usually for binary reward feedback.

Before we analyze the performance of Algorithm 4, we first introduce some definitions and notations that are adopted from (Agrawal and Goyal 2017; 2013).

Algorithm 4: Incentivized Thompson Sampling under Reward Drift

```

1 for  $t = 1, 2, 3, \dots, T$  do
2   The principal independently samples  $\theta_i(t)$  from
   distribution  $\mathcal{N}(\bar{\mu}_i(t), \frac{1}{n_i(t)+1})$  and selects arm
    $I_t = \arg \max_i \theta_i(t)$ ;
3   Steps 3-11 of Algorithm 1.
4 end

```

Definition 1. For each arm i , we denote two thresholds x_i and y_i such that $\mu_i \leq x_i \leq y_i \leq \mu_1$. $E_i^\mu(t)$ denotes the event $\bar{\mu}_i(t) \leq x_i$ and $E_i^\theta(t)$ the event $\theta_i(t) \leq y_i$. Also, let $p_{i,t} = \Pr(\theta_1(t) \geq y_i | \mathcal{F}_{t-1})$ where \mathcal{F}_{t-1} is the history of plays until time t .

Definition 2. For two arms i and j , if $\bar{\mu}_i(t) \neq \bar{\mu}_j(t)$, there exists a constant Δ_{ij} such that $|\bar{\mu}_i(t) - \bar{\mu}_j(t)| \geq \Delta_{ij}$. Let $\underline{\Delta} = \min \Delta_{ij}$.

We have the following result on the frequency $m_i(T)$ of compensation the players receive for pulling each arm i when considering the concentration inequality of Gaussian distribution (Abramowitz 1965).

Lemma 1. The expected frequency $\mathbb{E}(m_i(T))$ of compensation for pulling arm i is bounded by $\frac{2 \log T}{\Delta_i^2}$.

Proof. The proof is provided in the supplemental material. \square

Our analysis of regret generalizes that in (Agrawal and Goyal 2017; Feng, Parkes, and Xu 2019) to incorporate the effect of drift caused by compensation.

Theorem 3. For the incentivized Thompson Sampling algorithm, the expected regret $R(T)$ and compensation $C(T)$ can be bounded as follows:

$$\mathbb{E}(R(T)) \leq \sum_{i=2}^K ((4e^{11} + 21)P_i(T) + \frac{5}{\Delta_i^2} + Q_i(T) + \frac{\pi^2}{6}), \quad (9)$$

$$\mathbb{E}(C(T)) \leq 2 \max(l, 1)K \frac{\log T}{\underline{\Delta}^2}, \quad (10)$$

where $P_i(T) = \frac{18 \log(T\Delta_i^2)}{\Delta_i^2}$ and $Q_i(T) = \lceil \frac{9}{2\Delta_i^2} \left((1 + \frac{4\Delta_i l}{3\Delta_i^2}) \log T + \sqrt{1 + \frac{8\Delta_i l \log T}{3\Delta_i^2}} \right) \rceil$.

Proof. The analysis of compensation is straightforward, similar to that for the incentivized ε -Greedy algorithm.

By Lemma 1, the expected compensation $\mathbb{E}(C(T)) \leq 2 \max(l, 1) K \frac{\log T}{\Delta^2}$.

Consider now the regret for choosing suboptimal arm $i (i \neq 1)$. We can bound $\mathbb{E}(n_i(T))$ as follows:

$$\begin{aligned} & \mathbb{E}(n_i(T)) \\ & \leq \sum_{t=1}^T \Pr(I_t = i, E_i^\mu(t), E_i^\theta(t)) \\ & \quad + \sum_{t=1}^T \Pr(I_t = i, E_i^\mu(t), \overline{E_i^\theta(t)}) + \sum_{t=1}^T \Pr(I_t = i, \overline{E_i^\mu(t)}) \end{aligned}$$

The first two terms can be bounded by the results of (Agrawal and Goyal 2017), see the detail in supplemental material, since their analysis will not be affected by the reward drift. Specifically, by Lemma 4, the sum of first two terms is upper bounded by $cP_i(T) + \frac{5}{\Delta_i^2}$, where c is certain constant and $P_i(T) = \frac{18 \log(T\Delta_i^2)}{\Delta_i^2}$. As for the third term, the analysis is similar to that of UCB and ε -Greedy algorithm where the drift is bounded by $\mathcal{O}(\log T)$:

$$\begin{aligned} & \sum_{t=1}^T \Pr(I_t = i, \overline{E_i^\mu(t)}) \\ & \leq \sum_{t=1}^T \Pr(\overline{E_i^\mu(t)}) = \sum_{t=1}^T \Pr(\bar{\mu}_i(t) \geq x_i) \\ & = \sum_{t=1}^T \Pr\left(\hat{\mu}_i(t) + \frac{B_i(t)}{n_i(t)} \geq x_i\right) \\ & = \sum_{t=1}^T \Pr\left(\hat{\mu}_i(t) - \mu_i \geq \underbrace{\frac{\Delta_i}{3} - \frac{B_i(t)}{n_i(t)}}_{Y_i(t)}\right) \\ & \leq \sum_{t=1}^T \Pr(\hat{\mu}_i(t) - \mu_i \geq Y_i(t), n_i(t) \leq Q_i) \\ & \quad + \sum_{t=1}^T \Pr(\hat{\mu}_i(t) - \mu_i \geq Y_i(t), n_i(t) \geq Q_i) \\ & \leq Q_i + \sum_{t=1}^T e^{-2n_i(t)Y_i(t)^2} \leq Q_i + \frac{\pi^2}{6}. \end{aligned}$$

where the second last inequality is due to Hoeffding's inequality. We then choose Q_i such that, when $n_i(t) \geq Q_i$,

$$\frac{\Delta_i}{3} - \frac{B_i(t)}{n_i(t)} \geq \frac{\Delta_i}{3} - \frac{2l \log T}{\Delta^2 n_i(t)} \geq 0, \quad (11)$$

$$n_i(t)Y_i(t)^2 \geq \log T. \quad (12)$$

By (11), $Q_i \geq \frac{6l \log T}{\Delta_i \Delta^2}$. Since $n_i(t)Y_i(t)^2$ is non-increasing in $B_i(t)$, equation (12) requires

$$\frac{\Delta_i^2}{9} n_i + \frac{4l^2 \log^2 T}{\Delta^4} \frac{1}{n_i} \geq \left(1 + \frac{4\Delta_i l}{3\Delta^2}\right) \log T.$$

The above two equations lead to

$$Q_i \geq \left\lceil \frac{9}{2\Delta_i^2} \left(\left(1 + \frac{4\Delta_i l}{3\Delta^2}\right) \log T + \sqrt{1 + \frac{8\Delta_i l \log T}{3\Delta^2}} \right) \right\rceil. \quad \square$$

Discussion of Results

As can be seen from the above analysis, all three instantiations of the incentivized exploration algorithm attain $\mathcal{O}(\log T)$ regret and compensation upper bound under drifted reward. Our results match both the theoretical lower bound for regret in (Lai and Robbins 1985) and lower bound for compensation in (Wang and Huang 2018) without reward drift. Although explicit lower bounds of the regret and compensation with drifted feedback in our setting remain unknown, we argue that these lower bounds should be larger or equal to the lower bound without reward drift since non-drifting environment is a special case of the drifted reward feedback with drift function $f_t = 0$. On the other hand, the proposed incentive mechanism is still cost-efficient even the payment will lead to biased feedback, as the principal can reduce the regret from $\mathcal{O}(T)$ for the players' myopic choices to $\mathcal{O}(\log T)$ by paying merely $\mathcal{O}(\log T)$ in incentive.

In terms of sensitivity to unknown drift functions f_t , both incentivized ε -Greedy and Thompson Sampling attain $\mathcal{O}(l)$ regret and compensation, while the incentivized UCB attains $\mathcal{O}((l+1)^2)$ regret and $\mathcal{O}(l+1)$ compensation. This difference comes from two aspects: 1) UCB is deterministic given the history while ε -Greedy and Thompson Sampling have a randomized exploration phase which makes them less sensitive to the drift. 2) For UCB, the drift effect is bounded by the amount of compensation which affects the frequency of compensation and in turn shapes the amount of compensation, while for ε -Greedy and Thompson Sampling, the cumulative drift can be directly bounded by the frequency of compensation. Numerical experiments reported in the next section are consistent with these analytical results.

Numerical Examples

In this section, we carry out numerical experiments using synthetic data to complement the previous analysis of the incentivized MAB algorithms under reward drift, including UCB, ε -Greedy and Thompson Sampling.

We generate a pool of $K = 9$ arms with mean reward vector $\boldsymbol{\mu} = [0.9, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2, 0.1]$. In each iteration, after the player pulls an arm I_t , reward r_t is set to the arm's mean reward plus a random term drawn from $\mathcal{N}(0, 1)$, i.e. $r_t = \boldsymbol{\mu}_{I_t} + \mathcal{N}(0, 1)$. Because of the randomness in sample rewards, the greedy algorithm without exploration suffers a linear regret, e.g., we observe nearly 6000 regret for 20000 trials. For the reward drift under compensation, we consider a linear drifting function $b_t = lx_t$ where x_t is the compensation offered by the principle and coefficient $l \geq 0$. The player reveals drifted reward feedback $r_t + b_t$.

For the incentivized exploration, we first compare regret and compensation in a non-drifting environment ($l =$

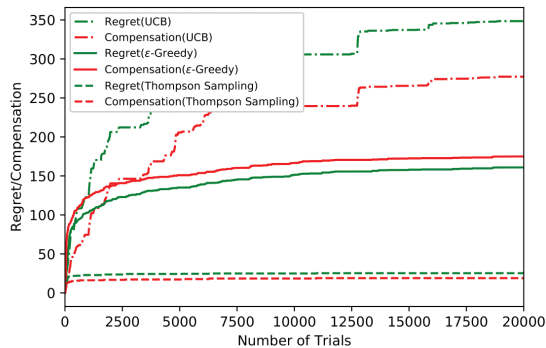


Figure 1: Regret and Compensation for UCB, ϵ -Greedy and Thompson Sampling without reward drift.




Figure 2: Regret and Compensation for UCB, ϵ -Greedy and Thompson Sampling with drift coefficient $l = 1.1$.

0) and a drifted reward environment ($l > 0$). In a non-drifting reward environment the player always gives unbiased feedback even offered with compensation. The result is shown in Fig. 1. As expected, all three instantiations of the incentivized MAB algorithms have a sub-linear regret and compensation. Thompson Sampling outperforms the other two both in the regret (which is consistent with observation from previous work (Vermorel and Mohri 2005; Chapelle and Li 2011)) and compensation.

In Fig. 2 we show the performance of the incentivized MAB algorithms under drifted reward with drift coefficient $l = 1.1$. We first observe that over the three algorithms Thompson Sampling still performs the best. While their relative performance are in same order as Fig. 1, the regret and compensation are worse than non-drifting setting, e.g., regret of UCB increases from 350 to 800 because of the biased feedback.

To better understand the effect of drifted reward, we vary the coefficient l from 0 to 1.1 and present the results in Table 1. We notice that the incentivized UCB incurs largest regret and compensation. This is due to the fact that, as the time goes, a larger UCB and uncertainty are assigned to those arms that are less explored but may in fact have small mean

l	0	0.05	0.1	0.4	0.7	0.9	1.1
UCB(R)	348.5	432.1	451.9	522.8	615.1	712.9	854.2
UCB(C)	277.2	292.9	349.5	375.6	408.0	473.0	422.7
ϵ -Greedy(R)	160.0	170.3	218.0	260.1	266.2	272.6	317.0
ϵ -Greedy(C)	185.9	217.4	130.4	167.6	102.8	161.8	115.2
TS(R)	25.3	28.2	33.4	37.1	46.3	63.6	74.5
TS(C)	18.9	23.7	20.9	29.3	22.9	29.1	25.3

Table 1: Regret (R) and Compensation (C) with different drift coefficients.

l	0	0.05	0.1	0.4	0.7	0.9	1.1
UCB(N)	1225	1639	1954	2172	2288	2912	3374
UCB(E)	0.4%	0.9%	0.5%	1.2%	1.9%	0.4%	3.1%
ϵ -Greedy(N)	273	329	304	303	276	293	308
ϵ -Greedy(E)	0.7%	1.5%	0.5%	1.0%	1.6%	0.4%	0.8%
TS(N)	60	79	58	98	131	109	106
TS(E)	0.7%	0.7%	1.6%	2.0%	0.1%	1.7%	0.7%

Table 2: Number of compensation (N) and relative error (E) of estimation of arm 1 with different drift coefficients.

rewards, and the resulting higher chance of those suboptimal arms being selected leads to larger regret and compensation. We also notice that the gap between regret and compensation of UCB increase faster compared to the other two. This is consistent with out theoretical analysis that the regret of UCB is in the order of $\mathcal{O}((l+1)^2)$ and compensation is in the order $\mathcal{O}(l+1)$.

We then exam the frequency of compensation, as well as the estimation error for arm 1 in terms of the relative error of the average drifted reward compared to the mean reward, and present the result in Table 2. We see that all three incentivized exploration algorithms achieve small estimation errors that are not sensitive to the drift coefficient l . This is expected, as the the expected compensation and thus the drift per time approaches 0 as T increases. However, while the incentivized ϵ -Greedy and Thompson Sampling have roughly a constant frequency of compensation across different l values, the incentivized UCB is more sensitive to the coefficient in the frequency of compensation. The constant frequency of compensation for ϵ -Greedy and Thompson Sampling can be seen from the proof of Theorem 2 and Lemma 1 that show the frequency does not depend on the drift. In contrast, seen from the proof of Theorem 1, the frequency of compensation for UCB depends on the drift through equation (6).

Conclusion

We propose and study multi-armed bandit algorithm with incentivized exploration under reward drift, where the player provides a biased reward feedback that is the sum of the true reward and a drift term that is non-decreasing in compensation. We analyze the regret and compensation for three instantiations of the incentivized MAB algorithm where the principal employs UCB, ϵ -Greedy and Thompson Sampling, respectively. Our results show that the algorithms achieve $\mathcal{O}(\log T)$ regret and compensation, and are therefore effective in incentivizing exploration. Our current analysis is based on the assumption that the reward drift is non-decreasing over the compensation. In the future work, we would like to study other assumptions about drift function and their corresponding impact on regret and compensation.

It is also important to explore if an algorithm can leverage the drifted reward to reduce the compensation.

References

- Abramowitz, M. 1965. Handbook of mathematical functions with formulas. *Graphs, and Mathematical Tables*.
- Agarwal, A.; Hsu, D.; Kale, S.; Langford, J.; Li, L.; and Schapire, R. 2014. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, 1638–1646.
- Agrawal, S., and Goyal, N. 2013. Further optimal regret bounds for thompson sampling. In *Artificial intelligence and statistics*, 99–107.
- Agrawal, S., and Goyal, N. 2017. Near-optimal regret bounds for thompson sampling. *Journal of the ACM (JACM)* 64(5):30.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3):235–256.
- Bouneffouf, D.; Bouzeghoub, A.; and Gançarski, A. L. 2012. A contextual-bandit algorithm for mobile context-aware recommender system. In *International Conference on Neural Information Processing*, 324–331. Springer.
- Bubeck, S., and Cesa-Bianchi, N. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning* 5(1):1–122.
- Chapelle, O., and Li, L. 2011. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, 2249–2257.
- Ehsani, Z., and Ehsani, M. H. 2015. Effect of quality and price on customer satisfaction and commitment in iran auto industry. *International Journal of Service Science, Management and Engineering* 1(5):52.
- Feng, Z.; Parkes, D. C.; and Xu, H. 2019. The intrinsic robustness of stochastic bandits to strategic manipulation. *arXiv preprint arXiv:1906.01528*.
- Frazier, P.; Kempe, D.; Kleinberg, J.; and Kleinberg, R. 2014. Incentivizing exploration. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 5–22. ACM.
- Gai, Y.; Krishnamachari, B.; and Jain, R. 2010. Learning multiuser channel allocations in cognitive radio networks: A combinatorial multi-armed bandit formulation. In *2010 IEEE Symposium on New Frontiers in Dynamic Spectrum (DySPAN)*, 1–9. IEEE.
- Han, L.; Kempe, D.; and Qiang, R. 2015. Incentivizing exploration with heterogeneous value of money. In *International Conference on Web and Internet Economics*, 370–383. Springer.
- Hirnschall, C.; Singla, A.; Tschitschek, S.; and Krause, A. 2018. Learning user preferences to incentivize exploration in the sharing economy. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Hoeffding, W. 1994. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*. Springer. 409–426.
- Immorlica, N.; Mao, J.; Slivkins, A.; and Wu, Z. S. 2018. Incentivizing exploration with unbiased histories. *arXiv preprint arXiv:1811.06026*.
- Immorlica, N.; Mao, J.; Slivkins, A.; and Wu, Z. S. 2019. Bayesian exploration with heterogeneous agents. In *The World Wide Web Conference*, 751–761. ACM.
- Jun, K.-S.; Li, L.; Ma, Y.; and Zhu, J. 2018. Adversarial attacks on stochastic bandits. In *Advances in Neural Information Processing Systems*, 3640–3649.
- Lai, T. L., and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* 6(1):4–22.
- Lee, G.-G., and Lin, H.-F. 2005. Customer perceptions of e-service quality in online shopping. *International Journal of Retail & Distribution Management* 33(2):161–176.
- Li, L.; Chu, W.; Langford, J.; and Schapire, R. E. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, 661–670. ACM.
- Liu, Y., and Ho, C.-J. 2018. Incentivizing high quality user contributions: new arm generation in bandit learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Lykouris, T.; Mirrokni, V.; and Paes Leme, R. 2018. Stochastic bandits robust to adversarial corruptions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, 114–122. ACM.
- Mansour, Y.; Slivkins, A.; and Syrgkanis, V. 2015. Bayesian incentive-compatible bandit exploration. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 565–582. ACM.
- Martensen, A.; Gronholdt, L.; and Kristensen, K. 2000. The drivers of customer satisfaction and loyalty: cross-industry findings from denmark. *Total Quality Management* 11(4-6):544–553.
- Radlinski, F.; Kleinberg, R.; and Joachims, T. 2008. Learning diverse rankings with multi-armed bandits. In *Proceedings of the 25th international conference on Machine learning*, 784–791. ACM.
- Razak, I.; Nirwanto, N.; and Triatmanto, B. 2016. The impact of product quality and price on customer satisfaction with the mediator of customer value. *Journal of Marketing and Consumer Research* 30(1):59–68.
- Sutton, R. S., and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Vermorel, J., and Mohri, M. 2005. Multi-armed bandit algorithms and empirical evaluation. In *European conference on machine learning*, 437–448. Springer.
- Wang, S., and Huang, L. 2018. Multi-armed bandits with compensation. In *NeurIPS*.
- Yue, Y., and Joachims, T. 2009. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, 1201–1208. ACM.