

Incentivized Exploration of Non-Stationary Stochastic Bandits

Sourav Chakraborty and Lijun Chen

Abstract—We study incentivized exploration for the multi-armed bandit (MAB) problem with non-stationary reward distributions, where players receive compensation for exploring arms other than the greedy choice and may provide biased feedback on the reward. We consider two different non-stationary environments: abruptly-changing and continuously-changing, and propose respective incentivized exploration algorithms. We show that the proposed algorithms achieve sublinear regret and compensation over time, thus effectively incentivizing exploration despite the nonstationarity and the biased or drifted feedback.

I. INTRODUCTION

The multi-armed bandit (MAB) problem is one of basic models for sequential decision-making under uncertainty, with diverse applications in areas such as clinical trials [1]–[3], financial portfolio design [4], recommendation systems [5], [6], search engine systems [7], and cognitive radio networks [8]. In the traditional MAB model, a decision maker iteratively selects an arm (or action) to pull at each time step, receives a certain reward from the environment, and decides on the arm for the next iteration. In the so-called stochastic MAB model, each arm’s reward distribution is unknown but remains fixed over time (hence, the ‘stationary’ bandit setting).

The objective of the decision maker (or the MAB algorithm) is to minimize the expected regret over the entire time horizon, defined as the expectation of the difference between the total reward obtained by pulling the best arm and the total reward obtained by the algorithm. Minimizing regret is achieved by balancing exploitation, the use of acquired information, with exploration, acquiring new information. If the decision maker always pulls the arm believed to be the best (i.e., exploitation only), they may miss the opportunity to identify another arm with a potentially higher expected reward. On the other hand, if the decision maker excessively explores various arms (i.e., exploration only), they will fail to accumulate as much reward as possible.

In this setup, the decision maker (the principal) and the player (the agent) who pulls the arm are assumed to be the same entity striving to balance exploitation and exploration. However, this may not always be the case in the real world. Many scenarios exist where the principal and the agent are different entities with different interests. The agent may select the *currently* best-performing arm in the face of uncertain reward (i.e., exploitation only), while the principal is interested in identifying the best-performing arm

in the long run (thus, the need to balance exploration and exploitation). Consider, for instance, an e-commerce system like Amazon. Amazon (the principal) would like the customers (the agents) to buy and try different products (arms) to identify the revenue-maximizing product (i.e., the best-performing arm in the long run) for a particular search query. However, customers are influenced by the current ratings and reviews of the products and behave myopically, i.e., selecting the currently highest-rated product (exploitation only). Such exploitation-only behavior can lead to significantly degraded performance due to inadequate exploration, as demonstrated in previous studies [9], [10]. The misaligned interests between the principal and the agents need to be reconciled to balance exploration and exploitation optimally.

Incentivized exploration has been introduced to the MAB problem to reconcile different interests between the principal and the agents [11]–[14]. The principal provides certain compensation to the agent to pull an arm other than the greedy choice currently having the best empirical reward, aiming to maximize the cumulative reward (or minimize the expected regret) while minimizing the total compensation to the agents. Early work on incentivized MAB models [13]–[17] assumed that the agents provide unbiased feedback or reward, independent of compensation received. However, this assumption does not always hold in the real world, and experimental studies such as [18], [19] show that agents are inclined to give higher evaluation or reward with an incentive (such as a coupon, gift card, or discount in the case of Amazon). The compensation might even be the primary driver of customer satisfaction [18], [20]. This drift in reward feedback may negatively impact the exploration-exploitation tradeoff, as a suboptimal arm can be mistakenly identified as the optimal one because of the drifted rewards. Liu et al. [21] investigated such an impact and showed that incentivized exploration based on their methods achieves optimal regret and compensation.

The authors in [21] considered the *stationary* bandit setting, i.e., the reward distribution of the arms does not change with time. In this paper, we consider the more challenging setting of *nonstationary* bandits, corresponding to an evolving environment where the reward distribution changes over time. Consider again the Amazon example: In the stationary setting, a product, say, a snow boot (an arm), is assumed to have the same value to Amazon (the principal) in terms of sales throughout the year. However, a snow boot will be more valuable in winter than the summer. A specific product might gain sudden popularity due to celebrity endorsement or lose popularity because of a certain controversy. Such scenarios are common in the real world and require the consideration

S. Chakraborty & L. Chen are with the Department of Computer Science, the University of Colorado, Boulder, Colorado, USA. Emails: {sourav.chakraborty, lijun.chen}@colorado.edu

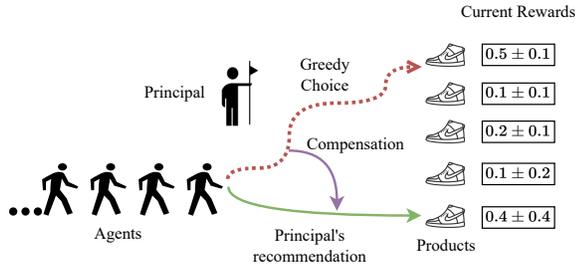


Fig. 1: Incentivized Exploration

TABLE I: Regret and compensation along with the corresponding theorems for the abruptly-changing (AC) and continuously-changing (CC) environments.

Env.	Algo. Scheme	Theorem	Regret	Compensation
AC	Alg. 3 + DUCB	1, 3	$\tilde{O}(T^{1/2})$	$\tilde{O}(T^{1/2})$
AC	Alg. 3 + SWUCB	2, 4	$\tilde{O}(T^{1/2})$	$\tilde{O}(T^{1/4})$
CC	Alg. 4 + UCB1	5, 6	$\tilde{O}(T^{2/3})$	$\tilde{O}(T^{2/3})$
CC	Alg. 4 + ϵ -Greedy	5, 6	$\tilde{O}(T^{2/3})$	$\tilde{O}(T^{1/3})$
CC	Alg. 4 + TS	5, 6	$\tilde{O}(T^{2/3})$	$\tilde{O}(T^{1/3})$

of nonstationary bandits. We aim to answer the following question: *Can we achieve effective incentivized exploration, despite the non-stationarity of the reward and the drift in reward feedback?*

Contributions. Specifically, we consider the incentivized exploration framework ([21], as illustrated in Fig. 1) where the agent receives from the principal a compensation that equals the difference in the estimated rewards between the principal’s recommended arm and the greedy choice, and provides biased feedback that is the sum of the true reward of an arm and a drift term that is a non-decreasing function of the compensation received for pulling the arm, but with changing reward distributions over time. We consider two non-stationary models and study the robustness of the proposed incentivized exploration algorithms in terms of regret and compensation. The first model assumes an abruptly changing environment where the rewards of each arm remain stationary until some breakpoint when they change abruptly. The second model considers a continuously changing environment where the rewards can vary continuously within a variation budget. We show that the regret and compensation bounds are sub-linear in time T , see Table I, and thus the proposed algorithms effectively incentivize exploration in the non-stationary environment.

Related Work. Early work on incentivized exploration and learning includes [11], [22], [23] that introduced a Bayesian incentivized model with discounted regret and compensation, and [12] that considered the non-discounted case and proposed an algorithm with $O(\sqrt{T})$ regret. In [13], the authors analyzed the non-Bayesian and non-discount reward case and demonstrated $O(\log T)$ regret and compensation. In [21], the authors considered biased user feedback under the

influence of incentives and showed that despite the reward drift, the proposed algorithms achieve $O(\log T)$ regret and compensation. Related work also includes [12], [24]–[26] on Bayesian Incentive Compatible (BIC) bandit exploration, where the principal wishes to persuade the agent to take action benefiting the principal, known as Bayesian Persuasion [27]. Additionally, see [28] for a review of the broad area of incentivized exploration.

The paper is organized as follows: In Section II, we introduce some preliminary concepts and results in stationary (Section II-A) and non-stationary MAB problems (Section II-B), as well as the incentivized exploration problem (Section II-C). In Section III, we present the proposed algorithms for incentivized exploration in abruptly-changing environments (Section III-A) and continuously changing non-stationary environments (Section III-B). In Section IV, we present numerical experiment results for the proposed algorithms. The proofs of all theoretical results can be found in the Appendix (Section VI).

II. PRELIMINARIES

A. Standard stochastic (stationary) MAB Problem

At each time step $t \in \{1, 2, \dots, T\}$, a decision maker chooses to arm a from a set of K arms based on the sequence of past arm pulls and received rewards and obtains a reward $X_t(a)$. The rewards for each arm are modeled by a sequence of independent and identically distributed (i.i.d.) random variables from an unknown distribution. Without loss of generality, the reward of each arm is assumed to be in $[0, 1]$. Denote by $\mu(a)$ the expectation of the reward of arm a , and a^* the optimal arm with the highest expected reward μ^* . The benchmark (or optimal) performance comprises pulling the optimal arm a^* at every time step. The regret R_T of an algorithm is defined as the difference between the benchmark performance and the total rewards collected by the algorithm:

$$R_T = \sum_{t=1}^T (\mu^* - X_t(a_t)), \quad (1)$$

where a_t is the arm the algorithm pulls at time t . A stochastic bandit algorithm’s performance is typically evaluated by how the *expected value* of R_T scales with the time horizon T , and the goal is to design algorithms that achieve sub-linear expected regret in T . The most common algorithms achieving a sub-linear regret are UCB1 ([29], [30]), Thompson Sampling ([31]), and ϵ -Greedy ([29], [10]).

B. Non-stationary MAB Problem

In the non-stationary setting, the rewards $X_t(a)$ for arm a are modeled by a sequence of independent random variables from potentially different distributions that are unknown and *may change* across time. Denote by $\mu_t(a_t)$ the expectation of the reward $X_t(a_t)$ for arm a_t at time step t . Similarly, let a_t^* be the arm with the highest expected reward, denoted by μ_t^* , at time t . The benchmark performance of the algorithm would be achieved by pulling the optimal arm a_t^* at every

time step t . The corresponding regret R_T is defined as:

$$R_T = \sum_{t=1}^T (\mu_t^* - X_t(a_t)). \quad (2)$$

The goal is to design algorithms that achieve sub-linear expected regret in T too. In this paper, we consider two commonly-used models of non-stationarity: (i) abruptly changing environment ([32], [33]) and (ii) continuously changing ([34]) environment. The algorithms discussed below for both environments achieve sub-linear expected regret.

a) Abruptly Changing Environment: The reward distributions remain fixed during certain periods and change at unknown time instants called breakpoints. Denote by β_T the total number of breakpoints that occur before time T . As shown in [35], standard bandit algorithms are not appropriate for this environment, and therefore several methods have been proposed. The most relevant to this paper are the two extensions of the UCB ([29]) algorithms: Discounted UCB (DUCB) [32] and Sliding Window UCB (SWUCB) [33].

The DUCB algorithm (**Algorithm 1**) uses a discount factor to emphasize recent rewards when calculating their average. Specifically, the algorithm uses a discount factor $\gamma \in (0, 1)$ to calculate the average of the observed rewards:

$$\bar{X}_t(\gamma, a) = \frac{1}{N_t(\gamma, a)} \sum_{\alpha=1}^t \gamma^{t-\alpha} \mathbf{1}(a_\alpha = a) X_\alpha(a). \quad (3)$$

Here, $N_t(\gamma, a)$ is the discounted frequency of arm a until time t . The algorithm further constructs an upper confidence bound $\bar{X}_t(\gamma, a) + c_t(\gamma, a)$ on the average reward with

$$c_t(\gamma, a) = 2\sqrt{\xi \log n_t(\gamma) / N_t(\gamma, a)} \quad (4)$$

as the discounted confidence radius (for some constant ξ tuned based on the context; see [33] for more details). Note that $n_t(\gamma) = \sum_{i=1}^K N_t(\gamma, a)$ is the sum of the discounted frequencies for all arms until time t . Notice that for $\gamma = 1$, DUCB recovers the UCB1 algorithm.

Algorithm 1 Discounted UCB

- 1: for t from 1 to K , pull arm $a_t = t$;
 - 2: for t from $K + 1$ to T , pull arm a_t which maximizes the upper confidence bound $\bar{X}_t(\gamma, a_t) + c_t(\gamma, a_t)$
-

With the SWUCB algorithm (**Algorithm 2**), instead of averaging rewards over the entire history with a discount factor, averages are computed based on a fixed-size horizon. At each time step t , SWUCB utilizes a local empirical average of the most recent τ arm pulls to construct an upper confidence bound $\bar{X}_t(\tau, a) + c_t(\tau, a)$ for the expected reward. The local empirical average is defined as:

$$\bar{X}_t(\tau, a) = \frac{1}{N_t(\tau, a)} \sum_{\alpha=t-\tau+1}^t \mathbf{1}(a_\alpha = a) X_\alpha(a)$$

with $N_t(\tau, a)$ the frequency of selecting arm a in the last τ arm pulls. The confidence radius is defined as:

$$c_t(\tau, a) = \sqrt{\xi \log(\min(t, \tau)) / N_t(\tau, a)}$$

with some constant ξ (see [33] for more details). Notably, in [33] the authors have demonstrated that both DUCB and SWUCB achieve a regret of $\tilde{O}(\sqrt{\beta_T T})$, where $\tilde{O}(\cdot)$ disregards logarithmic terms.

Algorithm 2 Sliding Window UCB

- 1: for t from 1 to K , pull arm $a_t = t$;
 - 2: for t from $K + 1$ to T , pull arm a_t which maximizes the upper confidence bound $\bar{X}_t(\tau, a_t) + c_t(\tau, a_t)$
-

b) Continuously Changing Environment: Here the number of changes in the mean rewards can potentially be infinite, but the total variation over a relevant time horizon is bounded by a *variation budget*; see, e.g., [34]. Specifically, for a time horizon of T , we define the variation budget V_T as a non-decreasing sequence of positive numbers $\{V_t\}_{t=1}^T$ such that $V_1 = 0$ and $KV_t \leq t$, with K the number of arms.

Recall that $\mu_t(a)$ is the expected regret of arm a at time t . Denote by $\mu(a) = \{\mu_t(a)\}_{t=1}^T$ the sequence of expected rewards of arm a , and $\mu = \{\mu(a)\}_{a=1}^K$ the sequence of expected rewards of all K arms. The set \mathcal{V} of permissible reward sequences for each arm can be written as:

$$\mathcal{V} = \left\{ \mu \in [0, 1]^{K \times T} : \sum_{t=1}^T \sup_{a \in [1, K]} |\mu_t(a) - \mu_{t+1}(a)| \leq V_T \right\}. \quad (5)$$

The above set of permissible reward sequences can capture various scenarios where the expected rewards may change continuously, in discrete shocks, or adhere to a certain rate of change.

For different permissible reward sequences, the achievable regrets may be different. We consider their supremum:

$$R_T^{\mathcal{V}} = \sup_{\mu \in \mathcal{V}} \left\{ \sum_{t=1}^T \mu_t^* - \mathbb{E} \left[\sum_{t=1}^T X_t(a_t) \right] \right\}. \quad (6)$$

In [34] the authors provided a near-optimal algorithm with a worst-case regret of $O(V^{1/3} T^{2/3})$.

C. Incentivized Exploration

As mentioned in Section I, the principal and the agents may have different interests in many real-world scenarios. The principal would like the agents to select the arms in such a way as to adequately explore different arms to maximize the accumulated rewards. An agent, however, influenced by the feedback of others, behaves myopically in the face of uncertainty, i.e., pulls the arm with the currently highest empirical reward (exploitation only).

Similar to [21], we consider a variant of the MAB problem where a principal aims to incentivize the agents to explore. At each time step t , an agent pulls one arm a_t based on the recommendation of the principal. The agent receives a reward r_t , which is then fed back to the principal and the agents. The principal uses a certain bandit algorithm to find the ‘optimal’ arm while balancing exploration and exploitation. When the principal wants to encourage agents to explore, they may offer compensation χ_t to the agents.

This compensation motivates the agents to follow a specific bandit algorithm that balances exploration and exploitation, ultimately maximizing their cumulative rewards.

However, because the agent receives compensation, their feedback from pulling the arm can become biased. This bias may introduce a deviation δ_t on top of the "true" reward $X_t(a_t)$. This deviation is influenced by some unknown, non-decreasing function f_t of the compensation χ_t . This function is assumed to possess the following characteristics.

Assumption 1 ([21]) *The reward drift function $f_t(x)$ is non-decreasing with $f_t(0) = 0$ and is Lipschitz continuous, i.e., there exists a constant l_t such that $|f_t(x) - f_t(y)| \leq l_t|x - y|$ for any x and y .*

Note that the received reward r_t is the biased feedback (i.e., equal to the sum $X_t(a_t) + \delta_t$), and the principal and agents cannot distinguish either $X_t(a_t)$ or δ_t from it.

Denote by g_t the greedy choice at time t , and note that the actual arm pulled, a_t , is the arm recommended by the principal. Let $\bar{X}_t(a)$ be the empirical average of the rewards of some arm a until time t . Along with the regret, the principal is also concerned with the total compensation he has to pay:

$$C_T = \sum_{t=1}^T (\bar{X}_t(g_t) - \bar{X}_t(a_t)). \quad (7)$$

We will characterize the efficacy of incentivized exploration in terms of both *expected* regret and *expected* compensation and aim to answer the following question: if and how can we design algorithms that achieve both sublinear regret and sublinear compensation? The authors in [21] have studied this important question in the setting of a *stationary* bandits and proposed algorithms that achieve $O(\log T)$ regret and compensation. In contrast, in this paper, we investigate the more challenging setting of *non-stationary* bandits.

III. INCENTIVIZED EXPLORATION IN NON-STATIONARY BANDITS

In this section, we design algorithms for the incentivized exploration for *nonstationary* bandits and show that they achieve sublinear regret and compensation.

A. Incentivized Exploration in the Abruptly-Changing Environment

Algorithm 3 describes a framework of incentivized exploration for the abruptly changing environment. At time t , the principal recommends an arm a_t (line 2) based on a non-stationary bandit algorithm (e.g., DUCB or SWUCB), and the greedy choice is denoted by g_t (line 3). The principal offers compensation χ_t (in line 5) to the agents, which is the difference between the empirical average of rewards from the greedy choice and the principal's recommendation when they differ. This compensation is provided if the principal's recommended arm doesn't align with the greedy choice. After receiving this compensation, the player's outcome is

Algorithm 3 Incentivized MAB under Reward Drift

```

1: for  $t \in [1, T]$  do
2:    $a_t \leftarrow$  Principal's Recommendation
3:    $g_t \leftarrow \arg \max_{a \in [1, K]} \bar{X}_t(a)$ 
4:   if  $a_t \neq g_t$  then
5:     Principal offers compensation of  $\chi_t \leftarrow \bar{X}_t(g_t) - \bar{X}_t(a_t)$  to the agent.
6:     Reward for pulling arm  $a_t$  is  $r_t \leftarrow X_t(a_t) + \delta_t$  where reward drift  $\delta_t \leftarrow f(\chi_t)$ 
7:   else
8:      $r_t \leftarrow X_t(a_t)$  is the reward with no compensation.
9:   end if
10: end for

```

affected by a bias δ_t , which is added to the "true" reward $X_t(a_t)$.

With equation (2), the expected regret is defined as

$$\mathbb{E} \left[\sum_{t=1}^T (\mu^* - X_t(a_t)) \right] = \sum_{a \neq a^*} (\mu^* - \mu(a)) \mathbb{E}[N_T(a)]. \quad (8)$$

Since the expected reward of an arm is in the range $[0, 1]$, we have $(\mu^* - \mu(a)) \leq 1$ for all a . Therefore, bounding the expected regret after T pulls essentially amounts to controlling the expected number of times a sub-optimal arm is pulled. In Theorem 1, we bound the expected number of times some sub-optimal arms $a \neq a_t^*$ are pulled when the principal uses DUCB algorithm to balance exploration and exploitation in Algorithm 3 (line 2) until time T .

Theorem 1 (Algorithm 3 + DUCB Regret Bound)

Given the time horizon T and the number of breakpoints β_T , the expected number of times some sub-optimal arms $a \neq a_t^$ are pulled is bounded as follows:*

$$\mathbb{E}[N_T(a)] \leq \tilde{\eta} \cdot \sqrt{T\beta_T} \log(T) \quad (9)$$

with some constant $\tilde{\eta} > 0$.

See the proof of Theorem 1 in the Appendix for the choice of the discount factor γ .

The following result is for the case when the principal uses the SWUCB algorithm instead in Algorithm 3 (line 2).

Theorem 2 (Algorithm 3 + SWUCB Regret Bound)

Given the time horizon T and the number of breakpoints β_T , the expected number of times some sub-optimal arms $a \neq a_t^$ are pulled is bounded as follows:*

$$\mathbb{E}[N_T(a)] \leq \tilde{\eta} \cdot \sqrt{\beta_T T \log(T)} \quad (10)$$

with some constant $\tilde{\eta} > 0$.

See the proof of Theorem 2 in the Appendix for the choice of the sliding window τ .

Remark 1 *The lower bound of the regret for an algorithm scheme for the abruptly changing environment is $\Omega(\sqrt{T})$ (see section 4 of [33]). Therefore, the proposed algorithm scheme*

is optimal up to some $\log T$ powers, besides the dependence on β_T .

Now, let us take a look at the total expected compensation. Consistent with the definition (7), in the non-stationary setting the total compensation is defined as follows: $C_T = \sum_{t=1}^T (\bar{X}_t(\gamma, g_t) - \bar{X}_t(\gamma, a_t))$ when using DUCB with discount factor γ , and $C_T = \sum_{t=1}^T (\bar{X}_t(\tau, g_t) - \bar{X}_t(\tau, a_t))$ when using SWUCB with sliding window τ .

Theorem 3 (Algorithm 3 + DUCB Compensation)

Given the time horizon T and the number of breakpoints β_T , the total expected compensation contributed by the arm a when the principal uses the DUCB algorithm is bounded as follows:

$$\mathbb{E}[C_T(a)] \leq \eta \cdot \beta_T^{3/2} \sqrt{T} (\log(T))^{3/2} \quad (11)$$

with some constant $\eta > 0$.

Theorem 4 (Algorithm 3 + SWUCB Compensation)

Given the time horizon T and the number of breakpoints β_T , the total expected compensation contributed by the arm a when the principal uses the SWUCB algorithm is bounded as follows:

$$\mathbb{E}[C_T(a)] \leq \eta \cdot (\beta_T)^{7/4} T^{1/4} (\log(T))^{3/4} \quad (12)$$

with some constant $\eta > 0$.

The proofs of the above theorems can be found in the Appendix.

B. Incentivized Exploration in the Continuously-Changing Environment

For the continuously changing environment, we divide the time horizon T into $\lfloor T/\sigma \rfloor$ batches of certain fixed size σ . At the start of each batch, the principal restarts a certain bandit algorithm to recommend arms to the agent. We

Algorithm 4 Restarting technique with a MAB algorithm

Require: σ, K, T

Ensure: $j = 1$

- 1: **while** $j \leq \lfloor T/\sigma \rfloor$ **do**
 - 2: $\alpha \leftarrow (j - 1)\sigma$
 - 3: **for** $t = 1, \dots, \min\{T, \alpha + \sigma\}$ **do**
 - 4: $a_t \leftarrow$ Principal's recommendation
 - 5: $g_t \leftarrow \arg \max_{a \in [1, K]} \bar{X}_t(a)$
 - 6: steps 4-9 from **Algorithm 3**
 - 7: **end for**
 - 8: Increment $j \leftarrow j + 1$ and return at line 1.
 - 9: **end while**
-

present Algorithm 4 where for a single batch the principal employs certain MAB algorithm such as UCB1, ϵ -Greedy or Thompson Sampling and recommends an arm to the agent (line 4) to balance exploration and exploitation. Based on the greedy choice of the agent (line 5), the compensation scheme is determined according to lines 4-9 of Algorithm 3. This process is repeated for $\lfloor T/\sigma \rfloor$ iterations.

Theorem 5 (Algorithm 4 Regret) Given the time horizon T and the variation budget V_T , if the principal employs UCB1, ϵ -greedy or Thompson Sampling for recommending arms to the agent in Algorithm 4, then the worst-case regret over the time horizon T is bounded as follows:

$$R_V^T \leq \eta \cdot V_T^{1/3} (K \log(T))^{1/3} T^{2/3} \quad (13)$$

with some constant $\eta > 0$.

See the proof of Theorem 5 in the Appendix for the choice of the batch size σ .

The overall compensation C_V^T for the entire time horizon T is calculated by summing up the compensation of each batch, following the definition (7).

Theorem 6 Given the time horizon T and the variation budget V_T , if the principal employs UCB1 for recommending arms to agents in Algorithm 4, then the worst-case total compensation is bounded as follows:

$$C_V^T \leq \eta_1 \cdot (KV_T \log(T))^{1/3} T^{2/3} \quad (14)$$

with some constant $\eta_1 > 0$. If the principal employs ϵ -greedy, then the worst-case total compensation is bounded as follows:

$$C_V^T \leq \eta_2 \cdot (KV_T \log(T))^{2/3} T^{1/3} \quad (15)$$

with some constant $\eta_2 > 0$. If the principal employs Thompson Sampling, then the worst-case total compensation is bounded as follows:

$$C_V^T \leq \eta_3 \cdot (KV_T \log(T))^{2/3} T^{1/3} \quad (16)$$

with some constant $\eta_3 > 0$.

See the Appendix for the proof of the above theorem.

IV. NUMERICAL EXPERIMENTS

We evaluate the performance of the algorithms proposed in Section III numerically.

A. Abruptly Changing Environment

Consider a setting of two arms (i.e., $K = 2$), indexed by 1 and 2), with the expected rewards 0.99 and 0.01, respectively. The expected rewards flip (i.e., swaps values) at every breakpoint that is evenly located at kT/p for $k = 1, 2, \dots, p-1$ for some integer $p > 0$. For instance, we have the breakpoints at $\lfloor T/3 \rfloor$ and $\lfloor 2T/3 \rfloor$ for $p = 3$,

In Algorithm 3, the discount factor $\gamma = 1 - (1/\gamma_C)\sqrt{\beta_T/T}$ for DUCB is chosen according to the analysis of Theorem 1 and the window size $\tau = \lfloor \tau_C \sqrt{T \log(T)/\beta_T} \rfloor$ for SWUCB according to Theorem 2. We run a hundred repetitions to obtain the average values of regret and compensation.

Figures 2 and 3 present the performance of Algorithm 3 with DUCB and SWUCB as submodules with $T = 5000$. Both algorithms outperform the UCB1 with Algorithm 3. The frequent changes force the UCB1 to make mistakes at the start of each breakpoint, as it considers the *entire history*. However, DUCB considers the decaying history, giving more importance to the recent past; and SWUCB considers a

Regret & Compensation (D-UCB/SW-UCB/UCB)

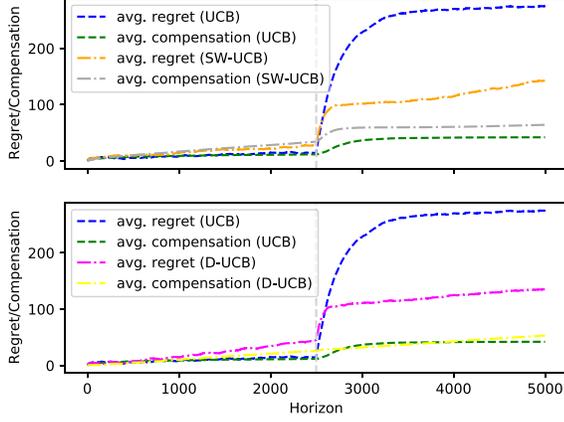


Fig. 2: (Upper) Regret and Compensation performance of DUCB with Algorithm 3 with $\gamma_C = 10$ (Below) Regret and Compensation performance of SWUCB with Algorithm 3 with $\tau_C = 0.9$, both with $T = 5000$ and $\beta_T = 1$

Regret & Compensation (D-UCB/SW-UCB/UCB)

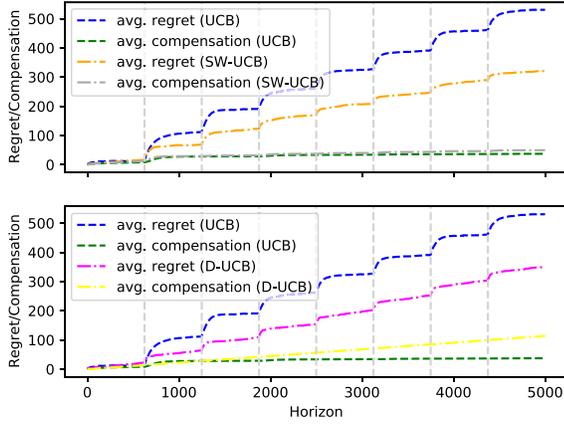


Fig. 3: (Upper) Regret and Compensation performance of DUCB with Algorithm 3 with $\gamma_C = 40$ (Below) Regret and Compensation performance of SWUCB with Algorithm 3 with $\tau_C = 1$, both with $T = 5000$ and $\beta_T = 1$

sliding window and thus adjusts quickly to changes in the reward distribution, leading to lower regret.

Table II summarizes the performance of SWUCB and DUCB with Algorithm 3, respectively, with varying breakpoints. We present the corresponding parameters τ_C and γ_C , which minimize the regrets. All the regret and compensation values are within the theoretical bounds. Besides, the regret is consistently lower than the UCB1 counterpart for DUCB and SWUCB, as all the parameters $\gamma, \gamma_C, \tau, \tau_C$ are tuned to minimize regret. The values considered are $\tau_C = [10, 20, 30, 40]$ and $\gamma_C = [0.9, 0.95, 1, 2]$ through experimentation.

In both cases the regret is growing in the order of $O(\sqrt{\beta_T})$ with varying breakpoints, as the theoretical analysis suggests. As for compensation, SWUCB increases more rapidly than DUCB, which explains the higher sensitivity of SWUCB to the number of breakpoints, which is in the order

TABLE II: The performance of Algorithm 3 with varying number of breakpoints β_T . The subscripts U, D, S stand for UCB1, DUCB, and SWUCB, respectively, with R as the regret and C as the compensation values.

β_T	γ_C	τ_C	R_U	R_S	R_D	C_U	C_S	C_D
1	15	1	275.2	135.1	142.7	42.1	53.2	64.2
2	10	1	364.2	203.5	205.7	42.5	70.7	92.3
3	15	1	430.4	239.5	247.1	41.6	81.2	82.8
4	15	0.95	394.8	264.1	259.7	41.4	95.1	89.2
5	10	1	423.7	288.9	302.4	39.8	100.8	112.3
6	25	1	481.8	330.1	279.1	38.5	107.9	67.6
7	30	0.95	484.2	339.0	299.7	38.6	117.1	59.2

of $O(\beta_T^{7/4})$ compared to $O(\beta_T^{3/2})$.

B. Continuously Changing Environment

Consider a setting of two arms, indexed by 1 and 2. The received (i.e., instantaneous) reward $X_t(a_t)$ for arm a_t at time t is modeled as a Bernoulli random variable with a changing expectation $\mu_t(a_t)$

$$X_t(a_t) = \begin{cases} 1 & \text{with prob. } \mu_t(a_t) \\ 0 & \text{with prob. } 1 - \mu_t(a_t) \end{cases} \quad (17)$$

for all $t \leq T$ and for any pulled arm $a_t \in [1, K]$. We have two sinusoidal evolution patterns (Fig. 4; inspired by [34]) with a variation budget $V_T = 3$ for $\mu_t(a_t)$ (the green dotted line is for arm 1 as μ_1 and the yellow dotted line is for arm 2 as μ_2 in Fig. 4). They describe different changing environments under the same variation budget. In the first (i.e., left half) instance, the variation budget is spent throughout the whole horizon, while in the second, the same variation budget is spent only over the first third of the whole horizon. In this experiment, at each time step t , the following happens in order: (i) the algorithm selects an arm $a_t \in [1, K]$, (ii) the binary rewards are generated based on (17), and (iii) the instantaneous reward $X_t(a_t)$ is observed. We denote by $a_t^* = \arg \max_{a \in [1, K]} \mu_t(a)$ as the optimal arm at time t . The instantaneous regret at time t is $X_t(a_t^*) - X_t(a_t)$. We run the experiment with multiple (in this case 2000 times) repetitions to obtain average values of regret and compensation.

Fig. 5 presents the performance of Algorithm 4 with UCB1, ϵ -Greedy and Thompson Sampling with regards to *regret* and *compensation*, whereas Fig. 4 presents the total rewards accumulated (i.e., averaged over 2000 iterations) for the same settings. Table III presents more data about the same performance with varying degrees of variation budget V_T . Note that the regret for all the algorithm schemes varies (i.e., increase) by a maximum of $V_T^{1/3}$ as suggested by the theoretical analysis. compensation increases more quickly for ϵ -Greedy and Thompson Sampling compared to UCB1, as UCB1 is comparatively less sensitive to V_T . Notice that ϵ -greedy does not drastically change its compensation values compared to Thompson Sampling, suggesting that the upper bound can be tightened.

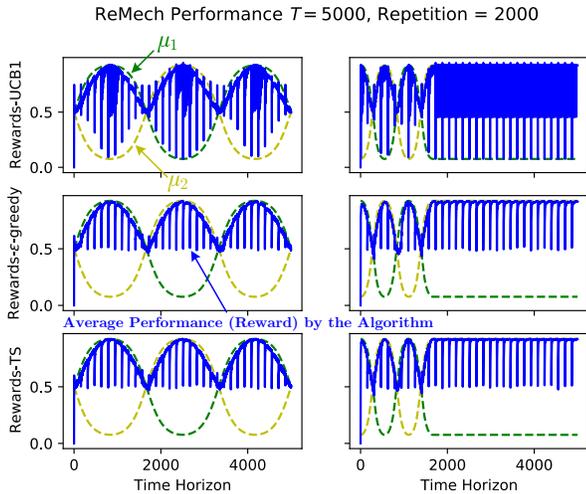


Fig. 4: Algorithm 4 (written as ReMech, shorthand for restarting mechanism, in the diagram) performance with $T = 5000$ with 2000 repetitions. The blue curve traces the total reward accumulated (averaged over all iterations) with ALGORITHM 4 at various time steps with UCB1, ϵ -greedy, and Thompson Sampling as respective submodules.

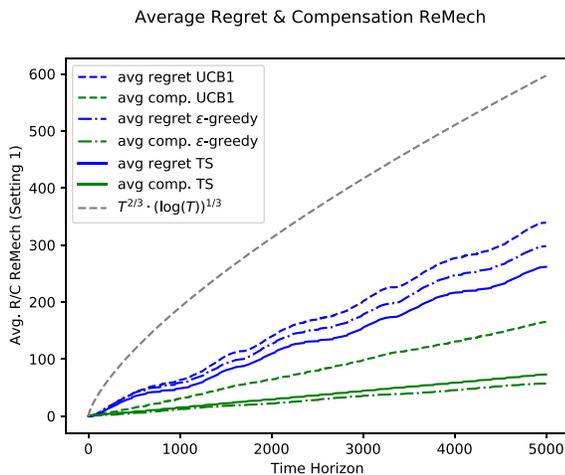


Fig. 5: Algorithm 4 performance with submodules of UCB1, ϵ -greedy, and Thompson Sampling, for a large horizon with $T = 5000$ with 2000 repetitions.

TABLE III: The performance of Algorithm 4 with different variation budgets. The subscripts $U, \epsilon G, T$ stand for UCB1, ϵ -Greedy and Thompson Sampling, respectively, with R as the regret and C as the compensation values.

V_T	3	6	9	12	15	18	24
R_U	156.1	175.9	185.7	191.4	198.3	207.5	210.3
C_U	88.9	107.4	119.8	127.2	135.0	145.6	149.8
$R_{\epsilon G}$	143.1	164.1	180.2	192.0	202.3	215.0	229.0
$C_{\epsilon G}$	37.3	52.4	64.1	73.6	80.7	88.7	99.5
R_T	125.1	147.2	163.8	177.8	185.9	197.8	211.6
C_T	48.4	69.0	84.9	97.5	107.5	118.1	132.2

We have studied the incentivized exploration for the MAB problem with non-stationary reward distributions, where the players receive compensation for exploring arms other than the greedy choice and may provide biased feedback on the reward. We consider two different environments that capture rewards' non-stationarity, and propose incentivized exploration algorithms accordingly. We show that the proposed algorithms achieve sublinear regret and compensation in time and thus are effective in incentivizing exploration despite the nonstationarity and the drifted feedback.

REFERENCES

- [1] J. C. Gittins, "Bandit processes and dynamic allocation indices," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 41, no. 2, pp. 148–177, 1979. [Online]. Available: <http://www.jstor.org/stable/2985029>
- [2] D. A. Berry and B. Fristedt, *Bandit problems : sequential allocation of experiments / Donald A. Berry, Bert Fristedt.* Chapman and Hall London ; New York, 1985.
- [3] W. H. Press, "From the Cover: Bandit solutions provide unified ethical models for randomized clinical trials and comparative effectiveness research," *Proceedings of the National Academy of Science*, vol. 106, no. 52, pp. 22 387–22 392, Dec. 2009.
- [4] E. Brochu, M. W. Hoffman, and N. de Freitas, "Portfolio Allocation for Bayesian Optimization," *arXiv e-prints*, p. arXiv:1009.5419, Sep. 2010.
- [5] D. Bouneffouf, A. Bouzeghoub, and A. L. Gançarski, "A contextual-bandit algorithm for mobile context-aware recommender system," in *ICONIP*, 2012.
- [6] L. Li, W. Chu, J. Langford, and R. E. Schapire, "A contextual-bandit approach to personalized news article recommendation," *Proceedings of the 19th international conference on World wide web - WWW '10*, 2010. [Online]. Available: <http://dx.doi.org/10.1145/1772690.1772758>
- [7] F. Radlinski, R. Kleinberg, and T. Joachims, "Learning diverse rankings with multi-armed bandits," in *Proceedings of the 25th International Conference on Machine Learning*, ser. ICML '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 784–791. [Online]. Available: <https://doi.org/10.1145/1390156.1390255>
- [8] Y. Gai, B. Krishnamachari, and R. Jain, "Learning multiuser channel allocations in cognitive radio networks: A combinatorial multi-armed bandit formulation," 05 2010, pp. 1 – 9.
- [9] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," 2012.
- [10] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: A Bradford Book, 2018.
- [11] P. Frazier, D. Kempe, J. Kleinberg, and R. Kleinberg, "Incentivizing exploration," in *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, ser. EC '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 5–22. [Online]. Available: <https://doi.org/10.1145/2600057.2602897>
- [12] Y. Mansour, A. Slivkins, and V. Syrgkanis, "Bayesian incentive-compatible bandit exploration," in *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, ser. EC '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 565–582. [Online]. Available: <https://doi.org/10.1145/2764468.2764508>
- [13] S. Wang and L. Huang, "Multi-armed bandits with compensation," in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018. [Online]. Available: <https://proceedings.neurips.cc/paper/2018/file/8bdb5058376143fa358981954e7626b8-Paper.pdf>
- [14] N. Immorlica, J. Mao, A. Slivkins, and S. Wu, "Bayesian exploration with heterogeneous agents," in *The Web Conference 2019*, May 2019. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/bayesian-exploration-with-heterogeneous-agents/>

- [15] C. Hirsenschall, A. Singla, S. Tschischek, and A. Krause, "Learning user preferences to incentivize exploration in the sharing economy," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/11874>
- [16] L. Han, D. Kempe, and R. Qiang, "Incentivizing exploration with heterogeneous value of money," *Lecture Notes in Computer Science*, p. 370–383, 2015. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-48995-6_27
- [17] Y. Liu and C.-J. Ho, "Incentivizing high quality user contributions: New arm generation in bandit learning," 2018. [Online]. Available: <https://aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16879>
- [18] A. Martensen, L. Gronholdt, and K. Kristensen, "The drivers of customer satisfaction and loyalty: Cross-industry findings from denmark," *Total Quality Management*, vol. 11, no. 4-6, pp. 544–553, 2000. [Online]. Available: <https://doi.org/10.1080/09544120050007878>
- [19] Z. Ehsani and M. Ehsani, "Effect of quality and price on customer satisfaction and commitment in iran auto industry," 2015.
- [20] L. Gwo-Guang and L. Hsiu-Fen, "Consumer perceptions of e-service quality in online shopping," *International Journal of Retail & Distribution Management*, vol. 33, pp. 161–176, 02 2005.
- [21] Z. Liu, H. Wang, F. Shen, K. Liu, and L. Chen, "Incentivized exploration for multi-armed bandits under reward drift," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, pp. 4981–4988, Apr. 2020. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/5937>
- [22] I. Kremer, Y. Mansour, and M. Perry, "Implementing the wisdom of the crowd," vol. 122, 06 2013, pp. 605–606.
- [23] Y.-K. Che and J. Hörner, "Recommender systems as mechanisms for social learning*," *Quarterly Journal of Economics*, vol. 133, pp. 871–925, 05 2018.
- [24] Y. Mansour, A. Slivkins, and V. Syrgkanis, "Bayesian incentive-compatible bandit exploration," 2019.
- [25] L. Cohen and Y. Mansour, "Optimal algorithm for bayesian incentive-compatible exploration," in *Proceedings of the 2019 ACM Conference on Economics and Computation*, ser. EC '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 135–151. [Online]. Available: <https://doi.org/10.1145/3328526.3329581>
- [26] M. Sellke and A. Slivkins, "The price of incentivizing exploration: A characterization via thompson sampling and sample complexity," 2021.
- [27] E. Kamenica and M. Gentzkow, "Bayesian persuasion," *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, October 2011. [Online]. Available: <https://www.aeaweb.org/articles?id=10.1257/aer.101.6.2590>
- [28] A. Slivkins, "Introduction to multi-armed bandits," 2021.
- [29] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, pp. 235–256, 2002.
- [30] T. L. Lai, H. Robbins *et al.*, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [31] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, Z. Wen *et al.*, "A tutorial on thompson sampling," *Foundations and Trends® in Machine Learning*, vol. 11, no. 1, pp. 1–96, 2018.
- [32] L. Kocsis and C. Szepesvari, "Discounted ucb pascal challenges workshop," *Venice, Italy (April 2006)*, 2006.
- [33] A. Garivier and E. Moulines, "On upper-confidence bound policies for non-stationary bandit problems," 2008.
- [34] O. Besbes, Y. Gur, and A. Zeevi, "Stochastic multi-armed-bandit problem with non-stationary rewards," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014. [Online]. Available: <https://proceedings.neurips.cc/paper/2014/file/903ce9225fca3e988c2af215d4e544d3-Paper.pdf>
- [35] C. Hartland, S. Gelly, N. Baskiotis, O. Teytaud, and M. Sebag, "Multi-armed bandit, dynamic environments and meta-bandits," 2006.

VI. APPENDIX

A. Material for Theorem 1

In this section we discuss the details pertaining to the incentivized exploration on abruptly changing environment,

specifically Theorem 1 (i.e., regret result for Algorithm 3 + DUCB). We start with laying out some fundamental results (i.e., lemma 1 and 2) required for the analysis later. After that, we present lemma 4 which directly leads to the main result.

Recall that $N_t(\gamma, a)$ is the discounted frequency (i.e., number of pulls) of arm a till time t and the sum of this quantity of all the arms is denoted by $n_t(\gamma)$. In the following lemma we upper bound the same.

Lemma 1 $n_t(\gamma) \leq \min(t, 1/(1-\gamma))$, for all $t \leq T$.

Proof: According to the definition, we have the following;

$$\begin{aligned} n_t(\gamma) &= \sum_{a=1}^K N_t(\gamma, a) \\ &= \sum_{a=1}^K \sum_{s=1}^t \mathbf{1}(a_s = a) \gamma^{t-s} \\ &= \sum_{s=1}^t \gamma^{t-s} \leq \sum_{s=1}^{\infty} \gamma^s = \frac{1}{1-\gamma}. \end{aligned} \quad (18)$$

Notice that $n_t(\gamma)$ cannot be more than t , we get the final upper bound. ■

Next, we look at the total discounted reward drift $D_t(\gamma, a)$ contributed by some arm a till some time t . Note that l_t is the Lipschitz constant corresponding to the drift function, at any given time t (Refer Assumption 1 in the main paper), and $N_t(a)$ as the total pulls (i.e., essentially $N_t(\gamma, a)$ for $\gamma = 1$) for arm a till time t . We define $l = \max_t l_t$ for the following lemma.

Lemma 2 $D_t(\gamma, a) \leq 2lN_t(a)\sqrt{\xi \log(n_t(\gamma))}$ for all arms $a \in [1, K]$ and $t \in [1, T]$.

Proof:

Notice that there is a drift in the reward from the agent only when the principal provides some compensation. The inequalities mentioned below quantify this event, and they must hold true simultaneously. Since the agent is assumed to be greedy in nature, the arm g_t (i.e., the arm with the highest empirical average till time t) seems more appealing than a_t (i.e., the principal's recommended arm at time t). The equality below represent this event.

$$\bar{X}_t(\gamma, g_t) \geq \bar{X}_t(\gamma, a_t), \quad (19)$$

However, the principal recommends arm a_t as per its algorithm which balances exploration and exploitation. The inequality below quantifies this event.

$$\bar{X}_t(\gamma, g_t) + c_t(\gamma, g_t) \leq \bar{X}_t(\gamma, a_t) + c_t(\gamma, a_t). \quad (20)$$

The above two inequalities imply $\bar{X}_t(\gamma, g_t) - \bar{X}_t(\gamma, a_t) \leq c_t(\gamma, a_t) - c_t(\gamma, g_t)$. Since the drift function is Lipschitz smooth (Assumption 1 in the main paper) we can upper bound the instantaneous drift $\delta_t(a_t)$ for the pulled arm a_t at time t as below,

$$\begin{aligned}
\delta_t(a_t) &\leq l_t (c_t(\gamma, a_t) - c_t(\gamma, g_t)) \\
&\leq l_t \left(2\sqrt{\frac{\xi \log(n_t(\gamma))}{N_t(\gamma, a_t)}} \right) = 2l_t \sqrt{\frac{\xi \log(n_t(\gamma))}{N_t(\gamma, a_t)}}.
\end{aligned} \tag{21}$$

Therefore, summing the instantaneous drift till some time t , for some arm a gives us its total drift contribution till said t , as shown below.

$$\begin{aligned}
D_t(\gamma, a) &= \sum_{s=1}^t \mathbf{1}(a_s = a) \delta_s(a) \\
&= \sum_{s=1}^t \mathbf{1}(a_s = a) 2l_s \sqrt{\frac{\xi \log(n_t(\gamma))}{N_s(\gamma, a)}} \\
&\leq 2l \sqrt{\xi \log(n_t(\gamma))} \sum_{s=1}^t \mathbf{1}(a_s = a) \frac{1}{\sqrt{N_s(\gamma, a)}} \\
&\leq 2l \sqrt{\xi(1-\gamma) \log(n_t(\gamma))} \sum_{s=1}^{N_t(a)} 1 \\
&\leq 2l N_t(a) \sqrt{\xi \log(n_t(\gamma))}
\end{aligned} \tag{22}$$

Recall from section 2-A that since the expected reward of an arm is in the range $[0, 1]$, we get $(\mu^* - \mu(a)) \leq 1$ for all arms $a \in [1, K]$. Therefore, bounding the expected regret after T pulls essentially amounts to controlling the expected number of times a sub-optimal arm is pulled. In the next lemma, we bound that expected number of total pulls $N_T(a)$ of some sub-optimal arm $a \neq a_t^*$ for the entire horizon $t \leq T$.

Lemma 3 *Let $\xi > 1/2$, $T > 1$ and $\gamma \in (0, 1)$. For any arm $a \in [1, K]$ and total number of breakpoints β_T , we have*

$$\mathbb{E}[N_T(a)] \leq \left(B(\gamma)T(1-\gamma) + C(\gamma) \frac{\beta_T}{1-\gamma} \right) \log \left(\frac{1}{1-\gamma} \right) \tag{23}$$

where

$$\begin{aligned}
B(\gamma) &= \frac{16(1-\gamma)\xi [T(1-\gamma)]}{F(\gamma) T(1-\gamma)} \\
&\quad + \frac{2 \left[-\log(1-\gamma) / \log(1 + 4\sqrt{1-1/2\xi}) \right]}{-\log(1-\gamma)(1-\gamma^{1/(1-\gamma)})}
\end{aligned} \tag{24}$$

and

$$F(\gamma) = \gamma^{1/(1-\gamma)} \left(\Delta_{\mu_T}(i) \sqrt{1-\gamma} - 4l \sqrt{-\xi \log(1-\gamma)} \right)^2 \tag{25}$$

and

$$C(\gamma) = \frac{(\gamma-1) \log((1-\gamma)\xi \log(n_K(\gamma)))}{\log(1-\gamma) \log(\gamma)} \tag{26}$$

Proof: Let us denote by $\Delta_{\mu_T}(a_t)$ the minimum of the difference of the expected reward $\mu_t(a_t^*)$ of the best arm at

time t and the expected reward $\mu_t(a_t)$ of the arm a_t for all $t \in [1, T]$ when a_t is not the optimal arm, i.e.,

$$\Delta_{\mu_T}(a_t) = \min_{t \in [1, T]; a_t \neq a_t^*} (\mu_t(a_t^*) - \mu_t(a_t)). \tag{27}$$

The quantity $N_T(a)$ for some arm a can be represented as the following.

$$\begin{aligned}
N_T(a) &= \sum_{t=1}^T \mathbf{1}(a_t = a \neq a_t^*) = 1 + \sum_{t=K+1}^T \mathbf{1}(a_t = a \neq a_t^*) \\
&= 1 + H + G
\end{aligned} \tag{28}$$

where $H = \sum_{t=K+1}^T \mathbf{1}(a_t = a \neq a_t^*; N_t(\gamma, a) < A(\gamma))$ and $G = \sum_{t=K+1}^T \mathbf{1}(a_t = a \neq a_t^*; N_t(\gamma, a) \geq A(\gamma))$ and where

$$A(\gamma) = \frac{16(1-\gamma)\xi \log(n_t(\gamma))}{\left(\Delta_{\mu_T} \sqrt{1-\gamma} - 4l \sqrt{(1-\gamma)\xi \log(n_t(\gamma))} \right)^2}. \tag{29}$$

The next few steps directly follow from the analysis in [33]. For the same definitions of $D(\gamma)$ and $\mathcal{T}(\gamma)$ from [33], we can bound $N_T(a)$ by:

$$\begin{aligned}
N_T(a) &\leq 1 + [T(1-\gamma)]A(\gamma)\gamma^{-1/(1-\gamma)} + \beta_T D(\gamma) + \\
&\quad \sum_{t \in \mathcal{T}(\gamma)} \mathbf{1}(a_t = a \neq a_t^*; N_t(\gamma, a) \geq A(\gamma))
\end{aligned} \tag{30}$$

Now, for some time t the event $\{G : \mathbf{1}(a_t = a \neq a_t^*; N_t(\gamma, a) \geq A(\gamma))\}$ will occur when the algorithm has the following inequality true, for the arm $a \neq a_t^*$ which is labelled as the event \mathcal{Z} ;

$$\mathcal{Z} : \bar{X}_t(\gamma, a) + c_t(\gamma, a) > \bar{X}_t(\gamma, a_t^*) + c_t(\gamma, a_t^*). \tag{31}$$

Let us denote by $\bar{Y}_t(\gamma, a)$ the discounted empirical average of arm a excluding the drift at each time till t (i.e., pure empirical average, different from $\bar{X}_t(\gamma, a)$). We have the following inequalities from \mathcal{Z} ;

$$\bar{Y}_t(\gamma, a) + \frac{D_t(\gamma, a)}{N_t(a)} + c_t(\gamma, a) > \bar{Y}_t(\gamma, a_t^*) + \frac{D_t(\gamma, a_t^*)}{N_t(a_t^*)} + c_t(\gamma, a_t^*). \tag{32}$$

Therefore, the upper bound for the difference between the expected reward of the optimal arm a_t^* and the current arm a is

$$\begin{aligned}
\bar{Y}_t(\gamma, a_t^*) - \bar{Y}_t(\gamma, a) &< \frac{D_t(\gamma, a)}{N_t(a)} + c_t(\gamma, a) \\
&= 2\sqrt{\xi \log(n_t(\gamma))} \left(l + \frac{1}{\sqrt{N_t(\gamma, a)}} \right).
\end{aligned} \tag{33}$$

Now, we can decompose the event G as the following: For \mathcal{Z} to happen, at least one of the events G_t^i has to occur:

$$\therefore \mathbf{1}(a_t = a \neq a_t^*; N_t(\gamma, a) \geq A(\gamma)) \subseteq G_t^1 \cup G_t^2 \cup G_t^3, \tag{34}$$

where

$$G_t^1 = \{\bar{X}_t(\gamma, a) > \mu_t(\gamma, a) + c_t(\gamma, a)\}, \quad (35)$$

$$G_t^2 = \{\bar{X}_t(\gamma, a_t^*) < \mu_t(\gamma, a_t^*) - c_t(\gamma, a_t^*)\}, \quad (36)$$

and

$$G_t^3 = \hat{\mu}_t(\gamma, a_t^*) - \hat{\mu}_t(\gamma, a) < 2\sqrt{\xi \log(n_t(\gamma))} \left(l + \frac{1}{\sqrt{N_t(\gamma, a)}} \right) \quad (37)$$

The event G_t^1 occurs when the algorithm *overestimates* the average reward the pulled arm a , the event G_t^2 occurs when the algorithm *underestimates* the average reward of the best arm a^* , and G_t^3 occurs when the difference between expected rewards for both the arms a and a_t^* are too small.

Note that for the choice of $A(\gamma)$ and the definition of $\Delta_{\mu_T}(a)$, the event G_t^3 never occurs, as

$$\begin{aligned} \bar{Y}_t(\gamma, a^*) - \bar{Y}_t(\gamma, a) &< 2\sqrt{\xi \log(n_t(\gamma))} \left(l + \frac{1}{\sqrt{N_t(\gamma, a)}} \right) \\ &\leq 2\sqrt{\xi \log(n_t(\gamma))} \left(l + \frac{1}{\sqrt{A(\gamma)}} \right) \\ &= \frac{\Delta_{\mu_T}(a)}{2} \end{aligned}$$

Therefore, we have the probability of the event G_t^3 as $\mathbb{P}[G_t^3] = 0$ and from [33] we get the corresponding probability for G_t^1 and G_t^2 as

$$\mathbb{P}[G_t^1] = \mathbb{P}[G_t^2] \leq \left[\frac{\log(n_t(\gamma))}{\log(1 + \eta)} \right] n_t(\gamma)^{-2\xi(1 - \frac{\eta^2}{16})} \quad (38)$$

We have the final bound, by taking $\xi > 1/2$ and $\eta = 4\sqrt{1 - (1/2\xi)}$, so as to make $2\xi(1 - \eta^2/16) = 1$:

$$\mathbb{E}[N_T(a)] \leq 1 + [T(1 - \gamma)]A(\gamma)\gamma^{-1/(1-\gamma)} + \beta_T D(\gamma) + Y$$

where

$$Y = \frac{1}{1 - \gamma} + \left[\frac{\log\left(\frac{1}{1-\gamma}\right)}{\log(1 + 4\sqrt{1 - (1/2\xi)})} \right] \frac{T(1 - \gamma)}{1 - \gamma^{1/(1-\gamma)}} \quad (39)$$

We obtain the statement of the lemma by substituting the values of $A(\gamma)$, $D(\gamma)$, and $n_t(\gamma)$. ■

Note that lemma 4 has done most of the heavy lifting, and just choosing an appropriate discount factor (i.e., $\gamma = 1 - \eta \cdot \sqrt{\beta_T/T}$) to minimize the expression gives us the result for Theorem 1.

B. Material for Theorem 2

In this section, we discuss the details pertaining to the incentivized exploration of an abruptly changing environment, specifically Theorem 2 (i.e., the regret result for Algorithm 3 + SWUCB). We start with laying out a fundamental result (i.e., Lemma 4) required for the analysis of Lemma 5, which directly leads to the main result.

To begin, we look at the total reward drift $D_t(\tau, a)$ contributed by some arm a until some time t with a sliding window of size τ . As mentioned in the analysis of Theorem 1, l_t is the Lipschitz constant corresponding to the drift function at any given time t (Refer Assumption 1 in the main paper), and $N_t(a)$ as the total pulls (i.e., essentially $N_t(\tau, a)$ for $\tau = t$) for arm a until time t . We define $l = \max_t l_t$ for the following Lemma.

Lemma 4 $D_t(\tau, a) \leq l N_t(a) \sqrt{\xi \log(\min(t, \tau))}$, for all arms $a \in [1, K]$ and $t \in [1, T]$.

Proof: The reasoning is similar to the analysis of Lemma 2. There is a drift in the reward from the agent only when the principal provides some compensation. The inequalities mentioned below, quantifying this event, must hold true simultaneously. Since the agent is assumed to be greedy in nature, the arm g_t (i.e., the arm with the highest empirical average until time t) seems more appealing than a_t (i.e., the principal's recommended arm at time t). However, the principal recommends arm a_t per its algorithm, balancing exploration and exploitation. The inequalities below represent the same.

$$\bar{X}_t(\tau, g_t) \geq \bar{X}_t(\tau, a_t) \quad (40)$$

$$\bar{X}_t(\tau, g_t) + c_t(\tau, g_t) \leq \bar{X}_t(\tau, a_t) + c_t(\tau, a_t) \quad (41)$$

The two inequalities imply $\bar{X}_t(\tau, g_t) - \bar{X}_t(\tau, a_t) \leq c_t(\tau, a_t) - c_t(\tau, g_t)$. Since the drift function is Lipschitz smooth (Assumption 1 in the main paper), we can upper bound the instantaneous drift $\delta_t(a_t)$ for the pulled arm a_t at time t as below,

$$\begin{aligned} \delta_t(a_t) &\leq l_t (c_t(\tau, a_t) - c_t(\tau, g_t)) \\ &\leq l_t \sqrt{\frac{\xi \log(\min(t, \tau))}{N_t(\tau, a_t)}} \end{aligned} \quad (42)$$

Therefore, summing the instantaneous drift $\delta_t(a)$ until some time t , for some arm a gives us its total drift contribution until said t , as shown below.

$$\begin{aligned} D_t(\tau, a) &= \sum_{s=1}^t \mathbf{1}(a_s = a) \delta_s(a) \\ &= \sum_{s=1}^t \mathbf{1}(a_s = a) l_s \sqrt{\frac{\xi \log(\min(t, \tau))}{N_s(\tau, a)}} \\ &\leq l \sqrt{\xi \log(\min(t, \tau))} \sum_{s'=1}^{N_t(a)} \frac{1}{\sqrt{N_{s'}(\tau, a)}} \\ &\leq N_t(a) l \sqrt{\xi \log(\min(t, \tau))} \quad (\text{Since } N_{s'}(\tau, a) \geq 1) \end{aligned} \quad (43)$$

As discussed in section 2-A, since the expected reward of an arm is in the range $[0, 1]$, we get $(\mu^* - \mu(a)) \leq 1$ for all arms $a \in [1, K]$. Therefore, bounding the expected regret after T pulls essentially amounts to controlling the expected number of times a sub-optimal arm is pulled. In the next lemma, we bound that expected number of total pulls $N_T(a)$ of some sub-optimal arm $a \neq a_t^*$ for the entire horizon $t \leq T$.

Lemma 5 *Let $\xi > 1/2$. For any integer $\tau > 0$ and any arm $a \in [1, K]$,*

$$\mathbb{E}[N_T(a)] \leq C(\tau) \frac{T \log(\tau)}{\tau} + \tau \beta_T + \log^2(\tau)$$

where

$$C(\tau) = \frac{(l\sqrt{\tau} + 1)^2 \xi [T/\tau]}{(\Delta_{\mu_T}(i))^2 T/\tau} + \frac{2}{\log(\tau)} \left\lceil \frac{\log(\tau)}{\log(1 + 4\sqrt{1 - (2\xi)^{-1}})} \right\rceil$$

Proof: As in the analysis of Lemma 2, let us denote by $\Delta_{\mu_T}(a_t)$ the minimum of the difference of the expected reward $\mu_t(a_t^*)$ of the best arm at time t and the expected reward $\mu_t(a_t)$ of the arm a_t for all $t \in [1, T]$ when a_t is not the optimal arm, i.e.,

$$\Delta_{\mu_T}(a_t) = \min_{t \in [1, T]; a_t \neq a_t^*} (\mu_t(a_t^*) - \mu_t(a_t)). \quad (44)$$

Let us consider some arm $a \in [1, K]$. The total pulls $N_T(a)$ can be represented as the following.

$$N_T(a) = 1 + \sum_{t=K+1}^T \mathbf{1}(a_t = a \neq a_t^*) \quad (45)$$

Which can be broken into two disjoint events, as mentioned below.

$$N_T(a) = 1 + \sum_{t=K+1}^T \mathbf{1}(a_t = a \neq a_t^*; N_t(\tau, a) < A(\tau)) + \sum_{t=1}^T \mathbf{1}(a_t = a \neq a_t^*; N_t(\tau, a) \geq A(\tau)) \quad (46)$$

where

$$A(\tau) = \frac{(l\sqrt{\tau} + 1)^2 \xi \log(\tau)}{(\Delta_{\mu_T}(a))^2} \quad (47)$$

The next few steps follow from [33]'s analysis. For the same definitions of $\mathcal{T}(\tau)$, and we can bound $N_T(a)$ by:

$$N_T(a) \leq 1 + [T/\tau]A(\tau) + \tau \beta_T + \sum_{t \in \mathcal{T}(\tau)} \mathbf{1}(a_t = a \neq a_t^*; N_t(\tau, a) \geq A(\tau)) \quad (48)$$

Now, for $t \in \mathcal{T}(\tau)$ the event $E : \{a_t = a \neq a_t^*; N_t(\tau, a) \geq A(\tau)\}$ will occur when the following inequality, labelled as \mathcal{Z} hold,

$$\mathcal{Z} : \bar{X}_t(\tau, a) + c_t(\tau, a) > \bar{X}_t(\tau, a_t^*) + c_t(\tau, a_t^*) \quad (49)$$

Let us denote by $\bar{Y}_t(\tau, a)$ the empirical average of arm a excluding the drift at each time until t (i.e., pure empirical average, different from $\bar{X}_t(\tau, a)$) for the sliding window size of τ . We have the following inequalities from \mathcal{Z} ;

$$\begin{aligned} \mathcal{Z} : \bar{Y}_t(\tau, a) + \frac{D_t(\tau, a)}{N_t(a)} + c_t(\tau, a) \\ > \bar{Y}_t(\tau, a_t^*) + \frac{D_t(\tau, a_t^*)}{N_t(a_t^*)} + c_t(\tau, a_t^*) \end{aligned} \quad (50)$$

Therefore, the upper bound for the difference between the expected reward of the optimal arm a_t^* and the current arm a is

$$\begin{aligned} \bar{Y}_t(\tau, a_t^*) - \bar{Y}_t(\tau, a) &< \frac{D_t(\tau, a)}{N_t(a)} + c_t(\tau, a) \\ &= \left(l\sqrt{N_t(\tau, a)} + 1 \right) \sqrt{\frac{\xi \log(\min(t, \tau))}{N_t(\tau, a)}} \\ &\leq (l\sqrt{\tau} + 1) \sqrt{\frac{\xi \log(\min(t, \tau))}{N_t(\tau, a)}} \end{aligned} \quad (51)$$

Now, we can decompose the event E as the following: for the event \mathcal{Z} to happen, at least one of the events E_t^i has to occur.

$$\{a_t = a \neq a_t^*; N_t(\tau, a) \geq A(\tau)\} \subseteq E_t^1 \cup E_t^2 \cup E_t^3 \quad (52)$$

Where

$$E_t^1 = \{\bar{X}_t(\tau, a) > \mu_t(\tau, a) + c_t(\tau, a)\} \quad (53)$$

$$E_t^2 = \{\bar{X}_t(\tau, a_t^*) < \mu_t(\tau, a_t^*) - c_t(\tau, a_t^*)\} \quad (54)$$

and

$$\begin{aligned} E_t^3 &= \{\bar{Y}_{\tau, t}(a_t^*) - \bar{Y}_t(\tau, a) \\ &< (l\sqrt{\tau} + 1) \sqrt{\frac{\xi \log(\min(t, \tau))}{N_t(\tau, a)}}\} \end{aligned} \quad (55)$$

The event E_t^1 represents the situation when the algorithm overestimates the average reward of arm a , E_t^2 when the algorithm underestimates the average reward of the best arm a_t^* , and E_t^3 is when the expected rewards for both the arms a and a_t^* have a small difference.

Note that for the choice of $A(\tau)$, the event E_t^3 never occurs, as shown in the following inequality;

$$\begin{aligned} (l\sqrt{\tau} + 1) \sqrt{\frac{\xi \log(\min(t, \tau))}{N_t(\tau, i)}} \\ \leq (l\sqrt{\tau} + 1) \sqrt{\frac{\xi \log(\tau)}{A(\tau)}} = \frac{\Delta_{\mu_T}(i)}{2} \end{aligned} \quad (56)$$

Therefore, we have the probability of the event E_t^3 as $\mathbb{P}[E_t^3] = 0$ and we get the corresponding probabilities for the E_t^1 and E_t^2 from [33] as

$$\mathbb{P}[E_t^1] = \mathbb{P}[E_t^2] \leq \left\lceil \frac{\log(\min(t, \tau))}{\log(1 + \eta)} \right\rceil \min(t, \tau)^{-2\xi(1 - \frac{\eta^2}{16})} \quad (57)$$

We finally have the bound, by taking $\xi > 1/2$ and $\eta = 4\sqrt{1 - (1/2\xi)}$, so as to make $2\xi(1 - \eta^2/16) = 1$:

$$\mathbb{E}[N_T(a)] \leq \underbrace{1 + \lceil T/\tau \rceil A(\tau) + \tau\beta_T}_M + 2 \underbrace{\sum_{t=1}^T \left\lceil \frac{\log(\min(t, \tau))}{\log(1 + \eta)} \right\rceil}_{N} \frac{1}{\min(t, \tau)} \quad (58)$$

Substituting $A(\tau)$ in M and expanding N we upper bound $\mathbb{E}[N_T(a)]$ by

$$\mathbb{E}[N_T(a)] \leq 1 + \lceil T/\tau \rceil \frac{(l\sqrt{\tau} + 1)^2 \xi \log(\tau)}{(\Delta_{\mu_T(i)})^2} + \frac{2T}{\tau} \left\lceil \frac{\log(\tau)}{\log(1 + \eta)} \right\rceil + \tau\beta_T + \log^2(\tau) \quad (59)$$

■

With the proper choice of τ , which minimizes the expression, we get the statement of Theorem 2.

C. Material for Theorem 5

Recall that in a continuously changing non-stationary environment, the mean rewards of the arms can change an arbitrary number of times but have a variation budget, limiting the total change throughout the horizon.

In line 4 of Algorithm 4, in the main paper, the principal employs some stochastic bandit algorithm (such as UCB1 or Thompson sampling) to recommend an arm to the agent. The standard instance-dependent regret results for some of the standard bandit algorithms have the form $O(\log T/\Delta_a)$ ([29]), where Δ_a is the difference in the expected reward of some arm a and the best arm a^* (also referred as the gap of arm a). For almost all intent and purposes, it is assumed that the minimum possible value of Δ_a (i.e., the denominator in the regret bound) is large enough to prevent the regret from being an arbitrarily large quantity and, therefore, make the regret bound meaningful. Assuming a well-behaved instance is reasonable for a stationary bandit model, however, in the continuously-changing non-stationary environment's case, it becomes a little too strong. Therefore, we need to give the rewards more freedom to vary but simultaneously impose certain assumptions to make them mathematically tractable.

We start by breaking the time horizon T into a sequence of batches of timestamps, where each batch (except possibly the last one) is of a fixed size σ . If we have total $m = \lceil T/\sigma \rceil$ batches, we define below the ‘‘gap’’ of a certain arm a in the current context.

Definition 1 *The minimum average difference $\Delta_j(a)$ between the mean rewards of some arm a in comparison with*

the instantaneous best arm a_t^ for each time step, within a single batch j , is defined as:*

$$\Delta_j(a) = \min_{j \in [1, m]} \frac{1}{\sigma} \sum_{t \in \mathcal{T}_j} (\mu_t^* - \mu_t(a))$$

where \mathcal{T}_j is the set of timestamps within the batch j . To be more precise, $\mathcal{T}_k = \{(k-1)\sigma + 1, 2, \dots, \min(T, k\sigma)\}$.

In the next couple of assumptions, we form a balance between providing the rewards more freedom to vary and simultaneously imposing certain limitations to make them mathematically tractable. Assumption 2 lower-bounds the minimum average difference between the mean rewards of the best overall arm and any other arm within a single batch, and assumption 3 restricts the number of times the difference between the expected rewards between any pair of different arms can go below a certain threshold.

Assumption 2 *There exists a constant $M \in (0, 1)$ such that $\Delta_j(a) \geq M$ for any arm $a \in [1, K]$ and all $j \in [1, m]$.*

Assumption 3 *For any given time epoch t , there exists $\alpha \in (0, 1)$ and some threshold value $\varepsilon \in [0, 1]$ such that the following is true*

$$\sum_{t=1}^T \sum_{a=1}^K \sum_{b=1; b \neq a}^K \mathbf{1}(\mu_t(a) - \mu_t(b) \leq \varepsilon) \leq t^\alpha \quad (60)$$

Given the variation of the rewards (and potentially gaps of different arms), we would henceforth consider the regret bound, which does not depend on the said gap of an arm, but only the T and K . For instance, UCB1 has a worst-case regret bound (see chapter 1 in [28]) of $O(\sqrt{KT \log T})$ for K arms. Next, we present the lemma that bounds the regret of algorithm 4 in the main paper, which directly leads to the main Theorem 5.

Lemma 6 *If the principal in Algorithm 4 (line 4) employs some stochastic bandit algorithm under reward drift with the worst-case regret of $\lambda\sqrt{TK \log T}$ for some constant $\lambda > 0$, and the batch size $\sigma = (\lambda T/V_T)^{2/3} (K \log T)^{1/3}$, Then, for $T \geq 2, K \geq 2$ and $V_T \in [1/K, T/K]$, the total regret is*

$$R_V^T \leq 2\lambda^{1/3} \cdot V_T^{1/3} (K \log T)^{1/3} T^{2/3} \quad (61)$$

Proof: We use the following proof structure. We break the horizon into sequences of batches of size σ each and then analyze the performance gap between the *single* best action and the arm returned by a dynamic oracle (i.e., returns the best action for each time t) in each batch. Then, we plug in the known performance of the principal's bandit algorithm relative to the single best action under reward drift. We sum them over the batches to establish the required regret bound.

For the time horizon $T \geq 1$ and a total of $K \geq 2$ arms, we have the variation budget $V_T \in [1/K, T/K]$. We break the horizon T into sequence of $m = \lceil T/\sigma \rceil$ batches denoted by \mathcal{T}_j for all $j \in [1, m]$, of size σ , except possibly the last one. We decompose the regret in some batch j , as the expected sum of differences between the arm with the

maximum expected reward μ_t^* at the time t (i.e., dynamic oracle result) and the expected reward μ_t of an arm chosen by the algorithm at t , with $p = \mathbb{E} \left[\max_{a \in [1, K]} \left\{ \sum_{t \in \mathcal{T}_j} X_t(a) \right\} \right]$ as:

$$\mathbb{E} \left[\sum_{t \in \mathcal{T}_j} (\mu_t^* - \mu_t) \right] = \underbrace{\sum_{t \in \mathcal{T}_j} (\mu_t^* - p)}_{J_1} + p - \underbrace{\mathbb{E} \left[\sum_{t \in \mathcal{T}_j} \mu_t \right]}_{J_2} \quad (62)$$

The first component J_1 represents the expected loss of the algorithm with respect to using a single action over batch j . The second component J_2 is the expected regret relative to the single best action in batch j . From (6) in [34], we know that $J_1 \leq 2V_j\sigma$. Since J_2 is, after all, the regret of a policy with respect to a single best action within a batch, we can plug in the worst-case performance of the bandit algorithm under reward drift. For some constant $\lambda > 0$, we have the following.

$$J_2 = \mathbb{E} \left[\max_{a \in [1, K]} \left\{ \sum_{t \in \mathcal{T}_j} X_t(a) \right\} - \mathbb{E} \left[\sum_{t \in \mathcal{T}_j} \mu_t^\pi \right] \right] \leq \lambda \sqrt{\sigma K \log(\sigma)} \quad (63)$$

The next step is to sum them over the entire horizon. Since there are m batches, the overall regret is over all the permissible reward sequences \mathcal{V} :

$$\begin{aligned} R_{\mathcal{V}}^T &= \sup_{\mu \in \mathcal{V}} \left\{ \sum_{t=1}^T \mu_t^* - \mathbb{E} \left[\sum_{t=1}^T \mu_t \right] \right\} \\ &\leq \sum_{j=1}^m \left(\lambda \sqrt{\sigma K \log(\sigma)} + 2V_j\sigma \right) \\ &\leq \left(\frac{2T}{\sigma} \right) \cdot \lambda \sqrt{\sigma K \log(\sigma)} + 2V_T\sigma \quad (\text{since } m \geq 1) \\ &= 2T \cdot \lambda \sqrt{\frac{K \log(\sigma)}{\sigma}} + 2V_T\sigma \end{aligned} \quad (64)$$

Selecting $\sigma = (\lambda T / V_T)^{2/3} (K \log(T))^{1/3}$, we get

$$R^\pi(\mathcal{V}, T) \leq 2\lambda^{1/3} \cdot V_T^{1/3} (K \log(T))^{1/3} T^{2/3} \quad (65)$$

Therefore, to prove Theorem 5, by virtue of Lemma 6, we need to show that UCB1, ε -greedy and Thompson Sampling algorithms follow the required regret bound (for a batch) of $O(\sqrt{\sigma K \log \sigma})$ under reward drift, for a batch size σ , and thus employable by the principal (i.e., in line 4 of Algorithm 4).

We know from [21] that the regret for UCB1, R^{UCB1} under reward drift, for a batch j , for some constant C , is

$$R^{UCB1} \leq \sum_{a \in [1, K]; \Delta_j(a) > 0} \frac{C(l+1)^2}{\Delta_j(i)} \log(\sigma) \quad (66)$$

As discussed before, the gap $\Delta_j(a)$ for any arm a can get arbitrarily small in the denominator, therefore a more general regret bound should be derived (see remark 1.13 in [28]), independent of the $\Delta_j(a)$ term. Let $\varepsilon \in (0, 1)$, then

- The regret contributed by all the arms with $\Delta_j(a) > \varepsilon$ is at most $\frac{CK(l+1)^2 \log(\sigma)}{\varepsilon}$.
- The regret contributed by all the arms with $\Delta_j(a) \leq \varepsilon$ is at most $\varepsilon \cdot \sigma^\alpha$ (from Assumption 2).

Therefore, the total regret for this batch j is at most:

$$R^{UCB1} \leq \varepsilon \cdot \sigma^\alpha + \frac{CK(l+1)^2 \log(\sigma)}{\varepsilon} \quad (67)$$

For $\varepsilon = \sqrt{\frac{CK(l+1)^2 \log(\sigma)}{\sigma}}$ and $\alpha = 1$, we have $R^{UCB1} \leq \sqrt{C} (l+1) \sqrt{\sigma K \log(\sigma)}$. This result is almost optimal as it almost matches the minimax lower bound of $O(\sqrt{\sigma})$ for any stochastic bandit algorithm ([28]).

Next, from [21] we know that the regret $R^\varepsilon G$ for ε -greedy under reward drift, for a batch j , for some constant C ,

$$R^\varepsilon G \leq \frac{l \cdot C}{M^2} \sum_{a \in [1, K]; \Delta_j(a) > 0} \frac{\log(\sigma)}{(\Delta_j(a))^2} \quad (68)$$

We employ the same technique as above to fix some $\varepsilon \in (0, 1)$.

- The regret contributed by all the arms with $\Delta_j(a) > \varepsilon$ is at most $\frac{KlC \log(\sigma)}{M^2 \varepsilon^2}$.
- The regret contributed by all the arms with $\Delta_j(a) \leq \varepsilon$ is at most $\varepsilon \cdot \sigma^\alpha$.

Therefore, the total regret for this batch j is at most:

$$R^\varepsilon G \leq \varepsilon \cdot \sigma^\alpha + \frac{KlC \log(\sigma)}{M^2 \varepsilon^2} \quad (69)$$

For $\varepsilon = \left(\frac{KlC \log(\sigma)}{M^2 \sigma} \right)^{1/3}$ and $\alpha = 3/4$, we have $R^\varepsilon G \leq (Cl/M^2)^{1/3} (K \log(\sigma))^{1/3} \sqrt{\sigma} \leq (Cl/M^2)^{1/3} \sqrt{3\sigma K \log(\sigma)}$ for $\sigma \geq 1$. The analysis is the same as above for Thompson Sampling.

In summary, we showed that the principal could use UCB1, ε -greedy, and Thompson sampling as a submodule (line 4 in Algorithm 4) to recommend arms to the agent, balancing the exploration-exploitation dilemma and achieving a sublinear regret.

D. Proof for Theorem 3

Note that the principal compensates an agent when the following inequalities are true.

The principal provides compensation when,

$$\bar{X}_t(\gamma, a) > \bar{X}_t(\gamma, a_t^*) \quad (70)$$

and

$$\bar{X}_t(\gamma, a) + \sqrt{\frac{\xi \log(n_t(\gamma))}{N_t(\gamma, a)}} < \bar{X}_t(\gamma, a_t^*) + \sqrt{\frac{\xi \log(n_t(\gamma))}{N_t(\gamma, a_t^*)}} \quad (71)$$

Thus, the compensation is provided to the agent even when the agent pulls the optimal arm and $N_t(\gamma, a_t^*) <$

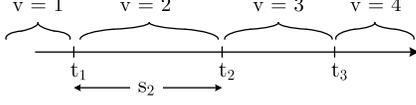


Fig. 6: Time intervals representing sets V and S for compensation analysis.

$N_t(\gamma, a)$. Therefore, the average number of times a player is compensated for pulling the best arm is upper bounded by $M = \max_{a \neq a_t^*} \mathbb{E}[N_T(a)]$. Since the rewards can change with a breakpoint, so can the best arm. We must consider each interval between breakpoints separately.

Let t_b be the timestamp at which the b th breakpoint occurs, $\forall b \in [1, \beta_T]$. Let $V = \{(t_{b-1}, t_b) : \forall b \in [1, \beta_T]\}$ and $S = \{s_v : s_v = |t_{b-1} - t_b| \forall v \in V; \forall b \in [1, \beta_T]\}$ (see Figure 6) be the set of intervals between breakpoints and the sizes of each interval, respectively. For convenience, consider $t_0 = 0$.

Let $v^+ \in V$ be the interval in which the principal pays the compensation a maximum number of times M_{v^+} when the player plays the interval's best arm i_v^* . Let $\hat{N}_{s_v}(a)$ be the number of times the suboptimal arm a was played in the interval v . Therefore,

$$v^+ = \arg \max_{v \in V} \left[\max_{a \in [1, K]: a \neq a_v^*} \mathbb{E} \left[\hat{N}_{s_v}(a) \right] \right] \quad (72)$$

Taking a^* to be the best arm in the interval v^+ , We get,

$$\begin{aligned} M_{v^+} &\leq \beta_T \max_{a \neq a^*} \mathbb{E} \left[\hat{N}_{s_v}(a) \right] \\ &\leq \beta_T \max_{a \neq a^*} \mathbb{E} \left[\hat{N}_T(a) \right] \end{aligned} \quad (73)$$

We get the total expected compensation as

$$\begin{aligned} \mathbb{E}[C_T] &\leq \sum_{i=1}^K \frac{\mathbb{E}[\hat{N}_T(a)]}{\sum_{j=1}^K N_j(\gamma, i)} \sqrt{\frac{\xi \log(n_t(\gamma))}{N_j(\gamma, i)}} \\ &\leq \sum_{j=1}^{M_{v^+}} \sqrt{\frac{\xi \log(n_t(\gamma))}{N_j(\gamma, i)}} + \sum_{i=1}^{K-1} \sum_{j=1}^K \sqrt{\frac{\xi \log(n_t(\gamma))}{N_j(\gamma, i)}} \\ &\leq \sqrt{\xi \log(n_t(\gamma))} \left(\sum_{j=1}^{\beta_T \max_{i \neq i^*} \mathbb{E}[\hat{N}_T(i)]} 1 + \sum_{i=1}^{K-1} \sum_{j=1}^K 1 \right) \\ &\leq \sqrt{\xi \log(n_t(\gamma))} \left(\sum_{i=1}^K \sum_{j=1}^{\beta_T \mathbb{E}[\hat{N}_T(i)]} 1 + \sum_{i=1}^{K-1} \sum_{j=1}^K 1 \right) \\ &\leq \sqrt{\xi \log(n_t(\gamma))} \left(\sum_{i=1}^K \sum_{j=1}^{(\beta_T+1)\mathbb{E}[\hat{N}_T(i)]} 1 \right) \\ &\leq \sqrt{\xi \log(n_t(\gamma))} \left(\sum_{i=1}^K (\beta_T + 1) \mathbb{E}[\hat{N}_T(i)] \right) \end{aligned} \quad (74)$$