

Ballooning Multi-Armed Bandits

by

Ganesh Ghalme, Swapnil Dhamal, Shweta Jain, Sujit Prakash Gujar, Narahari Yadati

in

International Conference on Autonomous Agents and Multi-agent Systems, AAMAS 2020

Report No: IIIT/TR/2020/-1



Centre for Visual Information Technology
International Institute of Information Technology
Hyderabad - 500 032, INDIA
January 2020

Ballooning Multi-Armed Bandits

Ganesh Ghalme^{*} Swapnil Dhamal[†] Shweta Jain[‡]
Sujit Gujar[§] Y. Narahari[¶]

January 29, 2020

Abstract

In this paper, we introduce *ballooning multi-armed bandits* (BL-MAB), a novel extension to the classical stochastic MAB model. In BL-MAB model, the set of available arms grows (or balloons) over time. In contrast to the classical MAB setting where the regret is computed with respect to the best arm overall, the regret in a BL-MAB setting is computed with respect to the best available arm at each time. We first observe that the existing MAB algorithms are not regret-optimal for the BL-MAB model. We show that if the best arm is equally likely to arrive at any time, a sub-linear regret cannot be achieved, irrespective of the arrival of other arms. We further show that if the best arm is more likely to arrive in the early rounds, one can achieve sub-linear regret. Our proposed algorithm determines (1) the fraction of the time horizon for which the newly arriving arms should be explored and (2) the sequence of arm pulls in the exploitation phase from among the explored arms. Making reasonable assumptions on the arrival distribution of the best arm in terms of the thinness of the distribution's tail, we prove that the proposed algorithm achieves sub-linear instance-independent regret. We further quantify explicit dependence of regret on the arrival distribution parameters. We reinforce our theoretical findings with extensive simulation results.

1 Introduction

The classical stochastic multi-armed bandit (MAB) model provides an elegant abstraction to a number of important sequential decision making problems. In this setting, the planner chooses (or pulls) a single arm in each discrete time instant from a fixed pool of finitely many arms for a finite number of time instants. Typically it is assumed that the number of arms is much smaller than the number of time instances. Each arm, when pulled, generates a reward from a fixed but a-priori unknown stochastic distribution corresponding to the pulled arm. However, the arms that are not pulled do not generate any reward to the planner. The planner's goal is to minimize the regret, i.e., the loss

^{*}Indian Institute of Science, India. ganeshg@iisc.ac.in

[†]Chalmers University of Technology, Sweden. swapnil.dhamal@gmail.com

[‡]Indian Institute of Technology, Ropar, India. shwetajains20@gmail.com

[§]International Institute of Information Technology. sujit.gujar@iiit.ac.in

[¶]Indian Institute of Science. narahari@iisc.ac.in

incurred in expected cumulative reward due to not knowing the reward distribution of the arms beforehand. The MAB problem encapsulates the classical exploration versus exploitation dilemma, in that the planner’s algorithm has to arrive at an optimal trade-off between exploration (pulling relatively unexplored arms) and exploitation (pulling the best arms according to the history of pulls thus far). This problem has been extensively studied in the literature. These studies include analyzing the lower bound on regret [1], design and analysis of asymptotically optimal algorithms [2; 3; 4; 5], empirical studies [6; 7; 8], and several extensions [9; 10]. We provide a detailed review of the relevant literature in Section 7.

The theoretical results in MAB are complemented by a wide variety of modern applications which can be seamlessly modelled in the MAB setup. Internet advertising [11; 12], crowdsourcing [13], clinical trials [14], wireless communication [15] represent a few of the many applications. Due to its wide range applications and an elegant theoretical foundation, many variants of the MAB problem have been proposed. In this paper, we propose a novel variant which we call Ballooning Multi-Armed Bandits (BL-MAB). In contrast to the classical MAB where the set of available arms is fixed throughout the run of an algorithm, the set of arms in BL-MAB grows (or balloons) over time. As the number of arms increases (potentially linearly) with time, it is clear that an optimal algorithm has to ignore (or drop) a few arms. Hence, in addition to achieving an optimal trade-off between the number of exploratory pulls and exploitation pulls, the algorithm must also ensure that it does not drop too many (or too few) arms.

To see that the traditional algorithms are not regret-optimal in the BL-MAB setting, consider the following thought experiment. Let a new arm arrive at each time instant in decreasing order of mean reward, and let the MAB algorithm run for a total of T time instants. The traditional MAB algorithms (such as UCB1, MOSS etc.) would pull the newly arrived arm at each time and thus would incur a regret of $O(T)$. Note here that the best arm appeared at the first time instant itself, however, as the set of arms is monotonically expanding over time, the algorithm could not sufficiently explore every arm. Observe that the regret in BL-MAB depends not only on the mean reward of the arms, but also on when they arrive. Hence, any BL-MAB algorithm ought to be aware of the arrival of the arms.

1.1 Motivation

We motivate the practical significance of BL-MAB with a few applications. In general, BL-MAB is directly applicable in any scenario where the set of options grows over time, and, the objective is to choose the best option available at any given time.

A contemporary example is provided by question and answer (Q&A) platforms such as Reddit, Stack Overflow, Quora, Yahoo! Answers, and ResearchGate, where the platform’s goal is to discover the highest quality answer that should be displayed in the most prominent slot, for a given question. Each answer post is modeled as a distinct arm of a BL-MAB instance, and the rewards are distributed according to a Bernoulli distribution parameterized by the quality of the posted answer. Note that the quality parameter is a priori unknown to the platform and hence needs to be learnt. For this, the platform employs certain endorsement mechanisms with indicators such as upvotes, likes, and shares (or re-posts). A user endorses the answer that is displayed to her, if she likes it. Each display of a posted answer corresponds to a pull of the corresponding arm. At each time instant, a new user observes the existing answer posts shown by the platform, decides whether to endorse them, and may also choose to post her own answer, thus increasing the number of available arms. Hence, the

number of available arms (answers) monotonically increases over time.

The problem of learning qualities of the answers on Q&A forums has been modeled under the MAB framework in various studies [16; 17; 18]. However, these studies resort to the existing MAB variations which are not well suited for Q&A forums. For instance, Ghosh and Hummel [16] model the problem with a classical MAB framework by limiting the number of arms via strategic choice of an agent, by assuming that a user incurs a certain cost for posting an answer and hence posts it only if she derives a positive utility by doing so. However, a user’s behavior on the platform may be driven by simple cognitive heuristics rather than a well calibrated strategic decision [19]. In another work [18], the number of arms is limited by randomly dropping some of the arms from consideration. The regret is then computed with respect to only the considered arms. That is, they do not account for the regret incurred due to the randomly dropped arms.

Some of the other applications of BL-MAB framework are in various websites that feature user reviews, such as Amazon and Flipkart (product reviews), Tripadvisor (hotel reviews), IMDB (movie reviews). As time progresses, the reviews for a product (or a hotel or a movie) keep arriving, and the website aims to display the most useful reviews for that product (or hotel or movie) at the top. The usefulness of a review is estimated using users’ endorsements for that review, similar to that in Q&A forums. BL-MAB is also applicable in scenarios where users comment on a video or news article on a video or news hosting website, where the website’s objective is to display the most popular or interesting comment on the top. The BL-MAB setting thus provides a natural framework to be considered in such type of application.

BL-MAB needs an independent investigation owing to a number of reasons. For instance, one of the MAB variants that holds some similarity with BL-MAB is sleeping multi-armed bandit (S-MAB) [20; 21], where a subset of a fixed set of base arms is available at each time instant. Though the S-MAB framework captures the availability of a small subset of arms at each time, it assumes that the set of base arms is fixed and is small as compared to the time horizon. In contrast, the BL-MAB framework allows for the number of available arms to increase, potentially linearly with time. Hence, an optimal sleeping bandits algorithm such as AUER[20] would end up giving a linear regret in BL-MAB setting.

Another MAB variant which is somewhat similar is the many-armed (potentially infinite) bandit [22; 23; 24], where the number of arms could be potentially equal to or greater than the time horizon. Berry et al. [24] consider the case of an infinite arm bandit with Bernoulli reward distribution. However, they consider that the optimal arm has a quality of 1, which is seldom the case in practical applications. Other investigations considering infinitely many arms [22; 23] make certain assumptions on the distribution of the near optimal arm to achieve sub-linear regret. Further, all the above works consider that all the arms are available in all time instants, and hence use the traditional notion of regret. In our case, the regret incurred by an algorithm in a given time instant is the difference between the quality of the best available arm during that time and the quality of the arm pulled by the algorithm (same as the notion of regret considered in sleeping bandits). The BL-MAB framework is thus an interesting blend of both the sleeping bandit model and the many-armed bandit model.

1.2 Our Contributions

- We introduce the BL-MAB model that allows the set of arms to grow over time.
- We show for the BL-MAB model that the instance independent regret will grow linearly with time, in

the absence of any distributional assumption on the arrival time of the highest quality arm. Specifically, we show that if the best arm is equally likely to arrive at any time, the sublinear expected regret cannot be achieved (Theorem 1).

- We propose an algorithm (BL-MOSS) which determines: (1) the fraction of the time horizon until which the newly arriving arms should be explored at least once and (2) the sequence of arm pulls during the exploitation phase. Our key finding is that BL-MOSS achieves sub-linear regret under practical and minimal assumptions on the arrival distribution of the best arm, namely, sub-exponential tail (Theorem 3) and sub-Pareto tail (Theorem 4). Note that we make no assumption on the arrival of the other arms. As the regret depends on the qualities of the arms and the sequence of their arrivals, it is interesting that with sub-exponential and sub-Pareto assumption on only the best arm’s arrival pattern, we can achieve sub-linear regret.
- We carry out a pertinent simulation study to empirically observe how the expected regret varies with the time horizon. We find a strong validation for our theoretically derived regret bounds.

The paper is organized as follows. In Section 2, we present the BL-MAB model. In Section 3, we first show that if the best arm arrives uniformly at random, one cannot achieve sub-linear regret. We hence define two distributional assumptions on the arrival time of the best arm which enables us to achieve sub-linear regret. Next, we present some preliminaries in Section 4, followed by our proposed algorithm and its theoretical analysis in Section 5.2. Section 6 presents our simulation results. We conclude the paper with related work (Section 7) and future directions (Section 8).

2 The Model

A classical MAB instance is given by the tuple $\langle K, (\mathcal{D}_i)_{i \in K} \rangle$. Here, K is a fixed set of arms and \mathcal{D}_i is the reward distribution corresponding to an arm i . Further, it is assumed that \mathcal{D}_i are mutually independent distributions. Denote by q_i , the mean of distribution \mathcal{D}_i . Consider that each of the distributions \mathcal{D}_i is supported over a finite interval and is unknown to the algorithm. Throughout the paper, without loss of generality, we consider that \mathcal{D}_i is supported over $[0, 1]$. Further, we will refer to q_i as the quality of the arm i . A MAB algorithm is run in discrete time instants, and the total number of time instants is denoted by time horizon T . In each time instant aka round, the algorithm selects a single arm and observes the reward corresponding to the selected arm. The arms which are not selected, do not give any reward. More precisely, a MAB algorithm is a mapping from the history of arm pulls and obtained rewards, to the set of arms.

At each time instant, a BL-MAB algorithm chooses a single arm from the set of available arms and receives a reward generated randomly according to the reward distribution \mathcal{D}_i of the chosen arm i . New arms may spring up at each time instant. Throughout the paper, we consider that at most one new arm arrives at each time, and the arms are never dropped. Let $K(t)$ denote the set of arms available at round t . In the BL-MAB model, this set of available arms grows by at most one arm per round, i.e., $K(t) \subseteq K(t+1)$ and $|K(t)| \leq |K(t+1)| \leq |K(t)| + 1$. A BL-MAB instance, therefore, is given by $\langle T, (K(t), (\mathcal{D}_i)_{i \in K(t)})_{t=1}^T \rangle$.

Similar to the notion of regret in the sleeping stochastic MAB model, we introduce the notion of regret in BL-MAB setting that takes into account the availability of the arms at each time t . Let i_t denote

the arm pulled by the algorithm and i_t^* be the best available arm at time t , i.e., $i_t^* = \arg \max_{i \in K(t)} q_i$. Further, let \mathcal{I} denote a BL-MAB instance and A be a BL-MAB algorithm. The instance-dependent regret of A is given by

$$\mathcal{R}_A(T, \mathcal{I}) = \mathbb{E} \left[\sum_{t=1}^T (q_{i_t^*} - q_{i_t}) \right].$$

Throughout the paper, we consider instance-independent regret given as $\mathcal{R}_A(T) = \sup_{\mathcal{I}} \mathcal{R}_A(T, \mathcal{I})$. Note that the instance-independent regret bound is a worst case regret bound over all the arrival sequences of the arms and all possible reward distributions. In the next section, we show, for the BL-MAB setting, that it is not possible to achieve sublinear instance-independent regret bound.

3 Lower Bound on Regret

As pointed out in Section 1, it is clear that UCB-style algorithms (which pull arms based on uncertainty) would pull each incoming arm at least once, leaving no rounds for exploitation. Hence, they incur linear regret in the ballooning bandit setup¹. However, it is not obvious that a different, more sophisticated algorithm (such as the one which randomly drops some arms) may not be able to achieve sub-linear regret. Our first result (Theorem 1) shows that no algorithm can attain sub-linear regret under a general BL-MAB setting.

Consider the following BL-MAB instance \mathcal{I} . Let there be a unique best arm i^* with quality $q_{i^*} = 1/2 + \varepsilon$ and all other arms $j \neq i^*$ have quality $q_j = 1/2$. A new arm arrives at each time, i.e., $|K(t)| = t$. Further, let the arrival of the best arm be uniformly distributed over time, i.e., $\mathbb{P}(t = i^*) = 1/T$ for all $t = 1, 2, \dots, T$. Let i_t^* denote the optimal arm till time t . Further, let G be the set of arms pulled by the algorithm, i.e., $G = \{i : N_{i,T} \geq 1\}$. We will show that for any fixed $G \subset T$, the expected regret is lower bounded by $\Omega(T)$. We first observe that an algorithm that pulls $|G|$ number of arms achieves the minimum regret when it pulls the first G arms.

Claim 1. The minimum regret is achieved when algorithm considers the first G arms.

The proof of Claim 1 is given in Appendix A.

Theorem 1. For the BL-MAB instance \mathcal{I} , the expected regret of any algorithm A is lower bounded by $\Omega(T)$.

Proof. The expected regret of a BL-MAB algorithm is given by

$$\begin{aligned} \mathcal{R}_A(T) &\geq \mathcal{R}_A(T, \mathcal{I}) = \mathbb{E}_A \left[\sum_{t=1}^T (q_{i_t^*} - q_{i_t}) \right] \\ &\geq P(i^* \in G) \sum_{i \in G} \mathbb{E}[N_{i,T}] \Delta(i^*, i) + \sum_{i \notin G} P(i^* = i) (T - i) \Delta(i^*, i) \end{aligned}$$

In the above expression, the first term represents the internal regret of the learning algorithm and the second quantity is the external regret. Here, $\Delta(i^*, i) = q_{i^*} - q_i = \varepsilon$ if $i^* < i$ and 0 otherwise. From Claim 1, we have that an algorithm will incur least regret if it pulls first $|G|$ arms. Further, from the

¹In particular, when $K(t) = t$.

classical result in [1], in order to separate the quality of the arms, we should have $\mathbb{E}[N_{i,T}] \geq \eta \cdot \log(T)$ for some positive problem dependent constant η , for all $i \in G$ and $i \neq i^*$. Hence, we have

$$\begin{aligned} \mathcal{R}_A(T) &\geq P(i^* \in G) \sum_{i=1}^{|G|-1} \eta \log(T) \varepsilon + \sum_{i=|G|+1}^T P(i = i^*)(T - i) \varepsilon \\ &= \left[\frac{\eta |G| (|G| - 1) \log(T)}{T} + \frac{(T - |G| - 1)(T - |G|)}{2T} \right] \cdot \varepsilon \\ &= \left[(1 + 2\eta \log(T)) |G|^2 - (2(T - \eta \log(T)) - 1) |G| + T^2 - T \right] \cdot \frac{\varepsilon}{2T} \end{aligned}$$

Note that the above expression is quadratic in $|G|$. For $T \leq 1/2 + \eta \log(T)$, the minimum occurs when the value of $|G|$ is the least (in the positive domain), which is 1. In this case we have

$$\begin{aligned} \mathcal{R}_A(T) &\geq \left[(1 + 2\eta \log(T)) - (2(T - \eta \log(T)) - 1) + T^2 - T \right] \cdot \frac{\varepsilon}{2T} \\ &= \Omega(T) \end{aligned}$$

For the case where $T > 1/2 + \eta \log(T)$, the minimum occurs when $|G| = \frac{2(T - \eta \log(T)) - 1}{2(1 + 2\eta \log(T))}$. Hence,

$$\begin{aligned} \mathcal{R}_A(T) &\geq \left[\frac{(2(T - \eta \log(T)) - 1)^2}{4(1 + 2\eta \log(T))} - \frac{(2(T - \eta \log(T)) - 1)^2}{2(1 + 2\eta \log(T))} + T^2 - T \right] \cdot \frac{\varepsilon}{2T} \\ &= \left[T^2 - T - \frac{(T - \eta \log(T) - 1/2)^2}{(1 + 2\eta \log(T))} \right] \cdot \frac{\varepsilon}{2T} \\ &> \left[\frac{(T - 1/2)}{2} \frac{2\eta \log(T)}{1 + 2\eta \log(T)} - 1/4 \right] \cdot \varepsilon = \Omega(T) \end{aligned}$$

Theorem 1 provides a strong impossibility result on the achievable instance-independent regret bound under BL-MAB setting. However, one can still achieve sub-linear regret by imposing appropriate structure on the BL-MAB instances. Observe that the regret depends on the arrival of arms, i.e., $(K(t))_{t=1}^T$, and their reward distributions $(\mathcal{D}_i)_{i \in K(t)}$. We impose restrictions on the arrival of the best arm $i^* = \arg \max_{i \in K(T)} q_i$ so that the probability that i^* arrives early is large enough; this would allow a learning algorithm to explore the best arm enough to estimate the true quality of that arm with high probability. As noted previously, the other arms may arrive arbitrarily. Further note that we make no assumption on the qualities of individual arms.

3.1 Arrival of the Best Arm

Let X be the random variable denoting the time at which the best arm arrives. Further, let $F_X(t)$ denote the cumulative distribution function of X . In our first result, we use the following Sub-exponential tail assumption on the arrival time of the best arm.

Sub-exponential tail. There exists a constant $\lambda > 0$ such that the probability of the best arm arriving later than t rounds, is upper bounded by $e^{-\lambda t}$, i.e., $F_X(t) > 1 - e^{-\lambda t}$.

Next, we consider a relaxed condition on the tail probabilities, i.e., when the tail does not shrink as fast as in the sub-exponential case. We consider the family of distributions whose tail is thicker than that of sub-exponential arrival distribution.

Sub-Pareto tail. There exists a constant $\beta > 0$ such that the probability of the best arm arriving later than t rounds, is upper bounded by $t^{-\beta}$, i.e., $F(t) > 1 - t^{-\beta}$.

The aforementioned assumptions naturally arise in the context of Q&A forums as observed in extensive empirical studies on the nature of answering as well as voting behavior of the users. Anderson et al. [25] observe that high reputation users hasten to post their answers early. One possible explanation for this phenomenon could be that the users who are motivated by the visibility that their answers receive, tend to be more active on the platform and also provide high quality answers early on, which explains their reputation score. Similarly, the most useful product reviews are likely to arrive early after the product release. Thus, it is reasonable to assume that the best arm arrives, with high probability, in early rounds.

Note that the uniform distribution is the limiting case of the sub-exponential case, when $\lambda = 0$. We show that, while the uniform distribution results in linear regret (Theorem 1), a sub-linear regret can be achieved for BL-MAB instances having the best arm arrival distribution with even slightly thinner tail than that of uniform distribution.

4 Preliminaries

We now present some essentials which will be useful for our analysis in the remainder of the paper.

4.1 Lambert W Function

Definition 1. For any $x > -e^{-1}$, the Lambert W function, $W(x)$, is defined as the solution to the equation $we^w = x$, i.e., $W(x)e^{W(x)} = x$.

Lambert W function satisfies the following properties [26] (proofs are provided in Appendix for completeness):

Property 1. The Lambert W function can be equivalently written as the inverse of the function $f(x) := xe^x$, i.e., $W(xe^x) = x$.

Property 2. For any $x \in [0, \infty)$, the Lambert W function is unique, non-negative, and strictly increasing.

Property 3. For any $x \geq e$, we have $\log(x)/2 < W(x) \leq \log(x)$.

It can be noted that it is easy to numerically approximate $W(x)$ for a given x , using Newton-Raphson's or Halley's method. Moreover, there exist efficient numerical methods for evaluating it to arbitrary precision [27].

4.2 The MOSS Algorithm

We use MOSS (Minimax Optimal Strategy in the Stochastic case) [28] as a black box learning algorithm. For a fixed number of k arms, the MOSS algorithm pulls an arm i_t at time t such that

$$i_t \in \arg \max_{i \in K} \left[\hat{q}_{i, N_{i,t}} + \sqrt{\frac{\max(\log(\frac{T}{k \cdot N_{i,t}}), 0)}{N_{i,t}}} \right].$$

Algorithm 1: BL-MOSS

Input: Time horizon T , Distributional parameter λ or β

1

$$\text{Set } \alpha := \begin{cases} \frac{W(2\lambda T)}{2\lambda T} & \text{under sub-exponential tail property} \\ T^{\frac{-2\beta}{1+2\beta}} & \text{under sub-Pareto tail property} \end{cases}$$

for $t = 1, 2, \dots, T$ **do** **Input:** A newly arriving arm at time t 2 **if** $|K(t)| \leq \lceil \alpha T \rceil$ **then**3 | MOSS ($K(t)$)4 **else**5 | MOSS ($K(\lceil \alpha T \rceil)$)6 **end**7 **end**

Here, $K = \{1, 2, \dots, k\}$ denotes the set of arms and $N_{i,t}$ is the number of times arm i was pulled before (and excluding) round t and $\hat{q}_{i,N_{i,t}}$ are the empirical estimates of the arm i from $N_{i,t}$ samples. Each arm is pulled once in the beginning, and ties are broken arbitrarily. The following result gives an upper bound on the expected regret of MOSS which is optimal up to a constant factor (it achieves the lower bound on regret given by [29]). Throughout the paper, we use the notation MOSS (K) to denote that the MOSS algorithm is run with set of arms K .

Theorem 2. [Audibert and Bubeck [28]] For any time horizon $T \geq 1$, the expected regret of Moss is given by $\mathcal{R}_{\text{Moss}}(T) \leq 6\sqrt{kT}$.

5 The B-MOSS Algorithm and Regret Analysis

5.1 The BL-MOSS Algorithm

We now present our algorithm, BL-MOSS (Algorithm 1), that uses MOSS as a black-box. The number of arms explored by BL-MOSS is dependent on the distribution of arrival of the best arm. In particular, BL-MOSS considers only the first $\lceil \alpha T \rceil$ arms in its execution ($\alpha \in (0, 1]$). Later in this section, we show how to derive the value of α for distributions with sub-exponential and sub-Pareto tails. Observe that the proposed BL-MOSS is a simple extension of MOSS and this algorithm is practically easy to implement.

5.2 Regret Analysis of BL-MOSS

For a given BL-MAB instance \mathcal{I} , let $j^* = \arg \max_{i \in K(\lceil \alpha T \rceil)} q_i$ and $i^* = \arg \max_{i \in K(T)} q_i$. Clearly, we have that $q_{i^*} \geq q_{j^*}$. As stated earlier, the regret of the algorithm can be decomposed into internal regret, i.e., the regret incurred by the learning algorithm considering only $\lceil \alpha T \rceil$ arms and external regret, i.e., the regret incurred by BL-MOSS due to the fact that BL-MOSS might have ignored the best arm. Write $\Delta(i, j) = q_i - q_j$ and let t_i be the time of arrival of arm i . Further, let i_t^* denote the best arm till time t . The instance-dependent regret $\mathcal{R}_{\text{BL-MOSS}}(T, \mathcal{I})$ is given as

$$\begin{aligned}
& \mathbb{P}(i^* = j^*) \left[\underbrace{\sum_{t=1}^{t_{j^*}^* - 1} \Delta(i_t^*, i_t) + \sum_{t=t_{j^*}^*}^T \Delta(j^*, i_t)}_{\mathcal{R}_{\text{BL-MOSS}}^{\text{int}}(T)} \right] \\
& + \mathbb{P}(i^* \neq j^*) \left[\underbrace{\sum_{t=1}^{t_i^* - 1} \Delta(i_t^*, i_t) + \sum_{t=t_i^*}^T \Delta(i^*, i_t)}_{\mathcal{R}_{\text{BL-MOSS}}^{\text{ext}}(T)} \right]
\end{aligned}$$

The first and the second terms respectively denote the internal regret and the external regret of BL-MOSS. We ignore the ceiling in $\lceil \alpha T \rceil$ throughout this section to avoid notation clutter.

Note that $\mathcal{R}_{\text{MOSS}(L)}(T) \leq \mathcal{R}_{\text{MOSS}(K)}(T)$ for all $L \subset K$. This is true for any time horizon T . From Theorem 2, we have the following observation about the internal regret of BL-MOSS.

Observation 1. For the value of α computed by BL-MOSS, we have $\mathcal{R}_{\text{BL-MOSS}}^{\text{int}}(T) \leq \mathcal{R}_{\text{MOSS}(\alpha T)}(T) \leq 6\sqrt{\alpha}T$.

In order to bound the overall regret, we begin with the following lemma which explicitly shows the relation between the expected regret of the algorithm and $F_X(\cdot)$. Recall that the random variable X denotes the time of arrival of the best arm.

Lemma 1. *The upper bound on the expected regret for any BL-MAB instance is given by $\mathcal{R}_{\text{BL-MOSS}}(T) \leq T(1 - (1 - 6 \cdot \sqrt{\alpha})F_X(\alpha T))$, with BL-MOSS exploring only the first αT arrived arms.*

Proof. For a given BL-MAB instance \mathcal{I} , let t_i denote the time at which arm i becomes available for the first time. Let i^* denote the best arm till T rounds, i.e., $i^* = \arg \max_{i \in K(T)} q_i$. Further, let j^* be the best arm among the arms considered by BL-MOSS, i.e., $j^* = \arg \max_{j \in K(\alpha T)} q_j$. Notice that $K(\alpha T) \subseteq K(T)$. This implies $q_{i^*} \geq q_{j^*}$.

$$\begin{aligned}
\mathcal{R}_{\text{BL-MOSS}}(T, \mathcal{I}) & \leq \mathbb{E} \left[\sum_{t=1}^{\alpha T} (q_{j^*} - q_{i_t}) + \sum_{t=\alpha T+1}^T (q_{i^*} - q_{i_t}) \right] && (\because q_{i^*} > q_{j^*}) \\
& = \mathbb{P}(i^* = j^*) \left[\sum_{t=1}^T (q_{j^*} - q_{i_t}) \right] \\
& + \mathbb{P}(i^* \neq j^*) \left[\sum_{t=1}^{\alpha T} (q_{j^*} - q_{i_t}) + \sum_{t=\alpha T+1}^T (q_{i^*} - q_{i_t}) \right] \\
& \leq 6\mathbb{P}(i^* = j^*)\sqrt{\alpha T \cdot T} + \sum_{t=1}^T (q_{i^*} - q_{i_t})\mathbb{P}(i^* \neq j^*) \\
& \hspace{10em} \text{(From Observation 1 and since } q_{i^*} \geq q_{j^*} \text{)} \\
& \leq 6T\sqrt{\alpha} \cdot \mathbb{P}(i^* = j^*) + \mathbb{P}(i^* \neq j^*)T && (\because \sum_{t=1}^T (q_{i^*} - q_{i_t}) \leq T) \\
& = 6T\sqrt{\alpha} \cdot \mathbb{P}(t_{i^*} \leq \alpha T) + (1 - \mathbb{P}(t_{i^*} \leq \alpha \cdot T))T \\
& = T(1 - (1 - 6 \cdot \sqrt{\alpha})\mathbb{P}(t_{i^*} \leq \alpha T))
\end{aligned}$$

$$= T(1 - (1 - 6 \cdot \sqrt{\alpha})F_X(\alpha T))$$

Note that the above inequality holds for any BL-MAB instance and hence we have $\mathcal{R}_{\text{BL-MOSS}}(T) = \sup_{\mathcal{I}} \mathcal{R}_{\text{BL-MOSS}}(T, \mathcal{I}) \leq T(1 - (1 - 6 \cdot \sqrt{\alpha})F_X(\alpha T))$ \square

Next, we show that under the sub-exponential tail property on X , BL-MOSS achieves sub-linear regret. We begin with the following lemma that lower bounds the probability of the arrival of the best quality arm in the initial αT rounds.

Lemma 2. *Let the arm arrival distribution of the best arm satisfy sub-exponential tail property for some $\lambda \geq 0$. Then for any $c > 0$ and $\alpha \geq \frac{W(\lambda T/c)}{\lambda T/c}$, we have that $F_X(\alpha T) > (1 - \alpha^c)$.*

Proof. Note that $\alpha \geq \frac{W(\lambda T/c)}{\lambda T/c} \implies \frac{\alpha \lambda T}{c} \geq W(\lambda T/c)$. Hence, by Property **P1** we have $W\left(\frac{\alpha \lambda T}{c} \cdot e^{\alpha \lambda T/c}\right) \geq W(\lambda T/c)$.

By Property **P2** we have $\frac{\alpha \lambda T}{c} \cdot e^{\alpha \lambda T/c} \geq \lambda T/c$. Equivalently, $\alpha \geq e^{-\alpha \lambda T/c}$. So, we have $1 - \alpha^c \leq 1 - e^{-\lambda(\alpha T)} < F_X(\alpha T)$. The last inequality follows from the sub-exponential tail property. \square

Theorem 3. *Let the arrival distribution of the best arm satisfy the sub-exponential tail property for some $\lambda \geq 0$, and let T be large enough such that $T > \frac{36c \log(36)}{\lambda}$ for some $c > 0$. Then with $\alpha = \frac{W(\lambda T/c)}{\lambda T/c}$, the upper bound on the expected regret of BL-MOSS, $\mathcal{R}_{\text{BL-MOSS}}(T)$, is $O\left(T \cdot \max\left(e^{-cW(\lambda T/c)}, e^{-\frac{W(\lambda T/c)}{2}}\right)\right)$. The upper bound on the expected regret is minimized when $c = 1/2$ and is given by $O\left(\sqrt{\frac{T \log(2\lambda T)}{2\lambda}}\right)$.*

Proof. From Lemma 2, we have $F_X(\alpha T) > 1 - \alpha^c$ for all $\alpha \geq \frac{W(\lambda T/c)}{\lambda T/c}$. Thus, from Lemma 1, we have $\mathcal{R}_{\text{BL-MOSS}}(T) < T(1 - (1 - 6 \cdot \sqrt{\alpha})(1 - \alpha^c))$.

Note that for achieving sub-linear regret, it is necessary that $(1 - 6 \cdot \sqrt{\alpha})$ is strictly positive, for which it is necessary that $\alpha < 1/36$. From Lemma 2, we also have $\alpha \geq \frac{W(\lambda T/c)}{\lambda T/c}$. Since such a feasible α may not exist for small values of T , we consider that T is large enough. It can be easily shown that $\frac{W(\lambda T/c)}{\lambda T/c} < 1/36 \iff T > \frac{36c \log(36)}{\lambda} \approx \frac{129c}{\lambda}$ (see Claim 2 in Appendix).

Thus, for $1/36 > \alpha \geq \frac{W(\lambda T/c)}{\lambda T/c}$, we have: $\mathcal{R}_{\text{BL-MOSS}}(T) < T(6 \cdot \sqrt{\alpha} + \alpha^c - 6 \cdot \alpha^{c+1/2})$. Recall that by definition, we have $\alpha \leq 1$. Thus when $c \in (0, 1/2]$, the term α^c dominates the other terms in the regret expression, whereas when $c > 1/2$, the term $\sqrt{\alpha}$ dominates. We analyze these cases separately.

Case 1 ($c \in (0, 1/2]$): In this case, the regret is given by $\mathcal{R}_{\text{BL-MOSS}}(T) = O(\alpha^c T)$. Note that the regret is minimized for the lowest feasible value of α , i.e., $\alpha = \frac{W(\lambda T/c)}{\lambda T/c}$, resulting in $\mathcal{R}_{\text{BL-MOSS}}(T) = O\left(T\left(\frac{W(\lambda T/c)}{\lambda T/c}\right)^c\right) = O\left(T \cdot e^{-cW(\lambda T/c)}\right)$. The last equality follows from the equivalent definition of Lambert W function (Property **P1**).

Case 2 ($c \in [1/2, \infty)$): In this case, the regret is given by $\mathcal{R}_{\text{BL-MOSS}}(T) = O(\sqrt{\alpha} T)$. Again, the regret is minimized when $\alpha = \frac{W(\lambda T/c)}{\lambda T/c}$. The regret in this case is given by $\mathcal{R}_{\text{BL-MOSS}}(T) = O\left(T \cdot \sqrt{\frac{W(\lambda T/c)}{\lambda T/c}}\right) = O\left(T \cdot e^{-\frac{W(\lambda T/c)}{2}}\right)$.

Further, we have that in Case 1, $e^{-cW(\lambda T/c)} > e^{-\frac{W(2\lambda T)}{2}}$ for any $c \in (0, 1/2)$ (see Claim 3 in Appendix). For Case 2, we have from Property P2 that, $W(\lambda T/c)$ is decreasing in c , which gives us that $e^{-\frac{W(2\lambda T)}{2}} < e^{-\frac{W(\lambda T/c)}{2}}$ for any $c \in (1/2, \infty)$. This shows that the minimum regret is achieved when $c = 1/2$, and the regret is given by $\mathcal{R}_{\text{BL-MOSS}}(T) = O\left(\sqrt{\frac{T \cdot W(2\lambda T)}{2\lambda}}\right) = O\left(\sqrt{\frac{T \log(2\lambda T)}{2\lambda}}\right)$. The last inequality follows from Property P3, since $2\lambda T \geq e \cdot T > \frac{36c \log(36)}{\lambda}$ where $c = 1/2$. \square

We now prove the sub-linear regret of BL-MOSS under the sub-Pareto tail property.

Lemma 3. *Let the arm arrival distribution of the best arm satisfy sub-Pareto tail property for some $\beta > 0$. Then for any $c > 0$ and $\alpha \geq T^{\frac{-\beta}{c+\beta}}$, we have that $F_X(\alpha T) > (1 - \alpha^c)$.*

Proof. First note that $\alpha \geq T^{\frac{-\beta}{c+\beta}} \iff \alpha^c \geq (\alpha T)^{-\beta}$. This implies that $(1 - \alpha^c) \leq 1 - (\alpha T)^{-\beta}$. Further, from the sub-Pareto tail property, we have that $1 - (\alpha T)^{-\beta} < F_X(\alpha T)$. \square

Theorem 4. *Let the arrival distribution of arms satisfy the sub-Pareto tail property for some $\beta > 0$, and let T be large enough such that $T > (36)^{\frac{c+\beta}{\beta}}$ for some $c > 0$. Then with $\alpha = T^{\frac{-\beta}{\beta+c}}$, the upper bound on the expected regret of BL-MOSS, $\mathcal{R}_{\text{BL-MOSS}}(T)$, is $O(\max(T^{\frac{c+\beta(1-c)}{c+\beta}}, T^{\frac{2c+\beta}{2(c+\beta)}}))$. The upper bound on the expected regret is minimized when $c = 1/2$ and is given by $O(T^{\frac{1+\beta}{1+2\beta}})$.*

Proof. From Lemmas 1 and 3, we have $\mathcal{R}_{\text{BL-MOSS}}(T) < T(1 - (1 - 6 \cdot \sqrt{\alpha})(1 - \alpha^c))$. For achieving sub-linear regret, it is necessary that $(1 - 6 \cdot \sqrt{\alpha})$ is strictly positive. So, we should have $\alpha < 1/36$. Further, from Lemma 3, we have $\alpha \geq T^{\frac{-\beta}{c+\beta}}$. So, for a feasible α to exist, it is necessary that $T^{\frac{-\beta}{c+\beta}} < 1/36 \iff T > (36)^{\frac{c+\beta}{\beta}}$, i.e., T is large enough. Thus, for $1/36 > \alpha \geq T^{\frac{-\beta}{c+\beta}}$, we have $\mathcal{R}_{\text{BL-MOSS}}(T) < T(6 \cdot \sqrt{\alpha} + \alpha^c - 6 \cdot \alpha^{c+1/2})$. As earlier, we analyze two cases.

Case 1 ($c \in (0, 1/2]$): In this case, the regret is given by $\mathcal{R}_{\text{BL-MOSS}}(T) = O(\alpha^c T)$. The minimum regret is obtained when $\alpha = T^{\frac{-\beta}{\beta+c}}$ and is given by $O(T^{1-\frac{c\beta}{c+\beta}})$.

Case 2 ($c \in [1/2, \infty)$): In this case, the regret is given by $\mathcal{R}_{\text{BL-MOSS}}(T) = O(\sqrt{\alpha} T)$. Again, the regret is minimum when $\alpha = T^{\frac{-\beta}{\beta+c}}$ and is given by $O(T^{\frac{2c+\beta}{2(c+\beta)}})$.

Furthermore, it is easy to see that in Case 1, $T^{\frac{1+\beta}{1+2\beta}} > T^{\frac{\beta+c(1-\beta)}{c+\beta}}$ for any $c \in (0, 1/2)$. Similarly, in Case 2, $T^{\frac{1+\beta}{1+2\beta}} > T^{\frac{2c+\beta}{2(c+\beta)}}$ for any $c \in (1/2, \infty)$. This shows that the minimum regret is achieved when $c = 1/2$. \square

5.3 Important Observations

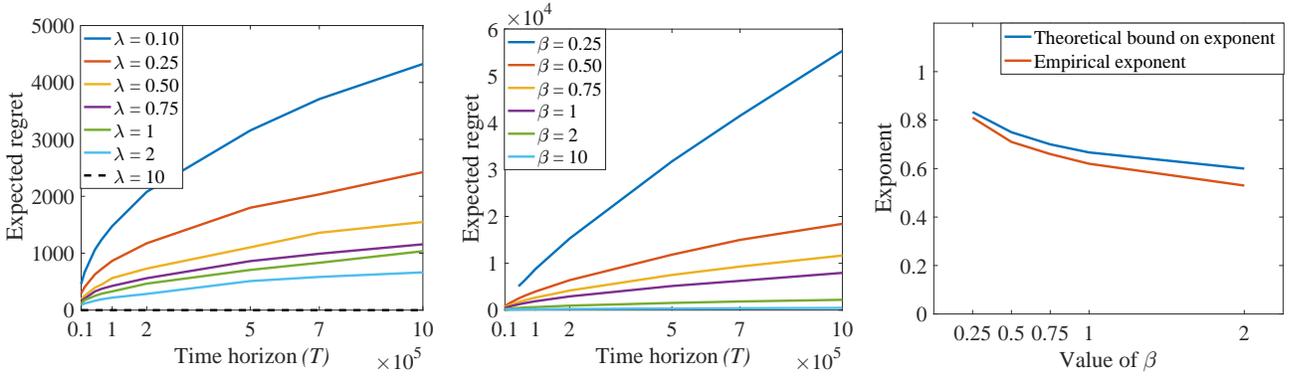
We conclude the section with some key observations.

- In the sub-exponential tail case, as $\lambda \rightarrow \infty$, we have $\frac{W(2\lambda T)}{2\lambda T} \rightarrow 0$. This implies that the upper bound on the expected regret goes to 0. Note that in this case, $\lceil \alpha T \rceil = 1$. Since BL-MOSS considers a single arm, the internal regret is zero. Further, we have $F_X(1) \rightarrow 1$, i.e., the first arm is optimal with probability approaching 1, the external regret is also zero. As $\lambda \rightarrow 0$, the tail bounds are trivial and

are satisfied by uniform distribution. From Theorem 1, we have that the regret in this case cannot be sub-linear.

- In the sub-Pareto tail case, following the similar argument as in the sub-exponential tail case, we have that as $\beta \rightarrow 0$, the regret $\mathcal{R}_{\text{BL-MOSS}}(T) \rightarrow O(T)$. On the other hand, as $\beta \rightarrow \infty$, the regret goes to $O(\sqrt{T})$. The larger value of β implies that the probability that the optimal arm arrives by $t = 2$ is close to 1; then we have that the regret of BL-MOSS is asymptotically optimal. Further, it asymptotically achieves the information theoretic lower bound (which is $O(\sqrt{T})$).
- One could also consider UCB1 instead of MOSS. While UCB1 is any-time algorithm, MOSS needs the time horizon as an input. However, the important distinction between the two algorithms is that the instance-independent regret of UCB1, which is $O(\sqrt{kT \log(T)})$, is greater than that of MOSS; hence we use MOSS in BL-MOSS. One can similarly use UCB1 to get any-time version of BL-MOSS with slightly more (up to $\sqrt{\log(T)}$) regret guarantee.

6 Simulation Study



(a) Expected regret as a function of time horizon for various sub-exponential tails

(b) Expected regret as a function of time horizon for various sub-Pareto tails

(c) Empirical exponents vs. theoretical bounds for time horizons up to 10^6

Figure 1

Simulation results

So far, we focused on deriving upper bounds on regret for distributions (on the arrival time of the best arm) having sub-exponential and sub-Pareto tail with different values of λ and β , respectively. In particular, for the case of sub-Pareto tail, we deduced that the extent of sublinearity of the regret (the exponent of T in the order of the regret) depends on the value of β . On the other hand, the upper bound on regret for the case of sub-exponential tail had the same order with respect to T for any reasonable value of λ , albeit with different multiplicative and additive terms for different values of λ . In this section, we aim to illustrate how the expected regret varies with the time horizon T , and how the empirical exponents compare with their theoretical bounds for different values of β and λ , for time horizons up to 10^6 rounds.

6.1 Simulation Setup

Note that in a traditional MAB setup, a simulation for a larger time horizon T'' could be conducted as an extension of a simulation for a smaller time horizon $T' < T''$. In other words, after obtaining the results for time horizon T' , the results for time horizon T'' can be obtained by running simulations for an additional $T'' - T'$ rounds. However, in the BL-MAB setup where new arms continue arriving with time and the desired time horizon is known, we have seen that the optimal value of α and hence $\lceil \alpha T \rceil$ depend on the time horizon. Owing to different values of $\lceil \alpha T \rceil$ for different time horizons T , the simulation for a time horizon T' are not extendable to time horizon $T'' > T'$. So even if we have simulation results for time horizon T' , it is necessary to run a fresh set of simulations for obtaining results for time horizon $T'' > T'$. In our simulation study, we consider the following values of time horizon: $\{1, 2, 5, 7\} \times 10^4$, $\{1, 2, 5, 7\} \times 10^5$, 10^6 .

We consider that a new arm arrives in each round, and the probability of an arm arriving at time t being the best arm is determined by the distribution function $F_X(t)$. Thereafter, this best arm (i^*) is assigned a quality (q_{i^*}) between 0 and 1 uniformly at random, and the rest of the arms are assigned quality parameters between 0 and q_{i^*} uniformly at random. Given a time horizon T , the value of α and hence $\lceil \alpha T \rceil$ are obtained based on our theoretical analysis. The arm to be pulled in a round is determined by Algorithm 1, wherein the pulled arm generates unit reward with probability equal to its quality, and no reward otherwise (i.e., as per Bernoulli distribution). The regret in each round is computed as the difference between the quality of the best arm available in that round and the quality of the pulled arm. The overall regret is the sum of the regrets over all rounds from 1 till T . Note that we are concerned with the regret irrespective of the numerical values of the arms' qualities. So, for a given instance of the arrival of the best arm, we consider the worst-case regret over 50 sub-instances, where the quality parameters assigned to the arms in different sub-instances are independent of each other. Also, since different instances would have the best arm arriving in different rounds, the expected regret is obtained by simulating over 1000 such random instances and averaging over the corresponding worst-case regret values.

Our primary objective is to observe how the expected regret varies with the time horizon T . In order to observe the influence of various sub-exponential and sub-Pareto tail distributions over the arrival time of the best arm, we conduct simulations for different values of parameters λ and β : $\{0.10, 0.25, 0.50, 0.75, 1, 2, 10\}$. The other objective is to determine the empirical exponent of the plots (i.e., the value of γ such that the expected regret is approximately a constant multiple of T^γ). To achieve this, we first estimate the constant factor ξ by dividing the expected regret for $T = 10^6$ by T^γ , for a given value of γ . We then compute the squared error when attempting to fit the expected regret with ξT^γ . Considering candidate values of γ to be between 0 and 1 with intervals of 0.01, we deduce the empirical exponent to be the value of γ which results in the least squared error. We also consider another method for determining the empirical exponent: we produce the line of best fit for the scatter plot of $\log(T)$ versus the log of the expected regret for that T ; the slope of this line gives the empirical exponent. The empirical exponents obtained using the two methods are almost identical (differing by less than 0.01).

6.2 Simulation Results

As mentioned at the end of our theoretical analysis, for the sub-exponential tail case when $\lambda \rightarrow \infty$, the upper bound on the expected regret goes to 0. In our simulations with the maximum observed

time horizon of 10^6 , the expected regret was observed to be uniformly zero, even for $\lambda = 10$ (see Figure 1(a)). Further, for other considered values of λ , the plots exhibit a prominent sub-linear nature. In particular, considering the maximum observed time horizon of 10^6 , the empirical exponents for different values of λ were consistently observed to be between 0.45 and 0.5 (Theorem 3 showed the order of the regret with respect to T , for reasonable values of λ , to be bounded by $\sqrt{T \log(T)}$, which is an exponent close to 0.5).

For the sub-Pareto tail case illustrated in Figure 1(b), note that we have no result for $\beta = 0.10$ because the value of T for obtaining a feasible α should be greater than 36^6 , which is beyond our maximum observed time horizon of 10^6 . Moreover, we have partial results for $\beta = 0.25$ because the value of T for obtaining a feasible α should be greater than 36^3 ; so the plot starts with $T = 0.5 \times 10^5$. It can be seen, in general, that the plots in Figure 1(b) follow a far less sub-linear nature and exhibit a much higher expected regret than those in Figure 1(a). This is intuitive from our analysis that the sub-exponential tail case is likely to result in a much lower regret than the sub-Pareto tail case. In particular, the empirical exponent for $\beta = 0.25$ was deduced to be 0.8, which is close to linear (its theoretical upper bound as per our analysis is 0.83). In general, considering the maximum observed time horizon of 10^6 , it can be seen from Figure 1(c) that the upper bound on the theoretical exponent (which is $\frac{1+\beta}{1+2\beta}$ from Theorem 4) and the empirical exponent are close to each other.

Note that the gap between the empirical exponents and the corresponding theoretical upper bounds could be attributed to the fact that it is difficult to find the worst-case distribution over the reward parameters of the arms. Hence, it is unlikely that the worst-case (or instance-independent) expected regret could be attained in the simulations with a random reward structure. Since the gap is not very significant, the simulation results suggest that the bounds derived in our regret analysis of BL-MOSS (in Section 5.2) are, in all probability, tight.

7 Additional Related Work

A standard stochastic MAB framework considers that the number of available arms is fixed (say k) and typically much less than the time horizon (say T). In the seminal work of Lai and Robbins [1], the authors showed that any MAB algorithm in such a setting must incur a regret of $\Omega(\frac{\log T}{D_{\text{KL}}})$ where D_{KL} is the Kullback-Leibler divergence between the best arm and the second best arm. Auer [30] proposed the UCB1 algorithm which attains a matching upper bound on the expected regret. However, the instance-independent (i.e., in adversarial case) regret of the variant of UCB1, (α, ψ) -UCB, is given by $O(\sqrt{kT \log T})$ [10]. The MOSS algorithm proposed by Audibert and Bubeck [28] achieves the instance-independent regret of $O(\sqrt{kT})$. Bubeck and Cesa-Bianchi [10] present a detailed survey on regret bounds of these algorithms.

The problem of learning qualities of the answers on Q&A forums was first modeled under MAB framework by Ghosh and Hummel [16] where generation of a new arm was considered as a consequence of strategic choice of an agent. Though this model captures strategic aspects of the contributors, there is an important practical issue with such modelling. Each agent, being a strategic attention seeker, is assumed to produce the effort just enough to satisfy incentive compatibility in the equilibrium. We do not assume an efforts-and-costs model and show that, even when the number of answers grows linearly with time if the qualities of arriving answers follow certain mild distributional assumption, the proposed algorithm achieves sub-linear regret.

Tang and Ho [17] consider a model with fixed number of arms but with a platform where agents provide biased feedback. On such Q&A forums, it is more relevant to consider the problem with increasing number of arms. A recent work by Liu and Ho [18] limits the growth of the bandit arms by randomly dropping some of the arms from consideration, and computing the regret with respect to only the considered arms. That is, they do not account for the regret incurred due to the randomly dropped arms.

8 Discussion and Future Work

In this paper, we presented a novel extension to the classical MAB model, which we call the Ballooning bandits model (BL-MAB). We showed that, in the absence of any distributional assumption on the arrival of the best quality arm, it is impossible to achieve a sub-linear regret. We proposed an algorithm for the BL-MAB model and provided sufficient conditions under which the proposed algorithm achieves sub-linear regret. In particular, when the arrival distribution of the best quality arm has a sub-exponential or sub-Pareto tail, our algorithm BL-MOSS achieves sub-linear regret by restricting the number of arms to be explored in an intelligent way. Our results indicate that, the number of arms to be explored depends on the distributional parameters, namely, λ (for sub-exponential case) and β (for sub-Pareto case), which must be known to the algorithm. It will be interesting to see how a learning algorithm can be designed to learn these parameters as well. As noted earlier, we only consider a structure on the arrival of the best arm. One could also consider a more sophisticated arrival process of the arms, for obtaining tighter regret guarantees.

References

- [1] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985. (Not cited.)
- [2] Peter Auer and Ronald Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010. (Not cited.)
- [3] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pages 1–39, 2012. (Not cited.)
- [4] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933. (Not cited.)
- [5] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002. (Not cited.)
- [6] Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pages 2249–2257, 2011. (Not cited.)
- [7] R Devanand and P Kumar. Empirical study of thompson sampling: Tuning the posterior parameters. In *AIP Conference Proceedings*, volume 1853, 2017. (Not cited.)
- [8] Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, Zheng Wen, et al. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018. (Not cited.)
- [9] Aleksandrs Slivkins. Introduction to multi-armed bandits. *arXiv preprint arXiv:1904.07272*, 2019. (Not cited.)
- [10] Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012. (Not cited.)
- [11] Moshe Babaioff, Yogeshwer Sharma, and Aleksandrs Slivkins. Characterizing truthful multi-armed bandit mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 79–88. ACM, 2009. (Not cited.)
- [12] Alessandro Nuara, Francesco Trovo, Nicola Gatti, and Marcello Restelli. A combinatorial-bandit algorithm for the online joint bid/budget optimization of pay-per-click advertising campaigns. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. (Not cited.)
- [13] Shweta Jain, Sujit Gujar, Satyanath Bhat, Onno Zoeter, and Y. Narahari. A quality assuring, cost optimal multi-armed bandit mechanism for expertsourcing. *Artificial Intelligence*, 254:44 – 63, 2018. ISSN 0004-3702. (Not cited.)
- [14] Sofía S Villar, Jack Bowden, and James Wason. Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 30(2):199, 2015. (Not cited.)

- [15] Setareh Maghsudi and Sławomir Stańczak. Joint channel selection and power control in infrastructureless wireless networks: A multiplayer multiarmed bandit framework. *IEEE Transactions on Vehicular Technology*, 64(10):4565–4578, 2014. (Not cited.)
- [16] Arpita Ghosh and Patrick Hummel. Learning and incentives in user-generated content: Multi-armed bandits with endogenous arms. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 233–246, 2013. (Not cited.)
- [17] Wei Tang and Chien-Ju Ho. Bandit learning with biased human feedback. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1324–1332, 2019. (Not cited.)
- [18] Yang Liu and Chien-Ju Ho. Incentivizing high quality user contributions: New arm generation in bandit learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. (Not cited.)
- [19] Keith Burghardt, Emanuel Alsina, Michelle Girvan, William Rand, and Kristina Lerman. The myopia of crowds: A study of collective evaluation on stock exchange. *Robert H. Smith School Research Paper No. RHS*, 2736568, 2016. (Not cited.)
- [20] Robert Kleinberg, Alexandru Niculescu-Mizil, and Yogeshwer Sharma. Regret bounds for sleeping experts and bandits. *Machine learning*, 80(2-3):245–272, 2010. (Not cited.)
- [21] Aritra Chatterjee, Ganesh Ghalme, Shweta Jain, Rohit Vaish, and Y Narahari. Analysis of thompson sampling for stochastic sleeping bandits. In *Uncertainty in Artificial Intelligence, UAI 2017*, 2017. (Not cited.)
- [22] Yizao Wang, Jean-Yves Audibert, and Rémi Munos. Algorithms for infinitely many-armed bandits. In *Advances in Neural Information Processing Systems*, pages 1729–1736, 2009. (Not cited.)
- [23] Alexandra Carpentier and Michal Valko. Simple regret for infinitely many armed bandits. In *International Conference on Machine Learning*, pages 1133–1141, 2015. (Not cited.)
- [24] Donald A Berry, Robert W Chen, Alan Zame, David C Heath, Larry A Shepp, et al. Bandit problems with infinitely many arms. *The Annals of Statistics*, 25(5):2103–2116, 1997. (Not cited.)
- [25] Ashton Anderson, Daniel Huttenlocher, Jon Kleinberg, and Jure Leskovec. Discovering value from community activity on focused question answering sites: a case study of stack overflow. In *Proceedings of the 18th ACM SIGKDD*, pages 850–858, 2012. (Not cited.)
- [26] Abdolhossein Hoorfar and Mehdi Hassani. Inequalities on the Lambert W function and hyperpower function. *J. Inequal. Pure and Appl. Math*, 9(2):5–9, 2008. (Not cited.)
- [27] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the Lambert W function. *Advances in Computational Mathematics*, 5(1):329–359, 1996. ISSN 1572-9044. (Not cited.)
- [28] Jean-Yves Audibert and Sébastien Bubeck. Regret bounds and minimax policies under partial monitoring. *Journal of Machine Learning Research*, 11(Oct):2785–2836, 2010. (Not cited.)
- [29] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 322–331. IEEE, 1995. (Not cited.)

[30] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002. (Not cited.)

Appendix

A Omitted Proofs

Claim 1. The minimum regret is achieved when algorithm considers the first G arms.

Proof. Suppose there exists an arm $i \notin G$ such that $i \leq |G|$, then there exists a corresponding arm $j \in G$ such that $j > |G|$. Construct a set $G' = G \cup \{i\} \setminus \{j\}$. As the best arm is uniformly distributed, both i and j have equal probability of being an optimal arm. Thus, we have $P(i^* \in G) = P(i^* \in G')$. Hence, the internal regret of the algorithm will be the same. The external regret will decrease as $j > |G| \geq i$. \square

Property 1. The Lambert W function can be equivalently written as the inverse of the function $f(x) := xe^x$, i.e., $W(xe^x) = x$.

Proof. The forward direction is straightforward. As $W(\cdot)$ is one to one function in the non-negative domain, we have $W(W(x)e^{W(x)}) = W(x)$. Let $y = W(x)$ then we have $W(ye^y) = y$. To show that $W(xe^x) = x$ implies $W(z)e^{W(z)} = z$, observe that $W(W(z_0)e^{W(z_0)}) = W(z_0)$. We get the required result by taking the inverse. \square

Property 2. For any $x \in [0, \infty)$, the Lambert W function is unique, non-negative, and strictly increasing.

Proof. Observe that $W(0) = 0$. Note that $f(x) = xe^x$ is continuous, one to one and strictly increasing. Hence, its inverse, $W(\cdot)$, is also increasing. \square

Property 3. For any $x \geq e$, we have $\log(x)/2 < W(x) \leq \log(x)$.

Proof. By definition, the Lambert W function satisfies $W(x)e^{W(x)} = x$. It is easy to see that $W(e) = 1$. Further we have,

$$\begin{aligned} 1 &= \frac{dW(x)}{dx} \cdot e^{W(x)} + \frac{dW(x)}{dx} \cdot e^{W(x)}W(x) \\ &= \frac{dW(x)}{dx}(x + e^{W(x)}) \\ \implies \frac{dW(x)}{dx} &= \frac{1}{x + e^{W(x)}} \end{aligned}$$

Let $f(x) = \log(x) - W(x)$. We have that $f(e) = 0$. We have that $\frac{df(x)}{dx} = \frac{1}{x} - \frac{1}{x + e^{W(x)}} = \frac{e^{W(x)}}{x(x + e^{W(x)})} = \frac{1}{x(1 + W(x))} > 0$. Hence we have that $f(\cdot)$ is increasing i.e. $f(x) > 0$ for all $x > e$. This shows that $W(x) \leq \log(x)$.

Now, let $g(x) = \frac{\log(x)}{2} - W(x)$. Here, we have $g(e) < 0$, and $\frac{dg(x)}{dx} = \frac{1}{2x} - \frac{1}{x + e^{W(x)}} = \frac{e^{W(x)} - e^{\log(x)/2}}{2x(x + e^{W(x)})} \leq 0$ (since $W(x) \leq \log(x)$ for $x \geq e$). So, $g(x) < 0$ for all $x \geq e$, implying that $\frac{\log(x)}{2} < W(x)$. This completes the proof. \square

Claim 2. $\frac{W(\lambda T/c)}{\lambda T/c} < 1/36 \iff T > \frac{36c \log(36)}{\lambda}$

Proof. We have the following equivalent inequalities.

$$\begin{aligned}
& \frac{W(\lambda T/c)}{\lambda T/c} < \frac{1}{36} \\
\iff e^{-W(\lambda T/c)} < \frac{1}{36} & \quad (\because W(x)e^{W(x)} = x) \\
\iff W(\lambda T/c) > \log(36) \\
\iff \frac{\lambda T}{c} > \log(36)e^{\log(36)} \\
\iff T > \frac{36c \log(36)}{\lambda}
\end{aligned}$$

The second to last inequality is obtained by applying the monotone increasing function $f(x) := xe^x$ on both sides, and then using Definition 1 of Lambert W function. \square

Claim 3. $e^{-cW(\lambda T/c)}$ is decreasing in c for $c \in (0, 1/2]$.

Proof. For $c_1 > c$, we have

$$\begin{aligned}
& \lambda T/c > \lambda T/c_1 \\
\iff W(\lambda T/c) > W(\lambda T/c_1) & \quad (\text{Property P2 of Lambert } W) \\
\iff e^{-W(\lambda T/c)} < e^{-W(\lambda T/c_1)} \\
\iff \frac{W(\lambda T/c)}{\lambda T/c} < \frac{W(\lambda T/c_1)}{\lambda T/c_1} & \quad (\because W(x)e^{W(x)} = x) \\
\iff cW(\lambda T/c) < c_1W(\lambda T/c_1) \\
\iff e^{-cW(\lambda T/c)} > e^{-c_1W(\lambda T/c_1)}
\end{aligned}$$

\square