
Non-Stationary Bandits under Recharging Payoffs: Improved Planning with Sublinear Regret

Orestis Papadigenopoulos

Department of Computer Science
The University of Texas at Austin
papadig@cs.utexas.edu

Constantine Caramanis

Department of Electrical and Computer Engineering
The University of Texas at Austin
constantine@utexas.edu

Sanjay Shakkottai

Department of Electrical and Computer Engineering
The University of Texas at Austin
sanjay.shakkottai@utexas.edu

Abstract

The stochastic multi-armed bandit setting has been recently studied in the non-stationary regime, where the mean payoff of each action is a non-decreasing function of the number of rounds passed since it was last played. This model captures natural behavioral aspects of the users which crucially determine the performance of recommendation platforms, ad placement systems, and more. Even assuming prior knowledge of the mean payoff functions, computing an optimal planning in the above model is NP-hard, while the state-of-the-art is a $1/4$ -approximation algorithm for the case where at most one arm can be played per round. We first focus on the setting where the mean payoff functions are known. In this setting, we significantly improve the best-known guarantees for the planning problem by developing a polynomial-time $(1 - 1/e)$ -approximation algorithm (asymptotically and in expectation), based on a novel combination of randomized LP rounding and a time-correlated (interleaved) scheduling method. Furthermore, our algorithm achieves improved guarantees – compared to prior work – for the case where more than one arm can be played at each round. Moving to the bandit setting, when the mean payoff functions are initially unknown, we show how our algorithm can be transformed into a bandit algorithm with sublinear regret.

1 Introduction

In the last two decades, the predominant rise of the social media industry has made the notion of a *newsfeed* an integral part of our lives. In a newsfeed, a user observes a structured sequence of content items (posts, photos etc.) particularly selected by the platform according to her/his preferences. Apart from social media, an analogous idea – potentially relabeled – also appears in different domains as, for example, “frequently bought together” in e-commerce, “shuffling similar songs” in music recommendation, or “recommended articles” in scholarly literature indexing databases. Whether it is measured in terms of click-rate or time devoted, the high-level objective of newsfeeds is fairly well-known: to maximize the user’s engagement with the platform. In many applications, however, achieving this objective is not as simple as identifying the user’s “favorite” content, given that her/his satisfaction can depend on the time passed since the same (or similar) content has been observed. As an example, a user’s engagement can worsen if a social media feed (resp., a music recommendation platform) constantly presents content from the same source (resp., same artist).

Motivated by such scenarios, researchers have recently studied online decision making problems capturing the notion of “recovering” payoffs, namely, scenarios where the payoff of an action drops (to zero) after each play and then slowly increases back to a baseline. In the context of online learning, these non-stationary models interpolate between multi-armed bandits, where the environment is assumed to be intact, and reinforcement learning, since the actions may now alter the future environment in a structured manner. These models are wide enough to capture many real-life applications, yet special enough to accept efficiently computable (near-)optimal solutions. In this direction, the following general model was first introduced (under slightly different assumptions) by Immorlica and Kleinberg [IK18] and has been recently studied by Simchi-Levi et al. [SLZZ21]:

Problem (Multi-armed bandits under recharging payoffs). *We consider a set of actions (or arms), where each arm i is associated with a mean payoff function p_i . For each τ , $p_i(\tau)$ is the expected payoff collected for playing action i , when i has been played before exactly τ rounds (we refer to τ as the delay of an arm at a specific round). For each arm, we assume that its payoff function: (a) is monotone non-decreasing in τ , and (b) has a known finite recovery time, namely, a specific delay after which the function remains constant. At each round, the player pulls a subset of at most k actions, observes the realized payoffs of the arms played (each computed with respect to the delay of the corresponding action), and collects their sum. The objective is to maximize the expected cumulative payoff collected within an unknown time horizon.*

We denote by k -RB an instance of the problem where at most k arms can be played at each round. Further, we distinguish between the *planning* setting, where the payoff functions are known to the player a priori (thus, payoffs can be considered deterministic), and the *learning* setting, where these are initially unknown and the user assumes semi-bandit feedback on the payoffs of the played actions. In the former case, the goal is to design an efficient algorithm which closely-approximates the optimal planning (since this is generally an NP-hard problem), while in the latter, the objective is to construct a bandit algorithm of sublinear regret, defined as the difference in expected payoff between a planning algorithm and its bandit variant due to the initial absence of information of the latter.

1.1 Background and Related Work

Immorlica and Kleinberg [IK18] first study the 1-RB problem and provide a $(1 - \epsilon)$ -approximation (asymptotically) for the planning setting, under the additional assumption that the payoff functions are weakly concave over the whole time horizon. On the other extreme, Basu et al. [BSS19] provide a long-run $(1 - 1/e)$ -approximation for the case where, after each play, an arm *cannot* be played again for a fixed number of rounds – a problem which can be cast as a special instance of 1-RB using Heaviside step payoff functions.¹ Shortly after the introduction of the problem, a number of special cases and variations have been studied [CCB20, PC21, PBG19, BPCS21, BCMTT20, CCB20, LCCBGB22, APB⁺21] (we address the reader to Appendix A for an overview).

In their recent work, Simchi-Levi et al. [SLZZ21] study the k -RB problem and prove the first $\mathcal{O}(1)$ -approximation guarantee under no additional assumptions. For large k , this result is the state-of-the-art. However, for small k , understanding the magnitude of the constant becomes the primary theoretical question. Specifically, the motivating case of newsfeeds (when, for instance, the content items are presented to a user sequentially) can be modeled as an instance of 1-RB. In this case, the approximation guarantee of [SLZZ21] becomes $1/4$, which is significantly weaker compared to both the $(1 - \epsilon)$ -approximation of [IK18] for concave functions and the $(1 - 1/e)$ -approximation of [BSS19] for the most “extreme” example of monotone convex functions, i.e., that of Heaviside step functions. The above observation indicates that either the approximability status of the problem is not well-understood, or that the problem does not gradually become “easier” by increasing the concavity of the payoff functions.

1.2 Technical Contributions

In our work, we resolve this discrepancy by designing a long-run $(1 - 1/e)$ -approximation algorithm (in expectation) for 1-RB, which improves the state-of-the-art due to [SLZZ21]. Simultaneously, our algorithm enjoys the same asymptotic guarantee of $(1 - \mathcal{O}(1/\sqrt{k}))$ for the general case of k -RB as in Simchi-Levi et al. [SLZZ21] with improved and explicit constants, as opposed to the guarantees of [SLZZ21] which do not come in a closed-form (see Table 1 for a summary of the best-known

¹Recall that the Heaviside step function is defined such that $H(\tau) = 1$ for $\tau \geq 0$, and $H(\tau) = 0$, otherwise.

Payoff functions	State-of-the-art	Reference	This work
Non-decreasing, concave, Lipschitz	$(1 - \epsilon)$ -PTAS for $k = 1$	[IK18]	–
Scaled/translated Heaviside step	$1 - \frac{1}{e}$ for $k = 1$	[BSSS19, PC21]	–
//	$\max_{a \in \mathbb{Z}_+} \frac{a}{a+1} \frac{k}{k+a}$ for any k	[SLZZ21]	$1 - \frac{k^k}{e^k k!}$
Non-decreasing (finite recovery)	$\frac{1}{4}$ for $k = 1$	[SLZZ21]	$1 - \frac{1}{e}$
//	$\max_{a \in \mathbb{Z}_+} \frac{a}{a+1} \frac{k}{k+a}$ for any k	[SLZZ21]	$1 - \frac{k^k}{e^k k!}$

Table 1: Summary of existing results (approximation guarantees) and our improvements.

approximation guarantees). Our algorithm is based on a novel combination of linear programming (LP) rounding and a time-correlated (interleaved) scheduling method, and is significantly simpler to implement compared to prior work. For the case where the mean payoff functions are initially unknown, we show how our algorithm can be transformed into a bandit algorithm with sublinear regret guarantees.

1.3 Key Ideas and Intuition

Henceforth, we refer to any instance where the payoff function of each action i has the form $q_i(\tau) = p_i \cdot H(\tau - d_i)$, where $p_i \in [0, 1]$, $d_i \in \mathbb{N}$, and $H(\cdot)$ the Heaviside step function, as the *Heaviside k -RB* problem. The fact that our $(1 - 1/e)$ -approximation for 1-RB matches the state-of-the-art bound obtained for Heaviside 1-RB [BSSS19, PC21] is no coincidence. Our solution technique shows that each k -RB problem is in fact (approximately) hiding a Heaviside k -RB problem.

Structural characterization of a natural LP relaxation. A key idea in both [IK18] and [SLZZ21] is to construct a concave relaxation of the optimal solution. Instead, we take a more direct approach and construct a natural LP relaxation (see Section 3) of the optimal average payoff (our construction uses a uniform upper bound on the recovery time of any arm). By carefully analyzing the structure of our LP, we prove that its extreme point solutions follow a particularly interesting *sparsity pattern*: there exist unique delays $\{\tau_i\}$ associated with the arms (which we refer to as “critical” delays), such that playing each arm i exactly once every τ_i rounds matches the average payoff of the relaxation. As we show, an exception to the above rule can be at most a single arm, which we refer to as *irregular*. The above observation already hints that our problem could potentially be reduced to an instance of Heaviside k -RB for which better approximation guarantees (compared to [SLZZ21]) are known to exist (at least for small k) [BSSS19, PC21].

Improved approximation guarantees for the planning k -RB. Constructing a planning where each arm is played at a rate indicated by its critical delay (as described above) is generally infeasible, since these rates stem from a relaxation of the problem. However, targeting these rates becomes the starting point of our approximation. In Section 4, we design a novel algorithm for the k -RB problem, called “Randomize-Then-Interleave”, which starts from computing an optimal solution to our LP relaxation. By combining the sparse structure of this solution with a *randomized rounding* step for determining a critical delay for the (possible) irregular arm, our algorithm produces a “proxy” instance of the Heaviside k -RB problem. Finally, the algorithm applies the *interleaved scheduling* method of [PC21] on the resulting Heaviside k -RB instance. As we prove, the above algorithm achieves a long-run $(1 - \mathcal{O}(1/\sqrt{k}))$ -approximation guarantee (in expectation), but with improved and explicit constants compared to [SLZZ21]. Importantly, for the particular case of k -RB, our guarantee is stronger than the one provided in the original work of [PC21].

As a first step, in Section 4.1 we study the performance of our algorithm under the simplifying assumption that the solution returned by our LP relaxation does not contain an irregular arm, that is, every arm is associated with a unique critical delay. For this case, the approximation guarantee of our algorithm follows by critically leveraging existing results on the *correlation gap* of the weighted rank function of uniform matroids due to [Yan11] (see Section 2). In Section 4.2, we relax this simplifying assumption by carefully studying the contribution of the irregular arm to the produced solution. Our strategy relies on the novel idea of considering parallel (fictitious) copies of the irregular arm, one for each possible critical delay. While the marginal probabilities of each copy being played under the corresponding delay at any round are consistent with the initial LP solution, these events are *mutually*

exclusive. Nevertheless, we show that expected payoff collected can only decrease by assuming that the parallel copies were instead independent. Having established that, the analysis reduces back to the “easier” case where an irregular arm does not exist.

A bandit adaptation with sublinear regret. Finally, we turn our focus to the learning setting, where the mean payoff functions are initially unknown, and the player has semi-bandit feedback on the realized payoffs of the played arms of each round. In this setting, we begin by analyzing the robustness of our planning algorithm under small perturbations of the payoff functions. Then, we provide sample complexity results to bound the number of samples required to get accurate estimates of the mean payoff functions. By combining the above elements, we transform our planning algorithm into a bandit one, based on an Explore-then-Commit (ETC) scheme, and prove that the latter achieves sublinear regret relative to its planning counterpart. We present these results in Section 5.

All the omitted proofs of our results have been moved to the Appendix.

2 Preliminaries

Problem definition and notation. We consider a set $\mathcal{A} = \{1, 2, \dots, n\}$ of n actions (or arms), where each arm $i \in \mathcal{A}$ is associated with a (mean) *payoff function* $p_i : \mathbb{N} \rightarrow [0, 1]$. For each $\tau \in \mathbb{N}$, $p_i(\tau)$ is the expected payoff for playing action $i \in \mathcal{A}$, when i has been played before exactly τ rounds (i.e., $\tau = 1$ if the action has been played in the previous round). We refer to the number of rounds since the last play of an arm as the *delay*. For each arm $i \in \mathcal{A}$, we assume that the function $p_i(\tau)$: (a) is monotone non-decreasing in τ , and (b) has a polynomial *recovery time* τ_i^{\max} such that $p_i(\tau) = p_i(\tau_i^{\max})$ for every $\tau \geq \tau_i^{\max}$. We assume knowledge of a universal upper bound τ^{\max} on the recovery time of all arms, i.e., $\tau^{\max} \geq \max_{i \in \mathcal{A}} \tau_i^{\max}$.² At each round, the player plays a subset of at most $k < n$ actions and collects the sum of their associated payoffs (each computed with respect to the delay of the corresponding action at the time it is played). The objective is to maximize the cumulative payoff collected within an (unknown) horizon of T rounds. We refer to the above instance as k -RB.

As a convention, we assume that the delay of each arm i is initially equal to 1 (i.e., all arms are played at time $t = 0$). Note that, even if we assume that all arms start with delay τ^{\max} (i.e., the arms have never been played before), our results do not change qualitatively.

For any non-negative integer $n \in \mathbb{N}$, we define $[n] = \{1, 2, \dots, n\}$. For any vector $\mathbf{w} \in \mathbb{R}^n$ and set $S \subseteq [n]$, we denote $\mathbf{w}(S) = \sum_{i \in S} w_i$. Finally, throughout this work we assume that, when comparing between different payoffs, ties are broken arbitrarily.

Continuous extensions and the correlation gap. Let $f : 2^{[n]} \rightarrow [0, \infty)$ be a set function defined over a finite set of n elements. For any vector $\mathbf{y} \in [0, 1]^n$, we denote by $\mathcal{D}(\mathbf{y})$ a distribution over $2^{[n]}$ with marginal probabilities $\mathbf{y} = (y_i)_{i \in [n]}$. We denote by $S \sim \mathcal{D}(\mathbf{y})$ a random subset $S \subseteq [n]$ sampled from the distribution $\mathcal{D}(\mathbf{y})$. Further, we denote by $S \sim \mathcal{I}(\mathbf{y})$ a random subset $S \subseteq [n]$ which is constructed by adding each element $i \in [n]$ to S , independently, with probability y_i .

We recall two canonical continuous extensions of a set function (see [CCPV07, Sch03]):

Definition 2.1 (Multi-linear extension). *For any vector $\mathbf{y} \in [0, 1]^n$, the multi-linear extension of a set function f is defined as*

$$F(\mathbf{y}) = \mathbb{E}_{S \sim \mathcal{I}(\mathbf{y})} [f(S)] = \sum_{S \subseteq [n]} f(S) \prod_{i \in S} y_i \prod_{i \notin S} (1 - y_i).$$

Definition 2.2 (Concave closure). *For any vector $\mathbf{y} \in [0, 1]^n$, the concave closure of a set function f is defined as*

$$f^+(\mathbf{y}) = \sup_{\mathcal{D}(\mathbf{y})} \mathbb{E}_{S \sim \mathcal{D}(\mathbf{y})} [f(S)] = \sup_{\boldsymbol{\alpha}} \left\{ \sum_{S \subseteq [n]} \alpha_S f(S) \mid \sum_{S \subseteq [n]} \alpha_S \mathbf{1}_S = \mathbf{y}, \sum_{S \subseteq [n]} \alpha_S = 1, \boldsymbol{\alpha} \succeq 0 \right\},$$

where $\mathbf{1}_S \in \{0, 1\}^n$ is an indicator vector such that $(\mathbf{1}_S)_i = 1$, if $i \in S$, and $(\mathbf{1}_S)_i = 0$, otherwise.

²Due to the polynomial recovery time assumption, oracle access to the payoff functions is not required.

For any non-negative weight vector $\mathbf{w} \in [0, \infty)^n$, of particular importance in the analysis of our algorithm is the function $f_{\mathbf{w},k} : 2^{[n]} \rightarrow [0, \infty)$, defined as

$$f_{\mathbf{w},k}(S) = \max\{\mathbf{w}(I) \mid I \subseteq S \text{ and } |I| \leq k\}. \quad (1)$$

We remark that $f_{\mathbf{w},k}$ corresponds to the *weighted rank function* of the (rank- k) uniform matroid under a weight vector \mathbf{w} and, hence, is non-decreasing submodular [Sch03].³ We denote by $F_{\mathbf{w},k}$ and $f_{\mathbf{w},k}^+$ the multi-linear extension and the concave closure of $f_{\mathbf{w},k}$, respectively.

The *correlation gap* [ADSY10] of a set function is defined as $\sup_{\mathbf{y} \in [0,1]^n} \frac{f^+(\mathbf{y})}{F(\mathbf{y})}$. The following result due to [Yan11] provides an upper bound on the correlation gap of $f_{\mathbf{w},k}(\cdot)$:

Lemma 2.3 (Correlation gap [Yan11]). *Let $f_{\mathbf{w},k} : 2^{[n]} \rightarrow [0, \infty)$ be the weighted rank function of the rank- k uniform matroid. Then, for any vector $\mathbf{y} \in [0, 1]^n$ we have*

$$f_{\mathbf{w},k}^+(\mathbf{y}) \geq F_{\mathbf{w},k}(\mathbf{y}) \geq \left(1 - \frac{k^k}{e^k k!}\right) \cdot f_{\mathbf{w},k}^+(\mathbf{y}) \approx \left(1 - \frac{1}{\sqrt{2\pi k}}\right) \cdot f_{\mathbf{w},k}^+(\mathbf{y}).$$

3 Structural Properties of a Natural LP Relaxation

We consider the following LP relaxation which applies to any planning instance of k -RB:

$$\begin{aligned} \text{maximize:} \quad & \sum_{\mathbf{x} \succeq \mathbf{0}} \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} p_i(\tau) \cdot x_{i,\tau} & (\mathbf{LP}_k) \end{aligned}$$

$$\text{s.t.:} \quad \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} x_{i,\tau} \leq k, \quad (C.1)$$

$$\sum_{\tau \in \mathbb{N}} \tau \cdot x_{i,\tau} \leq 1, \forall i \in \mathcal{A}. \quad (C.2)$$

In the above formulation, each variable $x_{i,\tau}$ denotes the fraction of time where arm i is played under delay τ – namely, exactly τ time steps after it was played for the last time. In (\mathbf{LP}_k) , constraint (C.1) follows by the fact that the total fraction of time where any action can be played (under any delay) cannot be more than k . Further, constraints (C.2) are valid constraints for the fractional plays of any individual arm under different delays.

We start by proving that the optimal solution of (\mathbf{LP}_k) is asymptotically an upper bound on the optimal average expected payoff:

Lemma 3.1. *For any instance of the k -RB problem, let V^* be the optimal value of (\mathbf{LP}_k) and $\text{OPT}(T)$ be the maximum payoff that can be collected in a time horizon of T rounds. Then $T \cdot V^* \geq \text{OPT}(T)$.*

Assuming knowledge of an upper bound $\tau^{\max} \geq \max_{i \in \mathcal{A}} \{\tau_i^{\max}\}$ on the maximum recovery time of any action, an optimal extreme point solution to (\mathbf{LP}_k) can be computed in polynomial time:

Fact 3.2. *There exists an optimal solution to (\mathbf{LP}_k) which is supported in $(x_{i,\tau})_{i \in \mathcal{A}, \tau \in [\tau_i^{\max}]}$. Further, given an upper bound τ^{\max} on the maximum recovery time of any action, an optimal extreme point solution to (\mathbf{LP}_k) can be computed in polynomial time.*

Let us introduce the notion of *supported* actions:

Definition 3.3 (Supported actions). *For any (potentially infeasible) solution \mathbf{x} to (\mathbf{LP}_k) , we say that an action $i \in \mathcal{A}$ is supported in \mathbf{x} , denoted by $i \in \mathcal{A}(\mathbf{x})$, if there exists $\tau \in \mathbb{N}$ such that $x_{i,\tau} > 0$.*

We introduce the following class of solutions:

Definition 3.4 (Delay-feasible solutions). *We say that a vector $\mathbf{x} = (x_{i,\tau})_{i \in \mathcal{A}, \tau \in \mathbb{N}}$ is delay-feasible, if for each supported action $i \in \mathcal{A}(\mathbf{x})$ there exists at most one non-zero variable x_{i,τ_i} , and for this variable it holds $x_{i,\tau_i} = 1/\tau_i$. Further, we say that \mathbf{x} is almost-delay-feasible, if it is delay-feasible, possibly with the exception of a single arm $\iota \in \mathcal{A}(\mathbf{x})$, for which there exist at most two non-zero variables $x_{\iota,\tau_{\iota,a}}, x_{\iota,\tau_{\iota,b}}$, such that $x_{\iota,\tau_{\iota,a}} < 1/\tau_{\iota,a}$ and $x_{\iota,\tau_{\iota,b}} < 1/\tau_{\iota,b}$. In the case of almost-delay-feasible solutions, we refer to ι as the irregular arm.*

³A function $f : 2^{[n]} \rightarrow \mathbb{R}$ is submodular, if for any $S, T \subseteq [n]$, it holds $f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$.

We remark that delay-feasible solutions subsume by definition almost-delay-feasible solutions and that a (almost-)delay-feasible solution is not necessarily a feasible solution to (\mathbf{LP}_k) .

As we prove in the next key-lemma, any extreme point solution of (\mathbf{LP}_k) is almost-delay-feasible:

Lemma 3.5 (Sparsity pattern of extreme point solutions). *Let \mathbf{x} be any extreme point solution of (\mathbf{LP}_k) . Then, \mathbf{x} is almost-delay-feasible.*

4 Improved Approximation Guarantees for Planning

For any instance of k -RB, our algorithm (see Algorithm 1 below), called ‘‘Randomize-Then-Interleave’’, constructs a feasible planning schedule without requiring knowledge of the time horizon. We remark that for the planning setting of the problem, we treat the payoff of any arm under any possible delay as deterministic.

The algorithm starts from computing an optimal extreme point solution \mathbf{x}^* to (\mathbf{LP}_k) , and then uses this solution to determine a *critical* delay τ_i^* for each supported arm $i \in \mathcal{A}(\mathbf{x}^*)$. By Lemma 3.5, we know that \mathbf{x}^* is almost-delay-feasible, which means that – possibly with the exception of a single irregular arm ι – for every arm $i \in \mathcal{A}(\mathbf{x}^*) \setminus \{\iota\}$, there exists a unique τ_i , such that $x_{i,\tau_i}^* = 1/\tau_i > 0$. Thus, Algorithm 1 sets this unique τ_i to be the critical delay of each arm $i \in \mathcal{A}(\mathbf{x}^*) \setminus \{\iota\}$. In the case where an irregular arm $\iota \in \mathcal{A}(\mathbf{x}^*)$ exists, then, by Definition 3.4, there exist at most two distinct $\tau_{\iota,a}, \tau_{\iota,b}$ with $x_{\iota,\tau_{\iota,a}}^* > 0$ and $x_{\iota,\tau_{\iota,b}}^* \geq 0$, in which case the critical delay τ_ι^* of ι is set to $\tau_{\iota,a}$ or $\tau_{\iota,b}$, with marginal probability $\tau_{\iota,a}x_{\iota,\tau_{\iota,a}}^*$ and $\tau_{\iota,b}x_{\iota,\tau_{\iota,b}}^*$, respectively. Notice that by constraints (C.2) of (\mathbf{LP}_k) , the above sampling process is well-defined, since $\tau_{\iota,a}x_{\iota,\tau_{\iota,a}}^* + \tau_{\iota,b}x_{\iota,\tau_{\iota,b}}^* \leq 1$. In the case where no τ_ι^* is sampled, the irregular arm is removed from $\mathcal{A}(\mathbf{x}^*)$. After computing the critical delays, for each arm $i \in \mathcal{A}(\mathbf{x}^*)$ the algorithm draws a random *offset* r_i independently and uniformly from $\{0, 1, \dots, \tau_i^* - 1\}$.

Algorithm 1: Randomize-Then-Interleave

```

/* Initialization phase */
Compute an optimal extreme point solution  $\mathbf{x}^*$  to  $(\mathbf{LP}_k)$ .
Let  $\iota \in \mathcal{A}$  be the irregular arm (if such an arm exists).
for each arm  $i \in \mathcal{A}(\mathbf{x}^*) \setminus \{\iota\}$  do
  | Let  $\tau_i^*$  be the unique  $\tau \in \mathbb{N}$  which satisfies  $x_{i,\tau}^* = 1/\tau$ .
if an irregular arm  $\iota \in \mathcal{A}(\mathbf{x}^*)$  exists then
  | Let  $\tau_{\iota,a}, \tau_{\iota,b} \in \mathbb{N}$  be such that  $x_{\iota,\tau_{\iota,a}}^* > 0$ ,  $x_{\iota,\tau_{\iota,b}}^* \geq 0$ , and  $\tau_{\iota,a} \neq \tau_{\iota,b}$ .
  | Set  $p_a \leftarrow \tau_{\iota,a}x_{\iota,\tau_{\iota,a}}^*$  and  $p_b \leftarrow \tau_{\iota,b}x_{\iota,\tau_{\iota,b}}^*$ .
  | Sample  $\tau_\iota^*$  from  $\{\tau_{\iota,a}, \tau_{\iota,b}, \infty\}$  with marginals  $p_a, p_b$ , and  $1 - p_a - p_b$ , respectively.
  | if  $\tau_\iota^* = \infty$  then
  |   | Remove  $\iota$  from  $\mathcal{A}(\mathbf{x}^*)$ .
for each arm  $i \in \mathcal{A}(\mathbf{x}^*)$  do
  | Sample an offset  $r_i$  uniformly at random from  $\{0, 1, \dots, \tau_i^* - 1\}$ .
/* Online phase */
for  $t = 1, 2, \dots$  do
  | Let  $C_t \subseteq \mathcal{A}$  be the subset of candidate arms, defined as  $C_t = \{i \in \mathcal{A}(\mathbf{x}^*) \mid t \bmod \tau_i^* \equiv r_i\}$ .
  | Play the maximum-payoff subset of  $k$  arms in  $C_t$ , namely,  $A_t = \operatorname{argmax}_{S \subseteq C_t, |S| \leq k} \sum_{i \in S} p_i(\sigma_{i,t})$ ,
  |   where  $\sigma_{i,t}$  (actual delay) is the time passed since arm  $i$  was last played before  $t$ .

```

After initialization, Algorithm 1 proceeds to its online phase where, at each round t , it first computes a subset of *candidate* arms C_t , namely, all arms $i \in \mathcal{A}(\mathbf{x}^*)$ which satisfy $t \bmod \tau_i^* \equiv r_i$. Then, the algorithm computes the maximum-payoff subset A_t of at most k candidate arms, computed using their actual delays (the number of rounds passed since last play), and plays these arms.

4.1 Approximation Analysis: Assuming Delay-Feasible LP Solutions

We now present an analysis of the approximation guarantee of Algorithm 1 for general k . In order to facilitate the presentation, we first focus on a special class of ‘‘easy’’ instances, where the optimal

extreme point solution \mathbf{x}^* of (\mathbf{LP}_k) is delay-feasible (i.e., there is no irregular arm). Consequently, we show how this assumption can be relaxed at the expense of additional technical effort.

Recall that we denote by $\mathcal{A}(\mathbf{x}^*)$ the set of arms that are supported in the solution \mathbf{x}^* . Notice that, since \mathbf{x}^* is delay-feasible by assumption, the only source of randomness in Algorithm 1 is due to the random offsets. Since each offset r_i is sampled independently and uniformly from $\{0, 1, \dots, \tau_i^* - 1\}$ and $i \in C_t$ if and only if $t \bmod \tau_i^* \equiv r_i$, the following fact is immediate:

Fact 4.1. *Let $\{\tau_i^*\}_{i \in \mathcal{A}(\mathbf{x}^*)}$ be the critical delays of the supported arms. At any round t , each arm $i \in \mathcal{A}(\mathbf{x}^*)$ belongs to C_t independently with probability $1/\tau_i^*$.*

Let us focus on the expected payoff collected by Algorithm 1 at any round $t \geq \tau^{\max}$. By our simplifying assumption that \mathbf{x}^* is delay-feasible, for each supported arm $i \in \mathcal{A}(\mathbf{x}^*)$, there is a unique critical delay τ_i^* , such that $x_{i, \tau_i^*}^* = 1/\tau_i^*$ is the only non-zero variable of \mathbf{x}^* corresponding to i .

We construct a vector $\mathbf{y} = [0, 1]^n$, such that $y_i = 1/\tau_i^*$ for each $i \in \mathcal{A}(\mathbf{x}^*)$, and $y_i = 0$, otherwise. Further, let us define a vector $\mathbf{w} \in [0, \infty)^n$ such that $w_i = p_i(\tau_i^*)$ for each $i \in \mathcal{A}(\mathbf{x}^*)$, and $w_i = 0$, otherwise. Recall that for any fixed round $t \geq \tau^{\max}$, by definition of C_t and the construction of the offsets, the events $\{i \in C_t\}_{i \in \mathcal{A}}$ are independent. Hence, we can think of C_t as a random set, where each arm $i \in \mathcal{A}$ belongs to C_t independently with probability y_i , namely, $C_t \sim \mathcal{I}(\mathbf{y})$.

Let $\underline{A}_t = \operatorname{argmax}_{S \subseteq C_t, |S| \leq k} \sum_{i \in S} p_i(\tau_i^*)$ be a set of at most k arms in C_t which maximizes the total payoff measured with respect to the critical delays. By construction of C_t , the actual delay $\sigma_{i,t}$ of each $i \in \mathcal{A}(\mathbf{x}^*)$ is (deterministically) lower bounded by τ_i^* , since any arm must be a candidate in order to be played (which happens exactly every τ_i^* rounds). Thus, for the expected collected payoff of any round $t \geq \tau^{\max}$ we have

$$\mathbb{E}_{C_t \sim \mathcal{I}(\mathbf{y})} \left[\max_{S \subseteq C_t, |S| \leq k} \sum_{i \in S} p_i(\sigma_{i,t}) \right] \geq \mathbb{E}_{C_t \sim \mathcal{I}(\mathbf{y})} \left[\max_{S \subseteq C_t, |S| \leq k} \sum_{i \in S} p_i(\tau_i^*) \right] = \mathbb{E}_{C_t \sim \mathcal{I}(\mathbf{y})} [\mathbf{w}(\underline{A}_t)]. \quad (2)$$

Using Equation (1) in combination with Definitions 2.1 and 2.2, we get

$$\mathbb{E}_{C_t \sim \mathcal{I}(\mathbf{y})} [\mathbf{w}(\underline{A}_t)] = \mathbb{E}_{C_t \sim \mathcal{I}(\mathbf{y})} [f_{\mathbf{w},k}(C_t)] = F_{\mathbf{w},k}(\mathbf{y}) \geq \left(1 - \frac{k^k}{e^k k!}\right) \cdot f_{\mathbf{w},k}^+(\mathbf{y}), \quad (3)$$

where the last inequality follows by Lemma 2.3.

The following technical lemma relates $f_{\mathbf{w},k}^+(\mathbf{y})$ with optimal value of (\mathbf{LP}_k) :

Lemma 4.2. *For any vectors $\mathbf{w} \in [0, \infty)^n$ and $\mathbf{y} \in [0, 1]^n$, let $f_{\mathbf{w},k}^+(\mathbf{y})$ be the concave closure of the weighted rank function of the rank- k uniform matroid with weights \mathbf{w} , evaluated at \mathbf{y} . Then, if \mathbf{y} satisfies $\|\mathbf{y}\|_1 \leq k$, it is the case that $f_{\mathbf{w},k}^+(\mathbf{y}) \geq \sum_{i \in [n]} w_i y_i$.*

Let V^* be the optimal value of (\mathbf{LP}_k) . Since $\|\mathbf{y}\|_1 \leq k$, then by applying Lemma 4.2 with vectors \mathbf{w} and \mathbf{y} as constructed above, we have that

$$f_{\mathbf{w},k}^+(\mathbf{y}) \geq \sum_{i \in [n]} w_i \cdot y_i = \sum_{i \in \mathcal{A}(\mathbf{x}^*)} p_i(\tau_i^*) \cdot x_{i, \tau_i^*}^* = V^*. \quad (4)$$

By combining Equations (2) to (4) with Lemma 3.1, it follows that, for any round $t \geq \tau^{\max}$, the expected payoff collected by Algorithm 1 is at least $(1 - k^k/e^k k!)$ -times the average optimal payoff. By using linearity of expectations together with the fact that all payoffs are bounded in $[0, 1]$, we can prove a long-run $(1 - k^k/e^k k!)$ -approximation guarantee (in expectation) for instances where \mathbf{x}^* is delay-feasible.

In the rest of this section, we show how the above analysis can be extended for the general case where \mathbf{x}^* is almost-delay-feasible (i.e., an irregular arm exists).

4.2 Approximation Analysis: Handling the Irregular Arm

We now extend the analysis of Section 4.1 and relax the assumption that the optimal extreme point solution of (\mathbf{LP}_k) is delay-feasible.

Let \mathbf{x}^* be an optimal solution to (\mathbf{LP}_k) and let $\iota \in \mathcal{A}(\mathbf{x}^*)$ be the irregular arm with $\tau_{\iota,a}^*$ and $\tau_{\iota,b}^*$ being the two potential critical delays such that $x_{\iota,\tau_{\iota,a}^*}^*, x_{\iota,\tau_{\iota,b}^*}^* \geq 0$ (with at least one of $x_{\iota,\tau_{\iota,a}^*}^*$ and $x_{\iota,\tau_{\iota,b}^*}^*$ being strictly positive, by Definition 3.4). Instead of conditioning on the sampled critical delay of the irregular arm ι , we construct two (fictitious) parallel copies ι_a and ι_b – one for each possible realization (and associated payoff). Let us define $\mathcal{A}' = \mathcal{A}(\mathbf{x}^*) \setminus \{\iota\} \cup \{\iota_a, \iota_b\}$ to be the set of supported arms of \mathbf{x}^* , where we replace the irregular arm ι with the two parallel copies ι_a and ι_b .

In the next lemma, we characterize the marginal probability that any arm in \mathcal{A}' is added to the set of candidate arms at any round. Notice that, as opposed to Fact 4.1, the randomness in this case is taken over both the sampling of a critical delay for the irregular arm as well as the random offsets.

Lemma 4.3. *For any arm $i \in \mathcal{A}'$ and at any round t , we have that $\mathbb{P}[i \in C_t \text{ and } \tau_i^* = \tau] = x_{i,\tau}^*$.*

Notice that, for the parallel copies of the irregular arm, the events $\{\iota_a \in C_t\}$ and $\{\iota_b \in C_t\}$ are now dependent and, in particular, *mutually exclusive* (since, eventually, at most one of them is realized). However, we are able to show that the expected payoff collected at any round $t \geq \tau^{\max}$ can only decrease if the two events were instead independent with the same marginal probabilities.

For convenience, we relabel the arms in \mathcal{A}' in non-increasing order of $p_i(\tau_i^*)$ (breaking ties arbitrarily). Note that by exchanging ι with the two copies ι_a and ι_b the cardinality of \mathcal{A}' can be at most $n + 1$.

Let us construct a $(n + 1)$ -dimensional vector $\mathbf{y} \in [0, 1]^{n+1}$, such that $y_i = x_{i,\tau_i^*}^*$ for any $i \in \mathcal{A}'$, and $y_i = 0$, otherwise. Further, we construct a weight vector $\mathbf{w} \in [0, \infty)^{n+1}$, such that $w_i = p_i(\tau_i^*)$ for any $i \in \mathcal{A}'$, and $w_i = 0$, otherwise. Let $\mathcal{D}(\mathbf{y})$ be a distribution over subsets of \mathcal{A}' , where every arm $i \in \mathcal{A}' \setminus \{\iota_a, \iota_b\}$ (except for ι_a and ι_b) is added to a set $S \sim \mathcal{D}(\mathbf{y})$, independently, with probability $x_{i,\tau_i^*}^*$. The arms ι_a and ι_b are also added to $S \sim \mathcal{D}(\mathbf{y})$ with marginal probabilities $x_{\iota_a,\tau_{\iota_a}^*}^*$ and $x_{\iota_b,\tau_{\iota_b}^*}^*$, respectively, yet not independently (since they are mutually exclusive). Therefore, by the above definitions and Lemma 4.3, the set of candidate arms at any fixed round $t \geq \tau^{\max}$ is distributed as $C_t \sim \mathcal{D}(\mathbf{y})$.

Let us fix any round $t \geq \tau^{\max}$. Similarly to Section 4.1, by using the fact that the actual delay $\sigma_{i,\tau}$ of each arm $i \in \mathcal{A}'$ is deterministically lower bounded by its critical delay τ_i^* , we can lower bound the expected payoff collected by our algorithm at time t as

$$\mathbb{E}_{C_t \sim \mathcal{D}(\mathbf{y})} \left[\max_{S \subseteq C_t, |S| \leq k} \sum_{i \in S} p_i(\sigma_{i,t}) \right] \geq \mathbb{E}_{C_t \sim \mathcal{D}(\mathbf{y})} \left[\max_{S \subseteq C_t, |S| \leq k} \sum_{i \in S} p_i(\tau_i^*) \right] = \mathbb{E}_{C_t \sim \mathcal{D}(\mathbf{y})} [f_{\mathbf{w},k}(C_t)], \quad (5)$$

where $\mathcal{D}(\mathbf{y})$ is defined as described above.

As we show in the following lemma, the RHS of the above inequality can only decrease, if we replace $C_t \sim \mathcal{D}(\mathbf{y})$ with $C_t \sim \mathcal{I}(\mathbf{y})$, namely, if we assume that the two copies ι_a and ι_b are added to C_t independently, but with the same marginals (y_{ι_a} and y_{ι_b} , respectively). As we show in the following lemma, the above property in fact holds for any submodular function (including $f_{\mathbf{w},k}$):

Lemma 4.4. *Let $f : 2^{[n]} \rightarrow \mathbb{R}$ be a submodular function over a ground set of n elements, where $n \geq 2$. Let $\mathcal{I}(\mathbf{y})$ be a distribution over $2^{[n]}$, where each element i is added in set $S \sim \mathcal{I}(\mathbf{y})$, independently, with probability y_i . Further, let $\mathcal{D}(\mathbf{y})$ be a distribution over $2^{[n]}$, where each element $i \in [n] \setminus \{\iota_a, \iota_b\}$ is added to a set $S \sim \mathcal{D}(\mathbf{y})$, independently, with probability y_i , except for two distinct elements ι_a, ι_b , whose addition to S is mutually exclusive but with the same marginals y_{ι_a} and y_{ι_b} , respectively. In this case, we have*

$$\mathbb{E}_{S \sim \mathcal{D}(\mathbf{y})} [f(S)] \geq \mathbb{E}_{S \sim \mathcal{I}(\mathbf{y})} [f(S)].$$

By combining inequality (5) with Lemma 4.4 for $n + 1$ elements (recall, the function $f_{\mathbf{w},k}$ is submodular), the expected reward collected at any round $t \geq \tau^{\max}$ can be lower bounded as

$$\mathbb{E}_{C_t \sim \mathcal{D}(\mathbf{y})} \left[\max_{S \subseteq C_t, |S| \leq k} \sum_{i \in S} p_i(\sigma_{i,t}) \right] \geq \mathbb{E}_{C_t \sim \mathcal{D}(\mathbf{y})} [f_{\mathbf{w},k}(C_t)] \geq \mathbb{E}_{C_t \sim \mathcal{I}(\mathbf{y})} [f_{\mathbf{w},k}(C_t)] = F_{\mathbf{w},k}(\mathbf{y}).$$

In addition, by combining Lemma 2.3 with Lemma 4.2 for the above choice of vectors \mathbf{w} and \mathbf{y} (again, it holds $\|\mathbf{y}\|_1 \leq k$), for any $t \geq \tau^{\max}$, it follows that

$$\mathbb{E}_{C_t \sim \mathcal{D}(\mathbf{y})} \left[\max_{S \subseteq C_t, |S| \leq k} \sum_{i \in S} p_i(\sigma_{i,t}) \right] \geq F_{\mathbf{w},k}(\mathbf{y}) \geq \left(1 - \frac{k^k}{e^k k!}\right) \cdot f_{\mathbf{w},k}^+(\mathbf{y}) \geq \left(1 - \frac{k^k}{e^k k!}\right) V^*,$$

where V^* is the optimal value of (\mathbf{LP}_k) .

Finally, using the fact that $V^* \geq \text{OPT}(T)/T$ by Lemma 3.1 and applying linearity of expectations over all rounds t , the following result follows directly:

Theorem 4.5. *For any instance of k -RB, where $\text{OPT}(T)$ is the optimal expected payoff that can be collected in T rounds, the total expected payoff collected by Algorithm 1 is at least*

$$\left(1 - \frac{k^k}{e^k k!}\right) \text{OPT}(T) - \mathcal{O}(k \cdot \tau^{\max}) \approx \left(1 - \frac{1}{\sqrt{2\pi k}}\right) \text{OPT}(T) - \mathcal{O}(k \cdot \tau^{\max}).$$

We remark that, although the asymptotic approximation guarantee of Algorithm 1 is the same as the one in [SLZZ21], the constants of our algorithm are significantly better. In particular, for $k \in \{1, 2, 3, 4, 5, 10\}$, the guarantees of [SLZZ21] and Algorithm 1 are $\{0.25, 0.33, 0.40, 0.44, 0.46, 0.57\}$ and $\{0.63, 0.72, 0.77, 0.80, 0.82, 0.87\}$, respectively.

5 Online Learning with Sublinear Regret

In this section, we consider the learning setting of k -RB where the payoff functions are initially unknown. At each round, the player plays a subset of at most k arms and observes the realized payoff of each individual arm played (semi-bandit feedback). Every time an arm $i \in \mathcal{A}$ is played under delay τ , the realized payoff is drawn independently from a distribution of mean $p_i(\tau)$, bounded in $[0, 1]$. For this setting, our goal is to show that there exists a bandit variant of Algorithm 1 with sublinear γ_k -approximate regret guarantee for k -RB, defined as

$$\text{Reg}_{\gamma_k}(T) = \gamma_k \text{OPT}(T) - R(T),$$

where $\gamma_k = 1 - k^k/e^k k!$ is the (multiplicative) approximation guarantee of Algorithm 1 and $R(T)$ is the total expected payoff collected by the bandit algorithm in T rounds.

Robustness of Algorithm 1. As a first step, we study the robustness of Algorithm 1 under small perturbations of the payoff functions; in particular, we analyze its approximation guarantee when it runs using estimates $\{\hat{p}_i\}_{i \in \mathcal{A}}$ instead of the actual payoffs $\{p_i\}_{i \in \mathcal{A}}$. We remark that the monotonicity of the (estimated) payoff functions is *not required* for the correctness of Algorithm 1, but only for proving its approximation guarantee.

Lemma 5.1. *Let \hat{p}_i be an estimate of the payoff function p_i for each $i \in \mathcal{A}$, such that $|\hat{p}_i(\tau) - p_i(\tau)| \leq \epsilon, \forall \tau \in \mathbb{N}$ for some $\epsilon \in (0, 1)$. The expected payoff collected in T rounds by Algorithm 1, when it uses the estimates instead of the actual payoff functions, is lower bounded by*

$$\gamma_k \text{OPT}(T) - \mathcal{O}(k \cdot \epsilon \cdot T + k \cdot \tau^{\max}).$$

Estimation of the payoff functions. By using standard concentration results [Hoe63], we bound the number of samples required for each arm-delay pair in order to get an ϵ -estimate \hat{p}_i for the mean payoff function p_i of every arm $i \in \mathcal{A}$.

Lemma 5.2. *For any $\epsilon, \delta \in (0, 1)$, let $\hat{p}_i(\tau)$ be the empirical average of samples drawn from $p_i(\tau)$ for every arm $i \in \mathcal{A}$ and delay τ , using $m = \frac{1}{2\epsilon^2} \ln\left(\frac{2\tau^{\max} n}{\delta}\right)$ samples, where n is the number of arms. Then, with probability $1 - \delta$, it holds $|\hat{p}_i(\tau) - p_i(\tau)| \leq \epsilon$ for every $i \in \mathcal{A}$ and $\tau \in [\tau^{\max}]$.*

The following observation provides an upper bound on the number of rounds required in order to get m samples for each arm-delay pair in any k -RB instance:

Fact 5.3. *In any k -RB instance, the number of rounds required to collect m independent samples from $p_i(\tau)$, for each arm $i \in \mathcal{A}$ and delay $\tau \in [\tau^{\max}]$, is upper-bounded by $\mathcal{O}(nm(\tau^{\max})^2/k)$.*

Bandit algorithm with sublinear regret. By combining the above elements, it is easy to design an Explore-then-Commit (ETC) variant of Algorithm 1, achieving sublinear regret in the bandit setting. For simplicity, we assume here that the time horizon T is known to the player – an assumption that can be relaxed using the doubling trick method [LS20].

Let $m = 1/2\epsilon^2 \ln(2\tau^{\max}n/\delta)$ for some $\epsilon, \delta \in (0, 1)$. In the first $\mathcal{O}(nm(\tau^{\max})^2/k)$ rounds, the bandit algorithm collects m independent samples from $p_i(\tau)$, for each $i \in \mathcal{A}$ and $\tau \in [\tau^{\max}]$. Then, by taking the empirical average of these samples, it constructs estimates $\hat{p}_i(\tau)$ for each i and τ . For the rest of the time horizon, the algorithm simulates Algorithm 1, using the empirical estimates $\hat{p}_i(\tau)$ in place of the (initially unknown) payoff functions.

By setting the parameters ϵ, δ accordingly and combining Lemmas 5.1 and 5.2 with Fact 5.3, we prove the following result:

Theorem 5.4. *There exists a learning adaptation of Algorithm 1 in the semi-bandit setting, for which the γ_k -approximate regret for T rounds can be upper-bounded as*

$$\mathcal{O}\left(n^{1/3} \cdot (k \cdot \tau^{\max})^{2/3} \ln^{1/3}(\tau^{\max}nT) \cdot T^{2/3} + k \cdot \tau^{\max}\right).$$

Conclusion and Further Directions

Following the recent line of work on non-stationary bandits with recharging payoffs, we revisited one of the most general formulations studied in the area and significantly improved the state-of-the-art for the planning setting. In particular, when at most one arm is played per round, we designed an algorithm that collects at least 63% (asymptotically and in expectation) of the optimal payoff, improving the best-known 25%, due to prior work. By providing results on the robustness and sample complexity of our algorithm, we transformed the latter into a bandit algorithm with sublinear regret. Our work leaves a number of interesting open questions. An approximation guarantee of $(1 - 1/e)$ has been proved to be the best achievable – under standard complexity assumptions – for a number combinatorial problems. Therefore, an immediate future direction would be to either provide a matching hardness result for the planning setting of the problem, or to further improve on the best-known approximation. Another interesting question would be to explore whether one can design a bandit adaptation of our algorithm with improved regret guarantees. Finally, the empirical evaluation of our algorithm on real (or artificial) data is another interesting direction.

Acknowledgements

This work was partially funded by the NSF IFML Institute (NSF 2019844), the NSF AI-EDGE Institute (NSF 2112471), and NSF 1704778.

References

- [ACBF02] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- [ADSY10] Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. Correlation robust stochastic optimization. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '10*, page 1087–1096, USA, 2010. Society for Industrial and Applied Mathematics.
- [AGO19] Peter Auer, Pratik Gajane, and Ronald Ortner. Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *Conference on Learning Theory*, pages 138–158. PMLR, 2019.
- [APB⁺21] Alexia Atsidakou, Orestis Papadigenopoulos, Soumya Basu, Constantine Caramanis, and Sanjay Shakkottai. Combinatorial blocking bandits with stochastic delays. In *International Conference on Machine Learning*, pages 404–413. PMLR, 2021.
- [BCB12] Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and non-stochastic multi-armed bandit problems. *Machine Learning*, 5(1):1–122, 2012.
- [BCMTT20] Nicholas Bishop, Hau Chan, Debmalya Mandal, and Long Tran-Thanh. Adversarial blocking bandits. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 8139–8149. Curran Associates, Inc., 2020.
- [Ber00] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2nd edition, 2000.
- [BGZ14] Omar Besbes, Yonatan Gur, and Assaf Zeevi. Stochastic multi-armed-bandit problem with non-stationary rewards. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, NIPS'14*, page 199–207, Cambridge, MA, USA, 2014. MIT Press.
- [BPCS21] Soumya Basu, Orestis Papadigenopoulos, Constantine Caramanis, and Sanjay Shakkottai. Contextual blocking bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 271–279. PMLR, 2021.
- [BSSS19] Soumya Basu, Rajat Sen, Sujay Sanghavi, and Sanjay Shakkottai. Blocking bandits. *Advances in Neural Information Processing Systems*, 32, 2019.
- [CCB20] Leonardo Cella and Nicolò Cesa-Bianchi. Stochastic bandits with delay-dependent payoffs. In *International Conference on Artificial Intelligence and Statistics*, pages 1168–1177. PMLR, 2020.
- [CCPV07] Gruiua Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a submodular set function subject to a matroid constraint. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 182–196. Springer, 2007.
- [Git79] John C Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–164, 1979.
- [GMS10] Sudipto Guha, Kamesh Munagala, and Peng Shi. Approximation algorithms for restless bandit problems. *Journal of the ACM (JACM)*, 58(1):1–50, 2010.
- [Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [IK18] Nicole Immorlica and Robert Kleinberg. Recharging bandits. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 309–319, 2018.
- [KNMS10] Robert Kleinberg, Alexandru Niculescu-Mizil, and Yogeshwer Sharma. Regret bounds for sleeping experts and bandits. *Machine learning*, 80(2):245–272, 2010.

- [KZ17] N Bora Keskin and Assaf Zeevi. Chasing demand: Learning and earning in a changing environment. *Mathematics of Operations Research*, 42(2):277–307, 2017.
- [LCCBGB22] Pierre Laforgue, Giulia Clerici, Nicolò Cesa-Bianchi, and Ran Gilad-Bachrach. A last switch dependent analysis of satiation and seasonality in bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 971–990. PMLR, 2022.
- [LCM17] Nir Levine, Koby Crammer, and Shie Mannor. Rotting bandits. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [LKKLM21] Liu Leqi, Fatma Kilinc Karzan, Zachary Lipton, and Alan Montgomery. Rebounding bandits for modeling satiation effects. *Advances in Neural Information Processing Systems*, 34, 2021.
- [LR⁺85] Tze Leung Lai, Herbert Robbins, et al. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [LS20] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [PBG19] Ciara Pike-Burke and Steffen Grunewalder. Recovering bandits. volume 32. 2019.
- [PC21] Orestis Papadigenopoulos and Constantine Caramanis. Recurrent submodular welfare and matroid blocking semi-bandits. *Advances in Neural Information Processing Systems*, 34, 2021.
- [PT99] Christos H Papadimitriou and John N Tsitsiklis. The complexity of optimal queuing network control. *Mathematics of Operations Research*, 24(2):293, 1999.
- [Sch03] Alexander Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
- [SGV20] Aadirupa Saha, Pierre Gaillard, and Michal Valko. Improved sleeping bandits with stochastic action sets and adversarial rewards. In *International Conference on Machine Learning*, pages 8357–8366, July 2020.
- [SLZZ21] David Simchi-Levi, Zeyu Zheng, and Feng Zhu. Dynamic planning and learning under recovering rewards. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 9702–9711. PMLR, 18–24 Jul 2021.
- [Tho33] 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.
- [TL12] Cem Tekin and Mingyan Liu. Online learning of rested and restless bandits. *IEEE Transactions on Information Theory*, 58(8):5588–5611, 2012.
- [Whi88] Peter Whittle. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 25(A):287–298, 1988.
- [Yan11] Qiqi Yan. Mechanism design via correlation gap. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 710–719. SIAM, 2011.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] See "Supplementary material".
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] See "Supplementary material".
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 2.
 - (b) Did you include complete proofs of all theoretical results? [Yes] Complete proofs are included in the supplementary material.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Related Work

Since the work of Thompson [Tho33] and later [LR⁺85], the (stochastic) multi-armed bandits (MAB) problem in stationary environments has been thoroughly studied (see [BCB12, LS20, ACBF02] and references therein). In an attempt to capture non-stationary environments, where the payoff distributions change over time, a number of generalizations of MAB – which are also special cases of reinforcement learning (RL) – have been proposed. Two particularly interesting classes are *rested* and *restless* bandits [TL12], where each action is associated with a state-machine and has a different payoff distribution depending on its current state. In rested bandits [Git79, TL12], the state of each action (hence, its payoff distribution) changes stochastically only when the arm is played, while in restless bandits [Whi88, TL12] the state changes at every round, independently of the player’s actions. Optimal algorithms for rested and restless bandits essentially result from solving Bellman’s equations [Ber00] – an approach which requires an exponentially large state space in the number of arms. In fact, Papadimitriou and Tsitsiklis [PT99] show that it is PSPACE-hard to optimally solve (or even approximate) restless bandits, even in the simpler case where state transitions are deterministic. One of the first attempts to provide constant approximation algorithms for special cases of restless bandits is due to Guha et al. [GMS10], who give a 2-approximation for a special case of restless bandits, which they call “monotone” bandits. Another interesting class of non-stationary bandits is that of *sleeping* bandits [KNMS10, SGV20], which generalizes the standard MAB setting, given that only an adversarially chosen subset of actions can be played at each round. Finally, we address the reader to [BGZ14, LCM17, KZ17, AGO19] for additional non-stationary bandit models.

Shifting our focus to the case of “recharging” (or “recovering”) payoffs, Immorlica and Kleinberg [IK18] first consider the 1-RB problem, under the additional assumption that the payoff functions are weakly concave and weakly increasing over the whole time horizon. For this problem, the authors provide a PTAS that achieves a $(1 - \epsilon)$ -approximation (asymptotically) to the optimal planning, and show how it can be adapted into a learning algorithm with sublinear regret guarantees. To achieve their approximation, they first compute an “optimal” playing frequency for each arm via a concave relaxation, and then construct an approximate planning schedule, in which the rate of playing each arm closely approximates this frequency. In this direction, they authors combine enumeration techniques with the novel technique of *interleaved rounding*. The main idea behind the latter is that each arm is first associated with a particular (randomly perturbed) sequence of real numbers and, then, the arms are played sequentially according to the order they appear on the real line.

Basu et al. [BSSS19] study a variant of stochastic multi-armed bandits where, after each play, an arm becomes unavailable for a known number of rounds (blocking time). As they show, for the case of one arm per round, a simple greedy heuristic, which plays, at each round, the arm of highest expected payoff among the non-blocked arms, yields (asymptotically) a $(1 - 1/e)$ -approximation. In order to prove this guarantee, the authors compare the payoff collected by their algorithm with the optimal solution of an LP-relaxation, by carefully characterizing the latter via analyzing its optimality conditions. Although not stated explicitly in [BSSS19], the “hard” constraint of not pulling a blocked arm is equivalent to choosing the payoff function of each arm to be a Heaviside step function parameterized by its blocking time and scaled by its baseline mean payoff. In other words, the payoff drops to zero for a fixed number of rounds (which, in this case coincides with the recovery time), and then rises up to a baseline. Since pulling a zero-payoff arm can only harm the overall solution, the two problems are (without loss of generality) equivalent, and [BSSS19] becomes a special case of 1-RB. Notice, however, that the greedy approach of [BSSS19] fails to provide any non-trivial approximation guarantee for the 1-RB setting.

Shortly after its introduction, the model of Basu et al. [BSSS19] has been also studied in contextual [BPCS21], adversarial [BCMTT20], and combinatorial environments [APB⁺21, PC21]. In [PC21], Papadigenopoulos and Caramanis study the blocking setting in the regime where more than one arms can be played at each round, subject to matroid constraints. For the planning setting, they provide a $(1 - 1/e)$ -approximation algorithm (asymptotically and in expectation) for any matroid. Their technique, called *interleaved scheduling* – which is similar in spirit with the *interleaved rounding* method of [IK18] – is essentially a time-correlated arm-sampling scheme, which provides the following two guarantees: (i) at any fixed round, each arm i is sampled with probability $1/\tau_i$ (the maximum possible in any optimal solution), where τ_i is the blocking time of the arm. (ii) Every time an arm is sampled, its payoff must have returned to its baseline (i.e., is never blocked). Notice that the algorithm of [PC21], which essentially chooses a maximum-payoff independent

set among the sampled arms of each round, also provides an alternative method for reaching the $(1 - 1/e)$ -approximation of [BSSS19] for the special case of one arm per round.

Cella and Cesa-Bianchi [CCB20] consider a similar model to 1-RB, where they assume that all the actions have identical payoff functions up to different scaling (i.e., baseline payoff), but different and unknown recovery times. For the planning problem, by focusing on a simple class of *periodic ranking policies*, the authors develop an algorithm, accompanied by a worst-case approximation guarantee that depends on characteristics of the input (payoff functions, recovery times, and baselines). Finally, other special cases and variations of the k -RB problem include payoff functions described by linear dynamics [LKLM21], sampled by Gaussian processes [PBG19], and more [LCCBGB22].

Simchi-Levi et al. [SLZZ21] provide the first $\mathcal{O}(1)$ -approximation the general case of the k -RB problem. Specifically, by focusing on a particular class of policies, which they call *purely periodic*, the authors provide a long-run $(1 - \mathcal{O}(1/\sqrt{k}))$ -approximation for any instance of k -RB. A key-element in both Immorlica and Kleinberg [IK18] and, subsequently, Simchi-Levi et al. [SLZZ21] is the construction of a concave relaxation of the optimal solution, based on the critical notion of *rate of return*: a piecewise linear function, which denotes the maximum payoff that can be achieved asymptotically by playing an arm at most a specific fraction of rounds. The authors in [IK18] provide an FPTAS for computing an $(1 - \epsilon)$ -approximate solution to their relaxation (for the case where $k = 1$), while, in [SLZZ21], such a solution is efficiently computed assuming finite recovery times.

B Structural Properties of a Natural LP Relaxation: Omitted Proofs

Lemma 3.1. *For any instance of the k -RB problem, let V^* be the optimal value of (\mathbf{LP}_k) and $\text{OPT}(T)$ be the maximum payoff that can be collected in a time horizon of T rounds. Then $T \cdot V^* \geq \text{OPT}(T)$.*

Proof. Let $n_{i,\tau}$ be the number of times arm i is played under delay $\tau \in \mathbb{N}$ in an optimal solution. By fixing $x_{i,\tau} = \frac{n_{i,\tau}}{T}$, we have that

$$\text{OPT}(T) = \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} p_i(\tau) \cdot n_{i,\tau} = T \cdot \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} p_i(\tau) \cdot \frac{n_{i,\tau}}{T} = T \cdot \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} p_i(\tau) \cdot x_{i,\tau}.$$

Further, note that in any instance where at most k arms can be played at each round, the total number of arm plays (independently of the delay) cannot be more than $k \cdot T$. Thus, we get

$$\sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} x_{i,\tau} = \frac{1}{T} \cdot \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} n_{i,\tau} \leq \frac{1}{T} \cdot k \cdot T \leq k.$$

Finally, consider any action $i \in \mathcal{A}$ which is played at time t under delay τ . The number of rounds that are occupied for the above specific play belong to the interval $I = \{t - \tau + 1, t\}$, with $|I| = \tau$ (note that for $\tau = 1$ only round t is occupied). Recalling our assumption that all actions start with delay 1 at round $t = 1$, then for any such interval we have that $I \subseteq [T]$. Combining the above with the fact that, for any action $i \in \mathcal{A}$, these intervals cannot overlap, we get that

$$\sum_{\tau \in \mathbb{N}} \tau \cdot n_{i,\tau} \leq T \implies \sum_{\tau \in \mathbb{N}} \tau \cdot x_{i,\tau} \leq 1, \quad \forall i \in \mathcal{A}.$$

Therefore, since \mathbf{x} is a feasible solution to (\mathbf{LP}_k) with value equal to $\text{OPT}(T)/T$, the proof follows directly. \square

Fact 3.2. *There exists an optimal solution to (\mathbf{LP}_k) which is supported in $(x_{i,\tau})_{i \in \mathcal{A}, \tau \in [\tau_i^{\max}]}$. Further, given an upper bound τ^{\max} on the maximum recovery time of any action, an optimal extreme point solution to (\mathbf{LP}_k) can be computed in polynomial time.*

Proof. Clearly, proving that any optimal solution to (\mathbf{LP}_k) is without loss of generality supported in $(x_{i,\tau})_{i \in \mathcal{A}, \tau \in [\tau_i^{\max}]}$ implies that we can restrict ourselves to a polynomial-size version of (\mathbf{LP}_k) , which is constructed by dropping all variables except for $(x_{i,\tau})_{i \in \mathcal{A}, \tau \in [\tau_i^{\max}]}$. In this way, an optimal extreme point solution to this restricted LP can be computed efficiently using a standard LP solver.

Let \mathbf{x} be any optimal solution to (\mathbf{LP}_k) such that there exists an arm i and delay $\tau' > \tau_i^{\max}$ satisfying $x_{i,\tau'} > 0$. Note that by definition of the recovery time, it must hold $p_i(\tau') = p_i(\tau_i^{\max})$.

Thus, we can construct a solution \mathbf{x}' which is identical to \mathbf{x} , except for the fact that $x'_{i,\tau'} = 0$ and $x'_{i,\tau_i^{\max}} = x_{i,\tau_i^{\max}} + x_{i,\tau'}$. Since $p_i(\tau') = p_i(\tau_i^{\max})$ the solutions \mathbf{x}' and \mathbf{x} have exactly the same value, while \mathbf{x}' trivially satisfies constraints (C.1). In addition, constraints (C.2) are still satisfied by \mathbf{x}' for any arm different than i . For arm i , we have

$$\begin{aligned} \sum_{\tau \in \mathbb{N}} \tau \cdot x'_{i,\tau} &= \sum_{\tau \in \mathbb{N} \setminus \{\tau', \tau_i^{\max}\}} \tau \cdot x_{i,\tau} + \tau_i^{\max} \cdot x'_{i,\tau_i^{\max}} \\ &= \sum_{\tau \in \mathbb{N} \setminus \{\tau', \tau_i^{\max}\}} \tau \cdot x_{i,\tau} + \tau_i^{\max} \cdot x_{i,\tau_i^{\max}} + \tau_i^{\max} \cdot x_{i,\tau'} \\ &\leq \sum_{\tau \in \mathbb{N}} \tau \cdot x_{i,\tau} \\ &\leq 1, \end{aligned}$$

where the first inequality follows by the fact that $\tau' > \tau_i^{\max}$ and the second by feasibility of \mathbf{x} .

By repeating the above process for every arm i and delay $\tau > \tau_i^{\max}$ with $x_{i,\tau'} > 0$, we can transform any solution to (\mathbf{LP}_k) to an equivalent solution of the desired form. \square

Lemma 3.5 (Sparsity pattern of extreme point solutions). *Let \mathbf{x} be any extreme point solution of (\mathbf{LP}_k) . Then, \mathbf{x} is almost-delay-feasible.*

Proof. Let $\mathcal{A}(\mathbf{x}) \subseteq \mathcal{A}$ be the set of supported actions in an extreme point solution \mathbf{x} of (\mathbf{LP}_k) , and let $|\mathcal{A}(\mathbf{x})| = n'$. Let \mathbf{LP}' be the LP that results by dropping from (\mathbf{LP}_k) all the variables and constraints that correspond to the non-supported actions in $\mathcal{A} \setminus \mathcal{A}(\mathbf{x})$. Then, the solution $\mathbf{x}' = (x_{i,\tau})_{i \in \mathcal{A}(\mathbf{x}), \tau \in \mathbb{N}}$ (i.e., the solution \mathbf{x} restricted to the arms of $\mathcal{A}(\mathbf{x})$) is an extreme point solution of this reduced LP. In order to see that, let us assume that \mathbf{x}' is not an extreme point and, hence, $\mathbf{x}' = \lambda \mathbf{x}'_1 + (1 - \lambda) \mathbf{x}'_2$ for some $\lambda \in (0, 1)$ and $\mathbf{x}'_1, \mathbf{x}'_2$ feasible solutions to \mathbf{LP}' . Then, by padding \mathbf{x}'_1 and \mathbf{x}'_2 with zero variables for any $i \in \mathcal{A} \setminus \mathcal{A}(\mathbf{x})$ and $\tau \in \mathbb{N}$, the resulting vectors, let \mathbf{x}_1 and \mathbf{x}_2 , respectively, are feasible solutions to (\mathbf{LP}_k) . However, it also holds that $\mathbf{x} = \lambda \mathbf{x}_1 + (1 - \lambda) \mathbf{x}_2$, contradicting the fact that \mathbf{x} is an extreme point of (\mathbf{LP}_k) .

Now let us focus on the reduced formulation \mathbf{LP}' . By construction, since the actions in $\mathcal{A}(\mathbf{x}')$ are all supported in \mathbf{x}' , it has to be that $\|\mathbf{x}'\|_0 \geq n'$. Let u be the number of variables of \mathbf{LP}' . By counting constraint (C.1), constraints (C.2) (one for each $i \in \mathcal{A}(\mathbf{x}')$), and u non-negativity constraints, we can see that the number of constraints in \mathbf{LP}' is $u + n' + 1$. Hence, since \mathbf{x}' is an extreme point solution of \mathbf{LP}' , it must contain at most $n' + 1$ non-zero variables (since at least $u - n' - 1$ of the inequalities which are tight in \mathbf{x}' must correspond to non-negativity constraints). This gives $\|\mathbf{x}'\|_0 \leq n' + 1$ and, thus, we can conclude that $\|\mathbf{x}'\|_0 \in \{n', n' + 1\}$.

We now distinguish between two cases:

(a) If $\|\mathbf{x}'\|_0 = n'$, then since every arm in the solution \mathbf{x}' is supported and $\mathcal{A}(\mathbf{x}') = n'$, by pigeonhole principle each arm $i \in \mathcal{A}(\mathbf{x}')$ has to be supported by a unique non-zero variable $x'_{i,\tau_i} > 0$ for some $\tau_i \in \mathbb{N}$. Further, since exactly n' out of the u non-negativity constraints are not tight, then n' out of the $n' + 1$ inequalities in (C.1) and (C.2) have to be met with equality. (i) In the case where these n' constraints come from (C.2), then for each $i \in \mathcal{A}(\mathbf{x}')$, there exists a unique delay τ_i and variable x'_{i,τ_i} , such that $x'_{i,\tau_i} = 1/\tau_i$, implying that \mathbf{x}' and, thus, \mathbf{x} is delay-feasible. (ii) In the case where only $n' - 1$ from the tight constraints in (C.1) and (C.2) come from (C.2), then for each corresponding arm $i \in \mathcal{A}(\mathbf{x}')$, there exists a unique delay τ_i and variable x'_{i,τ_i} , such that $x'_{i,\tau_i} = 1/\tau_i$. Therefore, the arm $\iota \in \mathcal{A}(\mathbf{x}')$ which is associated with the non-tight constraint in (C.2) has to be an arm which is supported by a single variable $0 < x_{\iota,\tau_\iota} < 1/\tau_\iota$. In this case, \mathbf{x}' and, thus, \mathbf{x} is almost-delay-feasible, and ι corresponds to the irregular arm.

(b) In the case where $\|\mathbf{x}'\|_0 = n' + 1$, by a similar argument as above, each arm $i \in \mathcal{A}(\mathbf{x}')$ has to be supported by a unique non-zero variable $x'_{i,\tau_i} > 0$ for some $\tau_i \in \mathbb{N}$, with the exception of a single arm $\iota \in \mathcal{A}(\mathbf{x}')$, which is supported by two non-zero variables $x_{\iota,\tau_{\iota,a}}, x_{\iota,\tau_{\iota,b}} > 0$. Further, since $\|\mathbf{x}'\|_0 = n' + 1$, it has to be that all the constraints in (C.1) and (C.2) are met with equality. Therefore, for any arm $i \in \mathcal{A}(\mathbf{x}') \setminus \{\iota\}$, there exists a unique delay τ_i and variable x'_{i,τ_i} , such that $x'_{i,\tau_i} = 1/\tau_i$. Again, the solution \mathbf{x} of (\mathbf{LP}_k) in this case is almost-delay-feasible, and ι corresponds to the irregular arm.

Since delay-feasible solutions subsume almost-delay-feasible by definition, it follows that, in any case, \mathbf{x} is almost-delay-feasible. \square

C Improved Approximation Guarantees for Planning: Omitted Proofs

Lemma 4.2. *For any vectors $\mathbf{w} \in [0, \infty)^n$ and $\mathbf{y} \in [0, 1]^n$, let $f_{\mathbf{w},k}^+(\mathbf{y})$ be the concave closure of the weighted rank function of the rank- k uniform matroid with weights \mathbf{w} , evaluated at \mathbf{y} . Then, if \mathbf{y} satisfies $\|\mathbf{y}\|_1 \leq k$, it is the case that $f_{\mathbf{w},k}^+(\mathbf{y}) \geq \sum_{i \in [n]} w_i y_i$.*

Proof. Let us consider the intersection of the n -dimensional hypercube $[0, 1]^n$ with the halfspace $H = \{\mathbf{z} \in \mathbb{R}^n \mid \sum_{i \in [n]} z_i \leq k\}$. It is easy to verify that every vertex \mathbf{v} of the resulting polytope lies in $\{0, 1\}^n$ and satisfies $\|\mathbf{v}\|_0 \leq k$. Then, since $\mathbf{y} \in [0, 1]^n \cap H$, by standard arguments in polyhedral combinatorics, \mathbf{y} can be expressed as a convex combination of characteristic vectors of vertices of $[0, 1]^n \cap H$, each corresponding to a subset of $[n]$ of cardinality at most k . This convex combination induces a probability distribution $\mathcal{D}'(\mathbf{y})$ over subsets of cardinality at most k with marginals \mathbf{y} . Therefore, by definition of $f_{\mathbf{w},k}^+$, we have

$$f_{\mathbf{w},k}^+(\mathbf{y}) = \sup_{\mathcal{D}(\mathbf{y})} \mathbb{E}_{S \sim \mathcal{D}(\mathbf{y})} [f_{\mathbf{w},k}(S)] \geq \mathbb{E}_{S \sim \mathcal{D}'(\mathbf{y})} [f_{\mathbf{w},k}(S)] = \mathbb{E}_{S \sim \mathcal{D}'(\mathbf{y})} [\mathbf{w}(S)],$$

where the last equality follows by the fact that $\mathcal{D}'(\mathbf{y})$ is only supported by sets of cardinality at most k .

Finally, using the fact that $\mathcal{D}'(\mathbf{y})$ respects the marginals \mathbf{y} , we can conclude

$$f_{\mathbf{w},k}^+(\mathbf{y}) \geq \mathbb{E}_{S \sim \mathcal{D}'(\mathbf{y})} [\mathbf{w}(S)] = \sum_{i \in [n]} w_i \mathbb{P}_{S \sim \mathcal{D}'(\mathbf{y})} [i \in S] = \sum_{i \in [n]} w_i y_i.$$

\square

Lemma 4.3. *For any arm $i \in \mathcal{A}'$ and at any round t , we have that $\mathbb{P}[i \in C_t \text{ and } \tau_i^* = \tau] = x_{i,\tau}^*$.*

Proof. Fix any arm $i \in \mathcal{A}'$ and round t . By the sampling process of critical delays using the solution \mathbf{x}^* , we have that $\mathbb{P}[\tau_i^* = \tau] = \tau \cdot x_{i,\tau}^*$. Notice that, when arm i is not the irregular arm, by Definition 3.4 this probability is identically one. Since the offset r_i is sampled uniformly at random from $\{0, 1, \dots, \tau_i^* - 1\}$ and $i \in C_t$ if and only if $t \bmod \tau_i^* = r_i$, we get that $\mathbb{P}[i \in C_t \mid \tau_i^* = \tau] = 1/\tau$. Hence, we can conclude that

$$\mathbb{P}[i \in C_t \text{ and } \tau_i^* = \tau] = \mathbb{P}[\tau_i^* = \tau] \cdot \mathbb{P}[i \in C_t \mid \tau_i^* = \tau] = (\tau \cdot x_{i,\tau}^*) \cdot \frac{1}{\tau} = x_{i,\tau}^*.$$

\square

Lemma 4.4. *Let $f : 2^{[n]} \rightarrow \mathbb{R}$ be a submodular function over a ground set of n elements, where $n \geq 2$. Let $\mathcal{I}(\mathbf{y})$ be a distribution over $2^{[n]}$, where each element i is added in set $S \sim \mathcal{I}(\mathbf{y})$, independently, with probability y_i . Further, let $\mathcal{D}(\mathbf{y})$ be a distribution over $2^{[n]}$, where each element $i \in [n] \setminus \{\iota_a, \iota_b\}$ is added to a set $S \sim \mathcal{D}(\mathbf{y})$, independently, with probability y_i , except for two distinct elements ι_a, ι_b , whose addition to S is mutually exclusive but with the same marginals y_{ι_a} and y_{ι_b} , respectively. In this case, we have*

$$\mathbb{E}_{S \sim \mathcal{D}(\mathbf{y})} [f(S)] \geq \mathbb{E}_{S \sim \mathcal{I}(\mathbf{y})} [f(S)].$$

Proof. By definition of the distribution $\mathcal{D}(\mathbf{y})$, since the events $\{\iota_a \in S\}$ and $\{\iota_b \in S\}$ are mutually exclusive, we have that $\mathbb{P}_{S \sim \mathcal{D}(\mathbf{y})} [\{\iota_a \in S\} \cap \{\iota_b \notin S\}] = \mathbb{P}_{S \sim \mathcal{D}(\mathbf{y})} [\iota_a \in S] = y_{\iota_a}$ and $\mathbb{P}_{S \sim \mathcal{D}(\mathbf{y})} [\{\iota_a \notin S\} \cap \{\iota_b \in S\}] = \mathbb{P}_{S \sim \mathcal{D}(\mathbf{y})} [\iota_b \in S] = y_{\iota_b}$. Further, we have that $\mathbb{P}_{S \sim \mathcal{D}(\mathbf{y})} [\{\iota_a \notin S\} \cap \{\iota_b \notin S\}] = 1 - \mathbb{P}_{S \sim \mathcal{D}(\mathbf{y})} [\{\iota_a \in S\} \cup \{\iota_b \in S\}] = 1 - y_{\iota_a} - y_{\iota_b}$.

Using the above, we can rewrite the LHS of the desired inequality as

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}(\mathbf{y})} [f(S)] &= \\ &\sum_{S' \subseteq [n] \setminus \{\iota_a, \iota_b\}} \prod_{i \in S'} y_i \prod_{\substack{i \notin S' \\ i \neq \iota_a, \iota_b}} (1 - y_i) \left((1 - y_{\iota_a} - y_{\iota_b}) f(S') + y_{\iota_a} f(S' + \iota_a) + y_{\iota_b} f(S' + \iota_b) \right). \end{aligned}$$

Now, for any fixed set $S' \subseteq [n] \setminus \{\iota_a, \iota_b\}$, the following identities can be easily verified:

$$\begin{aligned} (1 - y_{\iota_a} - y_{\iota_b}) f(S') &= (1 - y_{\iota_a})(1 - y_{\iota_b}) f(S') - y_{\iota_a} y_{\iota_b} f(S') \\ y_{\iota_a} f(S' + \iota_a) &= y_{\iota_a} (1 - y_{\iota_b}) f(S' + \iota_a) + y_{\iota_a} y_{\iota_b} f(S' + \iota_a) \\ y_{\iota_b} f(S' + \iota_b) &= y_{\iota_b} (1 - y_{\iota_a}) f(S' + \iota_b) + y_{\iota_a} y_{\iota_b} f(S' + \iota_b). \end{aligned}$$

By combining the above identities, we get that

$$\begin{aligned} &(1 - y_{\iota_a} - y_{\iota_b}) f(S') + y_{\iota_a} f(S' + \iota_a) + y_{\iota_b} f(S' + \iota_b) \\ &= (1 - y_{\iota_a})(1 - y_{\iota_b}) f(S') + y_{\iota_a} (1 - y_{\iota_b}) f(S' + \iota_a) + y_{\iota_b} (1 - y_{\iota_a}) f(S' + \iota_b) \\ &\quad + y_{\iota_a} y_{\iota_b} (f(S' + \iota_a) + f(S' + \iota_b) - f(S')) \\ &\geq (1 - y_{\iota_a})(1 - y_{\iota_b}) f(S') + y_{\iota_a} (1 - y_{\iota_b}) f(S' + \iota_a) + y_{\iota_b} (1 - y_{\iota_a}) f(S' + \iota_b) + y_{\iota_a} y_{\iota_b} f(S' + \iota_a + \iota_b), \end{aligned}$$

where, in the inequality, we use that $f(S') + f(S' + \iota_a + \iota_b) \leq f(S' + \iota_a) + f(S' + \iota_b)$, by submodularity of f .

By applying the above inequality for each $S' \subseteq [n] \setminus \{\iota_a, \iota_b\}$, we can conclude that

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}(\mathbf{y})} [f(S)] &= \sum_{S' \subseteq [n] \setminus \{\iota_a, \iota_b\}} \prod_{i \in S'} y_i \prod_{\substack{i \notin S' \\ i \neq \iota_a, \iota_b}} (1 - y_i) \left((1 - y_{\iota_a} - y_{\iota_b}) f(S') + y_{\iota_a} f(S' + \iota_a) + y_{\iota_b} f(S' + \iota_b) \right) \\ &\geq \sum_{S' \subseteq [n] \setminus \{\iota_a, \iota_b\}} \prod_{i \in S'} y_i \prod_{\substack{i \notin S' \\ i \neq \iota_a, \iota_b}} (1 - y_i) \left((1 - y_{\iota_a})(1 - y_{\iota_b}) f(S') + y_{\iota_a} (1 - y_{\iota_b}) f(S' + \iota_a) \right. \\ &\quad \left. + y_{\iota_b} (1 - y_{\iota_a}) f(S' + \iota_b) + y_{\iota_a} y_{\iota_b} f(S' + \iota_a + \iota_b) \right) \\ &= \sum_{S \subseteq [n]} \prod_{i \in S} y_i \prod_{i \notin S} (1 - y_i) f(S) \\ &= \mathbb{E}_{S \sim \mathcal{I}(\mathbf{y})} [f(S)], \end{aligned}$$

thus proving the lemma. \square

D Online Learning with Sublinear Regret: Omitted Proofs

Lemma 5.1. *Let \hat{p}_i be an estimate of the payoff function p_i for each $i \in \mathcal{A}$, such that $|\hat{p}_i(\tau) - p_i(\tau)| \leq \epsilon, \forall \tau \in \mathbb{N}$ for some $\epsilon \in (0, 1)$. The expected payoff collected in T rounds by Algorithm 1, when it uses the estimates instead of the actual payoff functions, is lower bounded by*

$$\gamma_k \text{OPT}(T) - \mathcal{O}(k \cdot \epsilon \cdot T + k \cdot \tau^{\max}).$$

Proof. Let \mathbf{x}^* and $\hat{\mathbf{x}}$ be the optimal solution of (\mathbf{LP}_k) in the case where the payoff functions in the objective are given by $\{p_i\}_{i \in \mathcal{A}}$ and $\{\hat{p}_i\}_{i \in \mathcal{A}}$, respectively. Let also V^* and \hat{V} be the optimal value of (\mathbf{LP}_k) in each of the aforementioned cases. We have

$$\hat{V} = \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} \hat{p}_i(\tau) \cdot \hat{x}_{i,\tau} \geq \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} \hat{p}_i(\tau) \cdot x_{i,\tau}^* \geq \sum_{i \in \mathcal{A}} \sum_{\tau \in \mathbb{N}} (p_i(\tau) - \epsilon) \cdot x_{i,\tau}^* \geq V^* - \epsilon \cdot k,$$

where in the first inequality we use the fact that $\hat{\mathbf{x}}$ is an optimal extreme point solution of (\mathbf{LP}_k) , when the estimates are used in the objective. In the second inequality, we use that $|\hat{p}_i(\tau) - p_i(\tau)| \leq \epsilon, \forall i \in \mathcal{A}, \tau \in \mathbb{N}$, while the third follows by constraint (C.1) of (\mathbf{LP}_k) .

We observe that Algorithm 1 can produce a feasible planning using any given extreme point solution of (\mathbf{LP}_k) (not necessarily the optimal). Let \hat{A}_t be the set of arms played at some round $t \geq \tau^{\max}$ by Algorithm 1 (using the estimates) and let \hat{C}_t be the corresponding set of candidate arms. By construction of the algorithm, it is the case that

$$\hat{A}_t = \operatorname{argmax}_{S \subseteq \hat{C}_t, |S| \leq k} \sum_{i \in S} \hat{p}_i(\hat{\sigma}_{i,t}),$$

where $\hat{\sigma}_{i,t}$ is the actual delay of arm i at time t in a run of Algorithm 1 (using the estimates).

Let $\hat{\tau}_i^*$ be the (sampled) critical delay of arm i in a run of Algorithm 1 (using the estimates). By working along the lines of Theorem 4.5, using monotonicity of $p_i(\tau)$ and the fact that $\hat{\sigma}_{i,t} \geq \hat{\tau}_i^*$, for the expected payoff collected at any round $t \geq \tau^{\max}$ (over the randomness of the payoff realizations), we have

$$\max_{S \subseteq \hat{C}_t, |S| \leq k} \sum_{i \in S} p_i(\hat{\sigma}_{i,t}) \geq \max_{S \subseteq \hat{C}_t, |S| \leq k} \sum_{i \in S} p_i(\hat{\tau}_i^*) \geq \max_{S \subseteq \hat{C}_t, |S| \leq k} \sum_{i \in S} \hat{p}_i(\hat{\tau}_i^*) - \epsilon \cdot k,$$

where in the last inequality we use the fact that \hat{p}_i is ϵ -close to p_i for each arm $i \in \mathcal{A}$.

Let γ_k be the multiplicative approximation ratio of Algorithm 1 for k -RB (given in Theorem 4.5). By the analysis of Theorem 4.5, it follows that

$$\mathbb{E} \left[\max_{S \subseteq \hat{C}_t, |S| \leq k} \sum_{i \in S} \hat{p}_i(\hat{\tau}_i^*) \right] \geq \gamma_k \cdot \hat{V} \geq \gamma_k \cdot V^* - \gamma_k \cdot \epsilon \cdot k \geq \gamma_k \cdot \frac{\operatorname{OPT}(T)}{T} - \gamma_k \cdot \epsilon \cdot k,$$

where the last inequality follows by Lemma 3.1.

Therefore, for any round $t \geq \tau^{\max}$, the expected payoff collected by Algorithm 1 (using the estimates) satisfies

$$\mathbb{E} \left[\max_{S \subseteq \hat{C}_t, |S| \leq k} \sum_{i \in S} p_i(\hat{\sigma}_{i,t}) \right] \geq \gamma_k \cdot \frac{\operatorname{OPT}(T)}{T} - \mathcal{O}(\epsilon \cdot k).$$

The proof follows by summing the above inequalities for every round $t \in [\tau^{\max}, T]$ and using linearity of expectation together with the fact that $\operatorname{OPT}(T)/T \leq k$. \square

Lemma 5.2. *For any $\epsilon, \delta \in (0, 1)$, let $\hat{p}_i(\tau)$ be the empirical average of samples drawn from $p_i(\tau)$ for every arm $i \in \mathcal{A}$ and delay τ , using $m = \frac{1}{2\epsilon^2} \ln\left(\frac{2\tau^{\max}n}{\delta}\right)$ samples, where n is the number of arms. Then, with probability $1 - \delta$, it holds $|\hat{p}_i(\tau) - p_i(\tau)| \leq \epsilon$ for every $i \in \mathcal{A}$ and $\tau \in [\tau^{\max}]$.*

Proof. For any arm $i \in \mathcal{A}$ and delay τ , let $\hat{p}_i(\tau)$ be the empirical average of m samples drawn independently from a distribution of mean $p_i(\tau)$ (and bounded in $[0, 1]$). By Hoeffding inequality [Hoe63], we have that

$$\mathbb{P}[|\hat{p}_i(\tau) - p_i(\tau)| > \epsilon] \leq 2 \exp(-2m\epsilon^2).$$

By setting $m = \frac{1}{2\epsilon^2} \ln\left(\frac{2\tau^{\max}n}{\delta}\right)$, we get that $|\hat{p}_i(\tau) - p_i(\tau)| \leq \epsilon$ with probability $1 - \frac{\delta}{\tau^{\max}n}$.

Therefore, by a simple union bound over all arms $i \in \mathcal{A}$ and delays $\tau \in [\tau^{\max}]$, we get

$$\mathbb{P}[\exists i \in \mathcal{A}, \tau \in [\tau^{\max}] : |\hat{p}_i(\tau) - p_i(\tau)| > \epsilon] \leq \delta,$$

thus, concluding the proof. \square

Fact 5.3. *In any k -RB instance, the number of rounds required to collect m independent samples from $p_i(\tau)$, for each arm $i \in \mathcal{A}$ and delay $\tau \in [\tau^{\max}]$, is upper-bounded by $\mathcal{O}(nm(\tau^{\max})^2/k)$.*

Proof. For any arm $i \in \mathcal{A}$ and delay τ , we can collect m independent samples from the distribution of mean $p_i(\tau)$, by playing arm i for m times – once every τ rounds – which can be achieved in $\mathcal{O}(m\tau)$ time steps. In order to collect m samples for any delay $\tau \in [\tau^{\max}]$ for a fixed arm i , we thus need $\mathcal{O}\left(\sum_{\tau \in [\tau^{\max}]} m\tau\right) = \mathcal{O}(m(\tau^{\max})^2)$ rounds in the worst case.

Assuming that we can play at most one arm per round, in order to collect m samples for each arm/delay pair, we thus need $\mathcal{O}(nm(\tau^{\max})^2)$ rounds in total. In a k -RB instance, where we can play at most $k < n$ arms per round, we can parallelize and, thus, expedite the above sampling process by a factor of k , leading to a total of $\mathcal{O}(nm(\tau^{\max})^2/k)$ exploration rounds. \square

Theorem 5.4. *There exists a learning adaptation of Algorithm 1 in the semi-bandit setting, for which the γ_k -approximate regret for T rounds can be upper-bounded as*

$$\mathcal{O}\left(n^{1/3} \cdot (k \cdot \tau^{\max})^{2/3} \ln^{1/3}(\tau^{\max} n T) \cdot T^{2/3} + k \cdot \tau^{\max}\right).$$

Proof. Let $R(T)$ be the expected payoff collected by our ETC-based bandit algorithm in T rounds. Recall that the γ_k -approximate regret for T rounds is defined as

$$\text{Reg}_{\gamma_k}(T) = \gamma_k \text{OPT}(T) - R(T).$$

For $\epsilon, \delta \in (0, 1)$ (to be specified later), let $m = \frac{1}{2\epsilon^2} \ln(2\tau^{\max} n/\delta)$. By Fact 5.3, we can collect m samples from the distribution of each arm-delay pair in the first $\mathcal{O}\left(\frac{n(\tau^{\max})^2}{\epsilon^2 k} \ln\left(\frac{\tau^{\max} n}{\delta}\right)\right)$ rounds. Since the mean payoffs are all bounded in $[0, 1]$ and we can play at most k arms per round, the regret accumulated for each of the exploration rounds is at most k , which implies that the total regret due to exploration can be upper bounded by $\mathcal{O}\left(\frac{n(\tau^{\max})^2}{\epsilon^2} \ln\left(\frac{\tau^{\max} n}{\delta}\right)\right)$.

Moving on to the exploitation phase, for every arm $i \in \mathcal{A}$ and delay $\tau \in [\tau^{\max}]$, let $\hat{p}_i(\tau)$ be the empirical estimate of $p_i(\tau)$ using m samples collected in the exploration phase. For the rest $T - \mathcal{O}\left(\frac{n(\tau^{\max})^2}{\epsilon^2 k} \ln\left(\frac{\tau^{\max} n}{\delta}\right)\right) \leq T$ rounds, by Lemma 5.2, we get that, with probability $1 - \delta$, it holds $|\hat{p}_i(\tau) - p_i(\tau)| \leq \epsilon$ for every arm i and delay τ . We refer to the above event as a “nice” sampling.

Thus, the total regret accumulated in the exploitation phase in the case where sampling is not nice (which happens with probability δ) is at most $T \cdot k \cdot \delta$. In the case where the sampling is nice, then Lemma 5.1 suggests that for the remaining T' rounds of the exploitation phase, we have

$$R(T') \geq \gamma_k \text{OPT}(T') - \mathcal{O}(k \cdot \epsilon \cdot T' + k \cdot \tau^{\max}),$$

which implies that the total regret accumulated in the exploitation phase can be upper bounded by $\mathcal{O}(k \cdot \epsilon \cdot T + k \cdot \tau^{\max} + T \cdot k \cdot \delta)$. By the above analysis, it follows that for any $\epsilon, \delta \in (0, 1)$, the total regret of our bandit algorithm satisfies

$$\text{Reg}_{\gamma_k}(T) \leq \underbrace{\mathcal{O}\left(\frac{n(\tau^{\max})^2}{\epsilon^2} \ln\left(\frac{\tau^{\max} n}{\delta}\right)\right)}_{\text{Exploration phase}} + \underbrace{\mathcal{O}\left(k \cdot \epsilon \cdot T + k \cdot \tau^{\max} + T \cdot k \cdot \delta\right)}_{\text{Exploitation phase}}.$$

By setting $\epsilon = \mathcal{O}\left(\sqrt[3]{\frac{n(\tau^{\max})^2 \ln(\tau^{\max} n T)}{k T}}\right)$ and $\delta = \frac{1}{T}$, we can upper bound the regret as

$$\text{Reg}_{\gamma_k}(T) \leq \mathcal{O}\left(n^{1/3} \cdot (k \cdot \tau^{\max})^{2/3} \ln^{1/3}(\tau^{\max} n T) \cdot T^{2/3} + k \cdot \tau^{\max}\right),$$

thus proving that our algorithm achieves a sublinear (in the time horizon) regret guarantee. \square

Societal Impact and Limitations

This is a theoretical work and, thus, any (positive or negative) societal impact depends on the application. A possible limitation of our work is the assumption that an upper bound on the recovery times of all arms is known a priori. Notice, however, this assumption (also made in the state-of-the-art prior work [SLZZ21]) is natural for settings where the payoff functions have a non-parametric form.