
A Last Switch Dependent Analysis of Satiation and Seasonality in Bandits

Pierre Laforgue¹

Giulia Clerici¹

Nicolò Cesa-Bianchi¹

Ran Gilad-Bachrach^{2,3,4}

¹ Dept. of Computer Science, Università degli Studi di Milano, Milan, Italy

² Department of Bio-Medical Engineering, Tel-Aviv University, Israel

³ Edmond J. Safra Center for Bioinformatics, ⁴ Sagol School of Neuroscience

Abstract

Motivated by the fact that humans like some level of unpredictability or novelty, and might therefore get quickly bored when interacting with a stationary policy, we introduce a novel non-stationary bandit problem, where the expected reward of an arm is fully determined by the time elapsed since the arm last took part in a switch of actions. Our model generalizes previous notions of delay-dependent rewards, and also relaxes most assumptions on the reward function. This enables the modeling of phenomena such as progressive satiation and periodic behaviours. Building upon the Combinatorial Semi-Bandits (CSB) framework, we design an algorithm and prove a bound on its regret with respect to the optimal non-stationary policy (which is NP-hard to compute). Similarly to previous works, our regret analysis is based on defining and solving an appropriate trade-off between approximation and estimation. Preliminary experiments confirm the superiority of our algorithm over both the oracle greedy approach and a vanilla CSB solver.

1 INTRODUCTION

As the range of applications of multi-armed bandits increases, new algorithms able to deal with a wider range of phenomena must be designed and analyzed. When interacting with humans, one can often observe behaviours such as satiation and seasonality, that are

hardly captured by stationary policies. Satiation typically occurs when making recommendations, e.g., for books, music, or movies (Kunaver and Požrl, 2017; Leqı et al., 2020). Stationary policies, which learn to recommend a single cuisine or musical genre, fail to capture the human desire for novelty (Dimitrijević et al., 1972; Kovacs et al., 2018). Another widespread type of nonstationary behaviour is seasonality. For example, certain music genres may be preferred during working hours, whereas chill playlists could be more appreciated on evenings or weekends (Schedl et al., 2018). Nonstationary behaviours also naturally occur in the medical domain. Consider people suffering from various mental disorders, such as post-traumatic stress disorder, depression, or anxiety. Some remedies include micro-interventions (mindfulness, positive psychology exercises), Cognitive Behavioural Therapy, or Dialectical Behavioral Therapy (Meinlschmidt et al., 2016; Owen et al., 2018; Schroeder et al., 2018; Fuller-Tyszkiewicz et al., 2019). However, studies have shown that a diminishing return effect exists when the same intervention is applied repeatedly (Paredes et al., 2014), motivating the modelling of satiation effects.

In this work, we introduce the first bandit model which captures nonstationary phenomena that include satiation and seasonal effects. Unlike previous works, where the state structure is fully determined by how long ago each arm was last played, our notion of state keeps track of the time elapsed since an action took part in a switch (i.e., when the arm being pulled changes). This allows to modulate satiation effects at a level of detail not within reach of previous models. In addition, our model drops many assumptions on the shape of the expected reward function (such as concavity, Lipschitz continuity, monotonicity), while only retaining boundedness which is plausible in human interactions. Dispensing with monotonicity enables the modelling of periodicity, which is key to capture seasonality.

Technical contributions. The backbone of our analysis goes along the lines of previous works: (1) we show that computing the optimal policy (with value OPT) in our setting is NP-hard; (2) we prove that OPT is well approximated by a certain class of simple policies; (3) we show how to learn the best policy in the approximating class. Also similarly to previous works, we initially consider cyclic policies (with bounded block length) as approximating class. One of the main technical hurdles when learning a cyclic policy in a nonstationary setting is the calibration problem. Namely, the expected reward of the cycle block (i.e., the sequence of arm pulls that is being repeated in the cycle) depends on the current state of all the arms appearing in the block. A simple, yet impractical solution, is to force the learning algorithm to play every block twice in a row, where the first play of a block calibrates the reward estimates computed in the second play. Our solution is radically different: rather than forcing the algorithm to artificially play additional arms, we calibrate an arm by referring to the first time the arm is pulled in the block. These calibration pulls are not used to compute reward estimates, because their reward depends on the last time the arms took part in a switch among the blocks previously played by the learning algorithm. By ignoring calibration pulls, we underestimate the block reward by at most K (the number of arms), which in typical instances is much smaller than the block length.

Our calibration approach reduces the problem of learning the best calibrated block to that of solving an instance of Combinatorial Semi-Bandits (CSB), which we can solve using algorithms whose regret is well understood. As both the regret bound and the approximation factor for OPT depend on the block length d , we can then choose d to optimize our performance, thus obtaining a regret bound against OPT of order $KT^{3/4}$, ignoring log factors. Note that this does imply polynomial-time convergence to OPT, as solving the decoding problem (an intermediate step in the regret minimization procedure) for our instance of CSB requires solving an integer linear program, which is NP-hard in general. The decoding problem amounts to compute the reward-maximizing block given the current estimates of the state-value function (which we represent with a $K \times d$ table of reward estimates). In practice, we can approximately solve it through a branch-and-bound approach using a LP relaxation to bound the value of the objective. We implement an approximate solver based on this approach and run an experiment on a simple instance of LSD where our algorithm is seen to outperform two natural baselines.

Related works. Previous works addressed various extensions of bandits where rewards depend on past

pulls. These include rested models (Gittins, 1979), such as Rotting Bandits (Bouneffouf and Féraud, 2016; Heidari et al., 2016; Cortes et al., 2017; Levine et al., 2017; Warlop et al., 2018; Seznec et al., 2019) and restless models (Whittle, 1988; Tekin and Liu, 2012). Similarly to Cella and Cesa-Bianchi (2020), our framework is neither rested nor restless, because the reward of arms that are pulled changes differently from the reward of arms that are not pulled. Unlike Cella and Cesa-Bianchi (2020), we do not assume a specific form for the reward function and so we can model satiation and seasonalities. Kleinberg and Immorlica (2018) also look at a similar model, but —unlike us— they assume the reward functions to be concave, increasing, and Lipschitz. Exploiting concavity (but not Lipschitzness) they prove that for every $0 < \varepsilon < 1$ there exists a periodic schedule of length $T \geq K/\varepsilon$ whose asymptotic average reward is at least $(1 - \varepsilon)\text{OPT}$. Pike-Burke and Grunewalder (2019) also assume that the expected reward is a function of the time since the arm is last pulled, but consider reward functions drawn from a Gaussian Process with known kernel. Another relevant work is Simchi-Levi et al. (2021), where the authors explore the possibility of pulling and collecting delay-dependent rewards from more than one arm at each time step. Finally, note that Basu et al. (2019) investigate an interesting related variant where an arm becomes unavailable for a certain amount of time after each pull. Although semantically close, we highlight that bandits with delayed feedback (Pike-Burke et al., 2018) are unrelated to our setting.

2 LSD BANDITS

Let \mathcal{A} be an action set, with cardinality $K = |\mathcal{A}|$. When necessary, individual actions are distinguished using superscripts, $\mathcal{A} = \{a^{(1)}, \dots, a^{(K)}\}$, whereas subscripts are used for the temporal indexing, so that a_t refers to the action played at time step t . In our setting, the expected reward of an arm a is assumed to be fully determined by the last time t it took part in a switch. That is, the last t when the tuple (a_{t-1}, a_t) contains a exactly once. Therefore, every arm $a \in \mathcal{A}$ is endowed with a state τ_a (τ_a actually depends on t , and we should use $\tau_a(t)$, but we drop this dependence whenever it is understood from the context) such that:

$$\tau_a(0) = 1 \quad \text{and} \quad \tau_a(t+1) = \delta_a(\tau_a(t), a_t), \quad (1)$$

where δ_a is the transition function given by:

$$\delta_a(\tau, a') = \begin{cases} \tau - 1 & \text{if } a' = a, \quad \tau \leq 0 \\ -1 & \text{if } a' = a, \quad \tau \geq 0 \\ 1 & \text{if } a' \neq a, \quad \tau \leq 0 \\ \tau + 1 & \text{if } a' \neq a, \quad \tau \geq 0 \end{cases} \quad (2)$$




Figure 1: Transition mechanisms (reduced to 6 and 4 states), and examples of expected reward function μ_a for LSD Bandits (left) and models with delay-dependent rewards (Kleinberg and Immorlica, 2018, e.g.) (right).

The transition graph (reduced to 6 states) is pictured in Figure 1 (bottom left). For each arm $a \in \mathcal{A}$, the state τ_a keeps track of the last time a switch of actions involving action a . A positive τ_a means that the last switch involving arm a was $[a, \text{not } a]$, and occurred τ_a time steps ago. In other words, arm a has not been played for the last τ_a rounds. A negative τ_a means that the last switch involving arm a was $[\text{not } a, a]$, and occurred $|\tau_a|$ time steps ago. In other words, arm a has been consistently played for the last $|\tau_a|$ rounds. We can now define Last Switch Dependent (LSD) Bandits, in which the expected reward of an arm a only depends on its last switch state τ_a .

Definition 1 (LSD Bandit). *A stochastic bandit with action set \mathcal{A} is a LSD bandit if for every action $a \in \mathcal{A}$ there exists an (unknown) function $\mu_a: \mathbb{Z} \rightarrow [0, 1]$, nondecreasing on \mathbb{Z}^- , such that the expected reward of arm a is given by $\mu_a(\tau_a)$, where τ_a is the last switch state of a , as defined in (1) and (2).*

Note that a LSD bandit is fully characterized by its reward functions $(\mu_a)_{a \in \mathcal{A}}$ and the noise distribution. In the rest of the paper, we also use the following shortcut notation: $\mu_i = \mu_{a^{(i)}}$, $\tau_i = \tau_{a^{(i)}}$, and $\boldsymbol{\tau}(t) = (\tau_a(t))_{a \in \mathcal{A}}$. Given a horizon T , the learner interacts with a LSD bandit as follows. First, the arm states are initialized to 1^a . Then, for all time steps t from 1 to T :

1. the learner chooses an action $a_t \in \mathcal{A}$,
2. the learner obtains a stochastic reward X_t with expected value $r_t := \mu_{a_t}(\tau_{a_t}(t))$,
3. states update: $\forall a \in \mathcal{A}, \tau_a(t+1) = \delta_a(\tau_a(t), a_t)$.

The goal is to maximize the expected sum of rewards,

^aIn case of a nondecreasing μ_i on \mathbb{Z}^+ , the initialization $\tau_i(0) = +\infty$ would be the most sensible. In the absence of monotonicity however, all positive states become equivalent choices for the initialization, $\tau_i(0) = 1$ is one of them.

and the performance is measured through the regret

$$R_T = \sum_{t=1}^T \mu_{a_t^*}(\tau_{a_t^*}(t)) - \mathbb{E} \left[\sum_{t=1}^T \mu_{a_t}(\tau_{a_t}(t)) \right],$$

where $[a_1^* \dots a_T^*]$ is the optimal sequence of actions, i.e., the sequence maximizing the expected sum of rewards over steps $1, \dots, T$. LSD Bandits generalize bandits with delay-dependent rewards (Kleinberg and Immorlica, 2018; Pike-Burke and Grunewalder, 2019; Cella and Cesa-Bianchi, 2020) in two ways.

First, we emphasize that the notion of last switch generalizes the notion of delay. Whereas both definitions coincide on the positive values, last switches also allow to model progressive satiation, while delays cannot. Indeed, in the case where the same arm is pulled consecutively, the last switch keeps updating towards the negative values, see (2), thus allowing to associate different rewards with different numbers of consecutive pulls. On the contrary, no matter if the arm is pulled two or ten times in a row, the delay remains equal to 0 (or -1 , depending on the convention), as only the fact that the same arm was pulled in the previous round matters. In other words, a bandit with delay-dependent rewards is a special case of LSD, where the reward function μ_a is constant over the negative τ 's, while for a general LSD bandit it is only assumed to be nondecreasing—see Figure 1, top. Assuming μ_a to be nondecreasing on \mathbb{Z}^- can be interpreted as a diminishing return requirement: the more the same arm is consecutively pulled, the smaller its reward gets (before resetting to the “default” value $\mu_a(1)$ when the series is interrupted). This is natural in most scenarios, and key to derive nonvacuous approximations (Table 1).

The second generalization regards the assumptions made on the values of μ_a on \mathbb{Z}^+ . In LSD bandits, μ_a is only supposed to be bounded. This assumption encompasses the framework of Cella and Cesa-Bianchi

(2020), who considers a parametric class of bounded expected reward functions that are increasing and have bounded memory (their maximum is attained after a certain arm-dependent delay). The Recharging Bandit model of Kleinberg and Immorlica (2018), instead, assumes expected reward functions that are concave, increasing, Lipschitz continuous, but not bounded, making their framework and ours incomparable. We nonetheless highlight several limitations of their setting: (1) concavity excludes simple examples, such as $\mu_a(\tau) = \mathbb{I}\{\tau \geq t_a\}$ for some predefined threshold t_a , (2) concavity prevents from modelling series of arm pulls, as a negative branch would force rewards to decrease unbounded, and (3) monotonicity prevents from modelling seasonality, see Figure 1, top. We finally highlight that relaxing the monotonicity of μ_a on \mathbb{Z}^+ unlocks many more applications than just modelling seasonality. Note that seasonality is understood with respect to the sequence of arm pulls. This matches the absolute time seasonality as soon as the learner is asked to make predictions at regular intervals.

Relaxing the concavity assumption has also important consequences on the difficulty of the problems which can be considered. For concave reward functions the oracle greedy policy provides a $1/2$ -approximation of the optimal policy (Kleinberg and Immorlica, 2018), yet Example 1 shows that oracle greedy can be arbitrarily bad in our setting, highlighting that LSD Bandits can be more difficult than Recharging Bandits.

Example 1. Consider the LSD bandit defined by the following reward functions: $\mu_1(\tau) = \epsilon + (1-\epsilon)\mathbb{I}\{\tau \geq 1\}$, and $\mu_2(\tau) = 0 \forall \tau$. An oracle greedy strategy, which pulls at each time step the arm with highest expected reward (assuming the knowledge of the μ_a) would always pull arm $a^{(1)}$ obtaining a reward of $1 + (T-1)\epsilon$ after T rounds. Instead, the optimal policy alternates between arms $a^{(1)}$ and $a^{(2)}$ and gets a $T/2$ overall reward. By making ϵ arbitrary close to 0, we can thus make oracle greedy arbitrarily bad. We conclude with a remark on why this example would be ruled out by concavity. In Kleinberg and Immorlica (2018), concavity is actually defined with respect to the origin, such that the reward obtained in case of consecutive pulls (here it is ϵ) is considered as an increment from 0. Concavity then prevents the next increments from being bigger, while here it is equal to $1 - \epsilon$, which is greater than ϵ as soon as $\epsilon < 1/2$. And for $\epsilon \geq 1/2$, one can note that oracle greedy is indeed a $1/2$ approximation of the optimal strategy, as revealed by the above computations.

2.1 Hardness Results

Going beyond Example 1, we emphasize the difficulty of solving LSD bandits by showing several hardness results. All missing proofs are deferred to Appendix A.

Proposition 1. *Computing the optimal policy for LSD bandits is NP-hard.*

The proof builds on (Basu et al., 2019, Theorem 3.1), and is detailed in Appendix A.1. Since computing the global optimal policy is NP-hard even with full knowledge of the problem instance, a common approach then consists in learning the optimal policy within a simpler class of approximating policies. For LSD bandits, a natural approximating class is the set of policies with cyclic play sequences. A play sequence a_1, a_2, \dots is cyclic if there exist t_0 and $d > 0$ such that $\forall t \geq t_0$ it holds that $a_{t+d} = a_t$. Cyclic policies can be shown to be optimal for 2-armed LSD bandits (Lemma 2), and give good constant factor approximations in general (Proposition 3).

Lemma 2. *For a LSD bandit with two arms, any deterministic policy induces a sequence of pulls which is cyclic from a certain time step onward. This holds in particular for the optimal deterministic policy.*

With 3 arms however, we can exhibit optimal policies which do not induce cyclic sequences. Note that this does not imply that no optimal policy is cyclic. Answering this question is however quite involved, and deferred to future work. In Proposition 3, we show a slightly weaker result, namely that cyclic policies can almost reach the best average reward, i.e., up to a term which is inversely proportional to the cycle length.

Proposition 3. *Let $(\mu_i)_{i=1}^K$ be an LSD bandit with K arms and constant expected rewards on \mathbb{Z}^- . Let $T \geq 0$ be the horizon, and $d \geq 0$ that divides T . Then, there exists a cyclic policy π with cycle length d (and $t_0 = 1$) such that*

$$\frac{1}{T} \sum_{t=1}^T r_\pi(t) \geq \frac{1}{T} \sum_{t=1}^T r^*(t) - \frac{K}{d} \quad (3)$$

where $r_\pi(t)$ and $r^*(t)$ are the expected rewards obtained at time step t by π and the optimal policy, respectively. When the μ_i are not constant on \mathbb{Z}^- , (3) holds with $K+2$ instead of K in the right-hand side. Note also that (3) requires $d \geq K$ (respectively $d \geq K+2$) to be nonvacuous, as we have: $0 \leq r^*(t) \leq 1$.

Proposition 3 is proved at the end of the section for constant μ_i on \mathbb{Z}^- . Thus, looking for the best cyclic policy simplifies the problem (as the search space is reduced), while maintaining a control on the approximation error via the cycle length. Unfortunately, this strategy is not viable either: it can be shown that finding the optimal cyclic policy for LSD bandits remains NP-hard, even when arms are totally ordered.

Proposition 4. *Finding the optimal cyclic policy for LSD bandit problems is NP-hard, even with separable reward functions of the form $\mu_i(\tau) = \mu_i^0 \cdot \gamma(\tau)$.*

μ_a	— on \mathbb{Z}^-	\nearrow on \mathbb{Z}^-	\sim on \mathbb{Z}^-
\nearrow on \mathbb{Z}^+	$-(K-1)$	$-(K+1)$	$-d$
\sim on \mathbb{Z}^+	$-K$	$-(K+2)$	$-d$

Table 1: Approximation errors by block of size d when μ_a is constant (—), nondecreasing (\nearrow), or non-monotone (\sim) on \mathbb{Z}^- and \mathbb{Z}^+ . The grey cell represents previous works. Our setting covers the blue ones. The red cells indicate vacuous approximations.

We conclude this section by proving Proposition 3. For $B = [a_1 \dots a_d]$, let $r(B | \tau_{\text{init}}) = \sum_{t=1}^d \mu_{a_t}(\tau_{a_t}(t))$ be the sum of the expected rewards obtained by playing the actions in B when τ is initialized to $\tau(1) = \tau_{\text{init}}$. Proposition 3 is essentially derived from the following key lemma, which exploits the boundedness assumption to relate $r(B | \tau)$ for different initial states.

Lemma 5. *Let $(\mu_i)_{i=1}^K$ be an LSD bandit with K arms and constant expected rewards on \mathbb{Z}^- . Then, for all block B and any initial states $\tau_{\text{init}}, \tau'_{\text{init}} \in \mathbb{Z}^K$, we have*

$$r(B | \tau'_{\text{init}}) \geq r(B | \tau_{\text{init}}) - K. \quad (4)$$

If the μ_i are not constant on \mathbb{Z}^- , a slight modification of B in the left-hand side is required to maintain a similar bound. Formally, for any block B of length d , there exists a block B' of length d such that for any initial states $\tau_{\text{init}}, \tau'_{\text{init}}$, we have

$$r(B' | \tau'_{\text{init}}) \geq r(B | \tau_{\text{init}}) - (K+2). \quad (5)$$

Proof. For simplicity, here we only prove the special case where the μ_i are constant on \mathbb{Z}^- . First note that after the first play of an arm a , the switch $[a, \text{not } a]$ occurs, and the arm goes to $\tau_a = 1$, independently of its initial state. Now, what happens during the first play of a ? Let t_a be the time step at which arm a is played for the first time. Even if there are many consecutive pulls of a , the rewards collected for τ_{init} and τ'_{init} are $\mu_a(\tau'_{\text{init},a} + t_a - 1), \mu_a(-1), \dots, \mu_a(-1)$ and $\mu_a(\tau_{\text{init},a} + t_a - 1), \mu_a(-1), \dots, \mu_a(-1)$ because the μ_i are constant on \mathbb{Z}^- . Thus, the only difference is $\mu_a(\tau'_{\text{init},a} + t_a - 1) - \mu_a(\tau_{\text{init},a} + t_a - 1) \leq 1$. Over the K arms, the total difference cannot exceed K , giving (4). Note that in full generality, K could be replaced by the number of different arms played in B . When the μ_i are not constant on \mathbb{Z}^- , a finer comparison of the first pulls is required. \square

The approximation errors by block of size d that can be achieved depending on the assumptions made on μ_a are summarized in Table 1. Relaxing the assumptions on \mathbb{Z}^+ comes with a low price, while the monotonicity on \mathbb{Z}^- is critical. We now provide a proof sketch of Proposition 3 when the μ_i are constant on \mathbb{Z}^- , and show an example where our analysis is essentially tight.

Proof sketch of Proposition 3. Note that there exists a time step $t_0 \leq T - d + 1$ such that $(1/d) \sum_{t=t_0}^{t_0+d-1} r^*(t) \geq (1/T) \sum_{t=1}^T r^*(t)$. Let $B^* = [a_{t_0}^* \dots a_{t_0+d-1}^*]$, and $\tau^*(t)$ be the sequence of states generated by the optimal policy, such that $\sum_{t=t_0}^{t_0+d-1} r^*(t) = r(B^* | \tau^*(t_0))$. The policy π consists in playing repeatedly block B^* . However, except for the first d steps, B^* is played by π with an initial state τ_{B^*} , i.e., the state reached by the system after a play of B^* , which might be different from $\tau^*(t_0)$. Applying Lemma 5, and assuming for simplicity that states were initialized in π to τ_{B^*} , we obtain

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T r_\pi(t) &= \frac{r(B^* | \tau_{B^*})}{d} \geq \frac{r(B^* | \tau^*(t_0)) - K}{d} \\ &\geq \frac{1}{T} \sum_{t=1}^T r^*(t) - \frac{K}{d}. \end{aligned}$$

If the μ_i are not constant on \mathbb{Z}^- , the repeated block is built using the second part of Lemma 5. \square

Example 2. *Consider the LSD bandit defined by the following reward functions*

$$\mu_i(\tau) = \mathbb{I}\{\tau \geq K - 1\} \quad \forall i \leq K.$$

The optimal policy consists in repeating the block $[a^{(1)} \dots a^{(K)}]$ and obtains an average reward of 1. Let $d = 2K - 1$, such that (up to permutations) we have $B^* = [a^{(1)} \dots a^{(K)} a^{(1)} \dots a^{(K-1)}]$. The average reward of B^* among the optimal sequence is 1, which is equal to the global average reward. Now, it is easy to check that playing repeatedly B^* yields an average reward of $K/(2K-1)$. On the other hand, the lower bound given by Proposition 3 is $1 - K/d = (K-1)/(2K-1)$, which matches the average reward as $K \rightarrow +\infty$.

3 PROPOSED APPROACH

We now introduce our approach to solve LSD bandits. It is based upon the CombUCB1 algorithm introduced in (Gai et al., 2012) to solve Combinatorial Semi-Bandits. Our adaptation is designed to cancel the interference created by the blocks previously played, and enjoys an approximation-estimation tradeoff which can be solved by an appropriate choice of the block size. In order to keep the exposition concise, we only focus here on the case where the μ_i are constant on \mathbb{Z}^- . Results for the general case can be found in the Supplementary Material.

Approximation. A general idea that we can retain from Section 2.1, and from Proposition 3 in particular, is that approximating the optimal sequence using a series of smaller blocks of length $d < T$ is reasonable.

However, the block B^* exhibited during the proof of Proposition 3 is seemingly impossible to estimate, as it requires the computation of the optimal sequence first. Another key challenge in LSD bandits is to handle the impact of the state: the same block played in different states may have different outcomes, making it difficult to identify “good blocks”. One natural way to tackle this problem is to play every block twice: the rewards obtained during the second play are then fully representative of the block value, if repeated as a cycle. Let τ_B be the state reached by the system after a play of block B from initial state $\mathbf{1}$. We can then introduce

$$B_{\text{double}}^* = \operatorname{argmax}_{|B|=d} r(B \mid \tau_B). \quad (6)$$

Playing cyclically B_{double}^* almost benefits from the same approximation guarantee as playing cyclically B^* (Proposition 6). However, solving (6) is nontrivial, as B influences both the reward-generating sequence and the initial state (it is essentially equivalent to compute the optimal cycle, which has been proven NP-hard in Proposition 4). Moreover, using an entire block simply to control the initial state of the system seems quite excessive, especially when d which has to grow to control the approximation error.

Instead, a cheaper way to *calibrate* the arms, i.e., to pull the arms before playing the real block such that the arms are in a controlled state independent of the past, consists of playing $B_\sigma = [a^{(\sigma(1))} \dots a^{(\sigma(K))}]$ for any permutation $\sigma \in \mathfrak{S}_K$. Indeed, even if no reward can be exploited from B_σ , it only takes K pulls to calibrate the system, while pre-playing the block is a $d > K$ long calibration. Let τ_σ be the state reached by the system after a play of block B_σ (note that it only depends on σ), we set

$$B_\sigma^* = \operatorname{argmax}_{|B|=d} r(B \mid \tau_\sigma). \quad (7)$$

A natural idea then consists in playing cyclically the block $[B_\sigma, B_\sigma^*]$ of size $K+d$, and approximation results (independent from σ) can be derived for this approach, see Proposition 6. Still, this strategy is not satisfactory for three reasons: (1) the approximation guarantee is worse than for double plays; (2) playing B_σ might still be highly inefficient: all arms are calibrated, while only those present in B_σ^* would require calibration; (3) in domains like song recommendation, this would amount to regularly play a representative song of each genre.

To remedy this issue, we introduce the following modified version of the block expected reward, which does not take into account the reward obtained if the arm is played for the first time in the block. Formally, for any block $B = [a_1 \dots a_d]$ we define

$$\tilde{r}(B) = \sum_{t=1}^d \mu_{a_t}(\tau_{a_t}(t)) \mathbb{I}\{\exists t_0 < t: a_{t_0} = a_t\}$$

where t is indexed with respect to B . Note that $\tilde{r}(B)$ is independent from the initial state, so that

$$\tilde{B}^* = \operatorname{argmax}_{|B|=d} \tilde{r}(B) \quad (8)$$

is well defined. The rationale behind \tilde{r} is the following: first pulls do not provide reliable rewards because they are influenced by the past. Therefore, we use them as a calibration for the future pulls. So, only the arms which are used are calibrated. Observe also that calibration and exploitation phases might be intertwined, which was impossible with B_σ . But most importantly, the loss incurred by maximizing \tilde{r} and not r is controllable, as we have for any block B and initial state τ_{init} : $r(B \mid \tau_{\text{init}}) - K \leq \tilde{r}(B) \leq r(B \mid \tau_{\text{init}})$. We can now state our complete approximation result.

Proposition 6. *Let $(\mu_i)_{i=1}^K$ be a LSD bandit with K arms and constant expected rewards on \mathbb{Z}^- . Let $T \geq 0$ be the horizon, and $d \geq 0$ that divides T . Let $r_{\text{double}}(t)$ be the expected rewards obtained at time step t by the policy playing cyclically B_{double}^* . We have*

$$\frac{1}{T} \sum_{t=1}^T r_{\text{double}}(t) \geq \left(1 - \frac{d}{T}\right) \left(\frac{1}{T} \sum_{t=1}^T r^*(t) - \frac{K}{d}\right). \quad (9)$$

Let $\sigma \in \mathfrak{S}_K$, and assume that $d + K$ divides T . Let $r_\sigma(t)$ be the expected rewards obtained at time step t by the policy playing cyclically $[B_\sigma, B_\sigma^*]$. We have

$$\frac{1}{T} \sum_{t=1}^T r_\sigma(t) \geq \frac{d}{d+K} \left(\frac{1}{T} \sum_{t=1}^T r^*(t)\right) - \frac{K}{d+K}. \quad (10)$$

Let $\tilde{r}^*(t)$ be the expected rewards obtained at time step t by the policy playing cyclically \tilde{B}^* . We have

$$\frac{1}{T} \sum_{t=1}^T \tilde{r}^*(t) \geq \frac{1}{T} \sum_{t=1}^T r^*(t) - \frac{K}{d}. \quad (11)$$

The proof of Proposition 6 is similar to that of Proposition 3, and essentially combine Lemma 5 with the definitions of B_{double}^* , B_σ , and \tilde{B}^* . Based on these results, using $\tilde{r}(B)$ seems by far the best option. First, it is immediate to check that (11) is tighter than both (9) and (10) as soon as $(1/T) \sum_{t=1}^T r^*(t) \geq K/d$, which is implicitly assumed for the bounds to be nonvacuous. Moreover, what cannot be seen from (9) is that computing a sequence of blocks with small regret against B_{double}^* requires to play each block twice, with no guarantee that the first play will provide any reward, thus dividing (9) by 2. On the contrary, as we see next, we can use the CombUCB1 algorithm of Gai et al. (2012) to estimate \tilde{B}^* with tight regret bounds.

Estimation. A Combinatorial Semi-Bandit (CSB)—see, e.g., (Audibert et al., 2014)—is an online learning problem where, at each time step, the learner has to select a subset of N actions in a universe of $L > N$ base actions, under some combinatorial constraints. The individual rewards for each selected action are then revealed, and the learner receives their sum as total reward^b. The regret is measured against the best subset of N base actions satisfying the constraints. Note for instance that a standard K -armed bandit is a particular instance of CSB, with $N = 1$, $L = K$, and no constraints. We now show that computing a series of blocks with small regret against \tilde{B}^* with respect to \tilde{r} can be reduced to a CSB problem.

In our case $N = d$ (the blocks contain d actions). The universe of base actions is however slightly more intricate to determine. While the analysis of CSB uses the fact that a block can only be filled with a subset of base actions, in our setting the same arm might be used several times in the same block. In addition to that, in CSB base actions have i.i.d. rewards, while the rewards of our arms also depend on the state. Our solution is to consider a universe of Kd^2 base actions: namely a base action is indexed by an arm $i \in \{1, \dots, K\}$, a state $\tau \in \{1, \dots, d\}$, and a position $t \in \{1, \dots, d\}$ in the block. The state coordinate ensures the i.i.d. nature of the rewards, while the time coordinate allows to remove the arm multiplicities, therefore making the map from a block to its representation one-to-one. This is needed to extract valid sequences from solutions to (12), where the combinatorial constraints (ensuring that only subsets deriving from a sequence of pulls can be selected), are made explicit. Note that the structure of the problem plays a critical role here, as we derive an action space of size $L = Kd^2$, as opposed to K^d in the absence of any structure. From a pure estimation point of view, a space of size Kd is enough, as the rewards are independent of the position in the sequence.

CombUCB1 is an algorithm introduced by Gai et al. (2012) to solve CSBs. Its analysis was later refined by Chen et al. (2013). Kveton et al. (2015, Theorem 6) prove a regret bound of order $\mathcal{O}(\sqrt{NLn \log n})$, where n is the number of rounds. This bound can thus be applied to our setting in a black box fashion, with $n = T/d$, $L = d$, and $N = Kd^2$. The resulting algorithm, which we call ISI-CombUCB1 for Initial States Independent CombUCB1, produces a sequence of blocks with small regret against \tilde{B}^* with respect to \tilde{r} . Combining this result with Proposition 6, we can bound the regret of ISI-CombUCB1 with respect to OPT as follows.

^bIn Chen et al. (2013) authors consider more general rewards, but the linear case is sufficient for our application

Theorem 7. *Let $(\mu_i)_{i=1}^K$ be an LSD bandit with K arms and constant expected rewards on \mathbb{Z}^- . Let $T \geq 0$ be the horizon, and choose $d \geq 0$ that divides T . Then ISI-CombUCB1, run with block size d and exploration parameter $\alpha = 1.5$, has regret bounded by*

$$R_T \leq \frac{KT}{d} + 47d\sqrt{KT \log \frac{T}{d}} + \left(\frac{\pi^2}{3} + 1\right) Kd^3.$$

Choosing $d \propto T^{1/4}$, we obtain $R_T = \tilde{\mathcal{O}}(KT^{3/4})$, where $\tilde{\mathcal{O}}$ is neglecting logarithmic factors.

Note that the second claim of Theorem 7 requires a horizon-dependent tuning. One can use the doubling trick to make the algorithm anytime without harming the bound. Regarding the optimality of Theorem 7, we recall that our approximation result in $\mathcal{O}(KT/d)$ is tight, see e.g., Example 2. As for the estimation part, Kveton et al. (2015) proved a lower bound for CombUCB1 that matches the upper bound up to a factor of $\sqrt{\log n}$, where n is the horizon. Instantiating this lower bound to our case, we obtain an overall lower bound of order $\Omega(KT/d + d\sqrt{KT})$. Theorem 7 is thus tight up to a polylogarithmic factor of $\sqrt{\log(T/d)}$.

We now give more details about the implementation of ISI-CombUCB1, which specializes CombUCB1 to the minimization of \tilde{r} . After an initialization step ensuring that each base action is played at least once, CombUCB1 maintains upper confidence bounds (UCBs) for each base action. At each round, the algorithm plays the admissible subset of actions with the biggest sum of UCBs, and updates the UCBs according to the individual rewards obtained. In ISI-CombUCB1, the optimization problem to select the best subset amounts to solve (8), but using the UCBs instead of the true expected rewards. As pointed out during the construction of the universe, in our case some actions share the same rewards, and it is actually enough to maintain only Kd UCBs (one for each arm and each delay). Although Theorem 7 applies to this version of ISI-CombUCB1 as well, the need of taking multiplicities into account makes Kd^2 appear in the bound. This memory-efficient version, which is preferred in practice, is summarized in Algorithm 1. Note that the initialization is such that pulling an uninitialized arm is always better than pulling any combination of initialized arms. As a consequence, the latter lasts at most Kd rounds, like in standard CombUCB1.

The key step in Algorithm 1 consists in computing the block B_t maximizing the current reward estimates. We now formally describe this optimization problem. Let $F \in \{0, 1\}^{K \times d}$, such that $F[i, t] = 1$ if and only if arm i is played for the first time in the block at time step t . Let $Y \in \{0, 1\}^{K \times d \times d}$, such that $Y[i, j, t] = 1$ if and only if arm i is played with delay j at time step t .

Let $Z = (F, Y)$ be the Kd^2 -sized representation we introduced earlier. Note that the column associated to $t = 1$ in Y is filled with zeros, as a pull here is by definition a first pull, and thus encoded in F . This notation of size $Kd(d+1)$ is however more convenient for coherence. The constraints for Z needed to describe a valid sequence of arm pulls are: Action Consistency (AC), i.e., at each time step, one and only one action is played, Unique First Pull (UFP), i.e., there is only one first pull per arm, First Pulls First (FPF), i.e., first pulls must precede any other pull of the same arm, and Time Consistency (TC), i.e., an arm can be pulled with delay j at time step t only if it was pulled at time step $t - j$, and not pulled since. These constraints write (for index limits that make sense)

$$\forall t \quad \sum_i F[i, t] + \sum_{i, j} Y[i, j, t] = 1 \quad (\text{AC})$$

$$\forall i \quad \sum_t F[i, t] \leq 1 \quad (\text{UFP})$$

$$\forall i, t \quad \sum_j Y[i, j, t] \leq \sum_{s=1}^{t-1} F[i, s] \quad (\text{FPF})$$

$$\forall i, j, t \quad Y[i, j, t] \leq F[i, t - j] + \sum_l Y[i, l, t - j] - \sum_s Y[i, j - s, t - s] \quad (\text{TC})$$

The objective function is the sum of the current UCBs for the second actions present in the block, and writes $\sum_{i, j, s} Y[i, j, s] U_t(i, j)$. Noticing that all the relations are linear, we can derive $c_t \in \mathbb{R}^{Kd^2}$, $G \in \mathbb{R}^{Kd^2 + K \times Kd^2}$, $h \in \mathbb{R}^{Kd^2 + K}$, $A \in \mathbb{R}^{d \times Kd^2}$, and $b \in \mathbb{R}^d$, such that the optimization problem writes as

$$\max_{z \in \{0, 1\}^{Kd^2}} c_t^\top z \quad \text{s.t.} \quad \begin{cases} Gz \leq h \\ Az = b \end{cases} \quad (12)$$

where z is a vector version of Z . Note that (12) is an integer linear program, which is NP-hard to solve in general. This is expected, as our approach enjoys a sublinear linear regret against OPT (which is NP-hard to compute), and is therefore bound to be intractable. In [Simchi-Levi et al. \(2021\)](#), a Fully Polynomial-Time Approximation Scheme (FPTAS) is used to address a similar problem, see Lemma 6 therein. However, we highlight that, although similar at first sight, our two ILPs are fundamentally different. While the authors try to select the best K arms at a fixed time step, we aim at selecting an optimal block, that takes into account the evolution of the rewards among time. This time constraint is specific to our problem, and prevents from using standard FPTASs. Instead, we propose a heuristics based on a branch-and-bound-like approach ([Land and Doig, 2010](#); [Clausen, 1999](#)), where we use

a LP relaxation to estimate the value of the objective. This amounts to testing every admissible (discrete) first action, and then keeping the one maximizing the relaxed objective (in which all subsequent actions are relaxed in $[0, 1]$). The same is repeated to choose the following $d - 1$ discrete actions. Overall, we solve d Linear Programs, of sizes $Kd, 2Kd, \dots, Kd^2$. Given a horizon T , and choosing d as in [Theorem 7](#), the total running time is $\mathcal{O}(K^{5/2}T^{9/4})$. For more details about the heuristic, the reader is referred to [Appendix B.3](#). Note that to ease readability we restricted ourselves to positive delays only in the core text. The complete resolution when negative delays are also considered is detailed in [Appendix B.2](#).

Although it is generally hard to derive approximation guarantees, we point out that: (1) this approach works well in practice, delivering the optimal solution in all the cases we tested; (2) as discussed in [Kveton et al. \(2015\)](#), if ISI-CombUCB1 is run with the approximate solver, [Theorem 7](#) can be adapted to bound the regret against the best block according to the solver’s approximation. Finally, note that our calibration approach is not affecting the complexity of finding the reward-maximizing block. Indeed, CombUCB1 is also required to solve an integer linear program similar to (12) to find the subset of base actions maximizing the block reward. \square

4 EXPERIMENTS

In this section, we benchmark ISI-CombUCB1 against vanilla CombUCB1 and OracleGreedy, which plays $\text{argmax}_i \mu_i(\tau_i(t))$ at each time step t , and breaks ties randomly. We measure the algorithms performance in terms of the total cumulative reward, averaged over ten repetitions. For both CombUCB1 and ISI-CombUCB1, the exploration parameter α is set to 1.5, as in ([Kveton et al., 2015](#)). Let d be the block size of CombUCB1, which consequently maintains Kd UCBs. Whenever an arm a is pulled with state $\tau \geq d$ (this might happen as the algorithm is not calibrated), we assume that the algorithm updates the estimate of $\mu_a(d)$ —it actually becomes an estimate of $\mu_a(\tau \geq d)$. For ISI-CombUCB1 to similarly maintain Kd UCBs, we consider blocks of size $d + 1$. Indeed, recall that rewards are ignored the first time an arm is pulled, so that reaching state $\tau_a = d$ requires (at least) $d + 1$ steps. This “extra” first pull does not provide any information (it is only used for calibration purposes), but it is nevertheless considered in the performance. Note finally that both algorithms use the same heuristic based on branch-and-bound and LP relaxations to compute the reward-maximizing block at each time step.

Algorithm 1 ISI-CombUCB1

input : number of arms K , block size d , horizon T

init : $\forall i \leq K, j \leq d, T_1(i, j) = 0, \bar{X}_1(i, j) = 0,$
 $U_1(i, j) = +\infty.$

for t from 1 to T/d **do**

// Play the best block for \tilde{r} based on the UCBs

Play $\hat{B}_t = [a_{t,1} \dots a_{t,d}]$ that maximizes in $(a_s)_{s \leq d}$

$$\sum_{s=1}^d U_t(a_s, \tau_{a_s}(s)) \mathbb{I}\{\exists s_0 < s : a_{s_0} = a_s\}$$

Get rewards $X_t(a_{t,1}, \tau_{a_{t,1}}(1)) \dots X_t(a_{t,d}, \tau_{a_{t,d}}(d))$

// Update the statistics

for $i \leq K$ and $j \leq d$ **do**

if arm i is played with delay j in block \hat{B}_t
 (counting the multiplicities) **then**

$$T_{t+1}(i, j) = T_t(i, j) + 1,$$

$$\bar{X}_{t+1}(i, j) = \frac{T_t(i, j)\bar{X}_t(i, j) + X_t(i, j)}{T_{t+1}(i, j)}$$

else

$$T_{t+1}(i, j) = T_t(i, j), \quad \bar{X}_{t+1}(i, j) = \bar{X}_t(i, j)$$

$$U_{t+1}(i, j) = \bar{X}_{T_{t+1}(i, j)}(i, j) + \sqrt{\frac{\alpha \log(t+1)}{T_{t+1}(i, j)}}$$

We consider the 5-armed LSD bandit with Bernoulli rewards and

$$\mu_1(\tau) = 0.95 \mathbb{I}\{\tau = 3\}, \quad \mu_2(\tau) = \begin{cases} 0.16 & \text{if } \tau = 6 \\ 0.96 & \text{if } \tau \geq 9 \\ 0.14 & \text{otherwise} \end{cases},$$

and $\mu_{3,4,5}(\tau) = 0.15$ for all τ , see Figure 2 (top). The empirical results for $d = 3$ are reported in Figure 2 (bottom). ISI-CombUCB1 converges towards the block $[a^{(1)}, a^{(i)}, a^{(i)}, a^{(1)}]$, where i is any arm in $\{3, 4, 5\}$. Here, the first pull of arm $a^{(1)}$ is needed to calibrate the state of this action, and to get an expected reward of $\mu_1(3)$ in the last pull of the block. On the first 8 time steps, OracleGreedy plays any constant arm, except at $t = 3$ and $t = 6$, where it plays $a^{(1)}$. At $t = 9$, it faces a choice between $\mu_1(3)$ and $\mu_2(9)$ and thus goes for the latter. Then, it will never be able to get $\mu_1(3)$ again, and will only get $\mu_2(9)$ every nine steps with any combination of constant actions in between. As for CombUCB1, it is unable to see the spike at $\mu_1(3)$. Precisely, whenever the algorithm plays arm $a^{(1)}$ with state $\tau > 3$, the expected reward is small, but still contributes to the UCB of $\mu_1(3)$. This decreases the UCB for the arm, and the algorithm pulls it less and less frequently. More experiments testing the calibration sequence approach are provided in the Appendices.

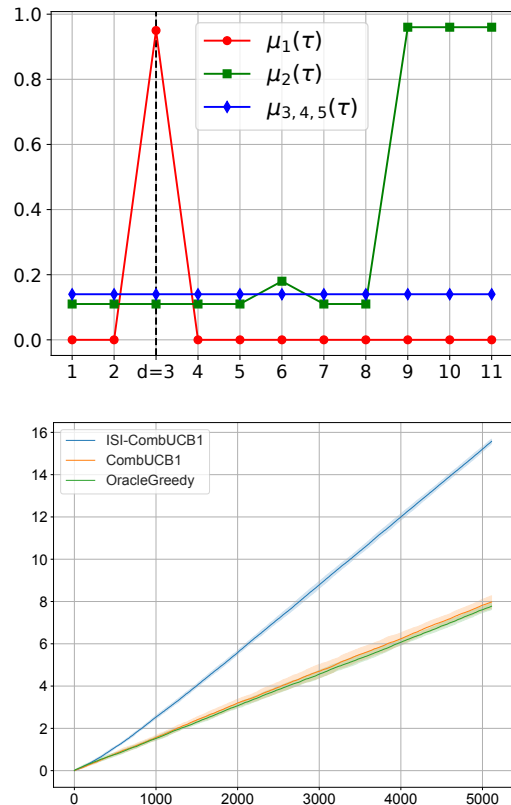


Figure 2: Reward functions with respect to τ (top) and cumulative rewards in thousands (bottom).

Acknowledgments. This work has been partially supported by the academic grant *Innovative Technologies, Interventions and Novel Approaches in the Field of Neuro-Wellness* from Joy Ventures and by the EU Horizon 2020 ICT-48 research and innovation action under grant agreement 951847, project ELISE (European Learning and Intelligent Systems Excellence).

References

- Andersen, M. S., Dahl, J., Vandenberghe, L., et al. (2013). CVXOPT: A Python package for convex optimization. Available at cvxopt.org.
- Audibert, J.-Y., Bubeck, S., and Lugosi, G. (2014). Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45.
- Bar-Noy, A., Bhatia, R., Naor, J., and Schieber, B. (2002). Minimizing service and operation costs of periodic scheduling. *Mathematics of Operations Research*, 27(3):518–544.
- Basu, S., Sen, R., Sanghavi, S., and Shakkottai, S. (2019). Blocking bandits. *arXiv preprint arXiv:1907.11975*.
- Bouneffouf, D. and Féraud, R. (2016). Multi-armed

- bandit problem with known trend. *Neurocomputing*, 205:16–21.
- Cella, L. and Cesa-Bianchi, N. (2020). Stochastic bandits with delay-dependent payoffs. In *International Conference on Artificial Intelligence and Statistics*, pages 1168–1177. PMLR.
- Chen, W., Wang, Y., and Yuan, Y. (2013). Combinatorial multi-armed bandit: General framework and applications. In *International Conference on Machine Learning*, pages 151–159. PMLR.
- Clausen, J. (1999). Branch and bound algorithms—principles and examples. *Department of Computer Science, University of Copenhagen*, pages 1–30.
- Cortes, C., DeSalvo, G., Kuznetsov, V., Mohri, M., and Yang, S. (2017). Discrepancy-based algorithms for non-stationary rested bandits. *arXiv preprint arXiv:1710.10657*.
- Dimitrijević, M. R., Faganel, J., Gregorić, M., Nathan, P., and Trontelj, J. (1972). Habituation: effects of regular and stochastic stimulation. *Journal of Neurology, Neurosurgery & Psychiatry*, 35(2):234–242.
- Fuller-Tyszkiewicz, M., Richardson, B., Lewis, V., Linardon, J., Mills, J., Juknaitis, K., Lewis, C., Coulson, K., O’Donnell, R., Arulkadacham, L., et al. (2019). A randomized trial exploring mindfulness and gratitude exercises as ehealth-based micro-interventions for improving body satisfaction. *Computers in Human Behavior*, 95:58–65.
- Gai, Y., Krishnamachari, B., and Jain, R. (2012). Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking*, 20(5):1466–1478.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–164.
- Heidari, H., Kearns, M. J., and Roth, A. (2016). Tight policy regret bounds for improving and decaying bandits. In *IJCAI*, pages 1562–1570.
- Holte, R., Mok, A., Rosier, L., Tulchinsky, I., and Varvel, D. (1989). The pinwheel: A real-time scheduling problem. In *Proceedings of the 22nd Hawaii International Conference of System Science*, pages 693–702.
- Kleinberg, R. and Immorlica, N. (2018). Recharging bandits. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 309–319. IEEE.
- Kovacs, G., Wu, Z., and Bernstein, M. S. (2018). Rotating online behavior change interventions increases effectiveness but also increases attrition. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–25.
- Kunaver, M. and Požrl, T. (2017). Diversity in recommender systems—a survey. *Knowledge-based systems*, 123:154–162.
- Kveton, B., Wen, Z., Ashkan, A., and Szepesvari, C. (2015). Tight regret bounds for stochastic combinatorial semi-bandits. In *Artificial Intelligence and Statistics*, pages 535–543. PMLR.
- Land, A. H. and Doig, A. G. (2010). An automatic method for solving discrete programming problems. In *50 Years of Integer Programming 1958-2008*, pages 105–132. Springer.
- Leqi, L., Kilinc-Karzan, F., Lipton, Z. C., and Montgomery, A. L. (2020). Rebounding bandits for modeling satiation effects. *arXiv preprint arXiv:2011.06741*. To appear in NeurIPS 2021.
- Levine, N., Crammer, K., and Mannor, S. (2017). Rotating bandits. *arXiv preprint arXiv:1702.07274*.
- Meinlschmidt, G., Lee, J.-H., Stalujanis, E., Belardi, A., Oh, M., Jung, E. K., Kim, H.-C., Alfano, J., Yoo, S.-S., and Tegethoff, M. (2016). Smartphone-based psychotherapeutic micro-interventions to improve mood in a real-world setting. *Frontiers in psychology*, 7:1112.
- Owen, J. E., Kuhn, E., Jaworski, B. K., McGee-Vincent, P., Juhasz, K., Hoffman, J. E., and Rosen, C. (2018). Va mobile apps for ptsd and related problems: public health resources for veterans and those who care for them. *Mhealth*, 4.
- Paredes, P., Gilad-Bachrach, R., Czerwinski, M., Roseway, A., Rowan, K., and Hernandez, J. (2014). Poptherapy: Coping with stress through popculture. In *Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare*, pages 109–117.
- Pike-Burke, C., Agrawal, S., Szepesvari, C., and Grunewalder, S. (2018). Bandits with delayed, aggregated anonymous feedback. In *International Conference on Machine Learning*, pages 4105–4113. PMLR.
- Pike-Burke, C. and Grunewalder, S. (2019). Recovering bandits. *Advances in Neural Information Processing Systems*, 32:14122–14131.
- Schedl, M., Zamani, H., Chen, C.-W., Deldjoo, Y., and Elahi, M. (2018). Current challenges and visions in music recommender systems research. *International Journal of Multimedia Information Retrieval*, 7(2):95–116.
- Schroeder, J., Wilkes, C., Rowan, K., Toledo, A., Paradiso, A., Czerwinski, M., Mark, G., and Linehan, M. M. (2018). Pocket skills: A conversational mobile

web app to support dialectical behavioral therapy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–15.

Seznec, J., Locatelli, A., Carpentier, A., Lazaric, A., and Valko, M. (2019). Rotting bandits are no harder than stochastic ones. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2564–2572. PMLR.

Simchi-Levi, D., Zheng, Z., and Zhu, F. (2021). Dynamic planning and learning under recovering rewards. In *International Conference on Machine Learning*, pages 9702–9711. PMLR.

Tekin, C. and Liu, M. (2012). Online learning of rested and restless bandits. *IEEE Transactions on Information Theory*, 58(8):5588–5611.

Warlop, R., Lazaric, A., and Mary, J. (2018). Fighting boredom in recommender systems with linear reinforcement learning. In *Neural Information Processing Systems*.

Whittle, P. (1988). Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 25(A):287–298.

A TECHNICAL PROOFS

In this section, we gather all the technical proofs omitted in the main body of the paper.

A.1 Proof of Proposition 1

First note that focusing on the case where the μ_i are constant on \mathbb{Z}^- is enough to prove NP-hardness. Within this setting, we actually prove a slightly more general result than Proposition 1, namely that for any $\tau_{\text{init}} \in \mathbb{N}^K$, computing

$$B^* = \operatorname{argmax}_{|B|=d} r(B \mid \tau_{\text{init}}) \quad (13)$$

is NP-hard as soon as d is greater than a certain value d_0 precised later on. In particular, when $\tau_{\text{init}} = \mathbf{1}$, and $d = T$, solving (13) amounts to find the best policy. For some T large enough, we get $d \geq d_0$, and we obtain that computing the optimal policy is NP-hard (Proposition 1). But more interestingly, this result also highlights that the inner optimization problem of a CombUCB1 approach is itself NP-hard. Note that the same result holds for the inner optimization problem of ISI-CombUCB1, as the hardness can be similarly proved with \tilde{r} instead of r in (13).

Our reduction is largely adapted from the one in (Basu et al., 2019). Indeed, our framework encapsulates the setting described by Equations (14), which is the one used by the authors to show the hardness of their problem. The main differences are: 1) we need to take into account the initial state τ_{init} , which can be handled by considering a larger block size d , and 2) in (Basu et al., 2019), any empty slot in the scheduling is automatically filled with a pull of the 0 arm (it is the only arm which can be played), while we have to invoke a finer argument, namely that at any empty slot, playing the 0 arm is the best possible action, since all actions yield a null expected reward, but playing the 0 arm allows every other arms to recharge.

As in (Basu et al., 2019), we thus consider the Pinwheel Scheduling Problem (PSP), see (Holte et al., 1989), which can be stated as follows. Given a set of actions $\{1, \dots, K\}$, associated to delays $(d_i)_{i=1}^K \in \mathbb{N}^K$, the PSP consists in deciding whether it is possible or not to design an action scheduling, i.e., a mapping from \mathbb{N} to $\{1, \dots, K\}$, such that each arm i is played at least once in any sequence of d_i consecutive pulls. A PSP instance with such a scheduling is called a YES instance. Otherwise, it is referred to as a NO instance. Furthermore, a PSP instance is said to be dense if $\sum_{i=1}^K 1/d_i = 1$. In particular, this implies for YES instances that each arm i is played exactly every d_i pulls. It has been shown in (Bar-Noy et al., 2002) that PSP on dense instances is NP-complete. In the next paragraph, we show that PSP on dense instances reduces to particular instances of our problem, described by Equations (14), proving the hardness of our problem.

Given a dense PSP instance, we construct an instance of our problem as follows. For every arm $i \leq K$, we set:

$$\mu_i(\tau) = \mathbb{I}\{\tau \geq d_i\}. \quad (14)$$

We set an additional arm $K + 1$, such that $\mu_{K+1}(\tau) = 0$ for all τ . Given an initial state τ_{init} , we want to find the best block of actions $B^* = \operatorname{argmax}_{|B|=d} r(B \mid \tau_{\text{init}})$. Then, we have two possibilities:

- if the PSP is a YES instance, the PSP schedule gives a 1 reward at each time step. Taking into account the fact that the initial state might not exactly suit the schedule, we obtain that $r(B^* \mid \tau_{\text{init}}, \text{YES}) \geq d - K$.
- if the PSP is a NO instance, we can use the same argument as in (Basu et al., 2019) and prove that $r(B^* \mid \tau_{\text{init}}, \text{NO}) \leq d - \left\lfloor \frac{d}{\prod_{i=1}^K d_i} \right\rfloor$.

Hence, for $d \geq d_0 := (K + 1) \prod_{i=1}^K d_i$, if we were able to compute B^* , and therefore $r(B^* \mid \tau_{\text{init}})$, we could discriminate YES from NO instances, depending on whether $r(B^* \mid \tau_{\text{init}}) \geq d - K$ or not. Solving PSP on dense instances thus reduces to solving (13) for the particular choice of reward functions explicited in (14), therefore proving the hardness of the latter problem. \square

A.2 Proof of Lemma 2

Assume without loss of generality that $a^{(1)}$ is the first action played. If there is no switch of actions, then $a^{(1)}$ is played indefinitely, which means that the sequence is cyclic. If there is only one (respectively two) switch(es)

of actions, then $a^{(2)}$ (respectively $a^{(1)}$) is played indefinitely after some time step, from which the sequence is cyclic. If there are three (or more) switches of actions, then the sequence witnesses twice the switch of actions $[a^{(1)}, a^{(2)}]$, and thus visits twice the state $(\tau_1, \tau_2) = (1, -1)$. Since the policy is deterministic, the same cycle will be repeated. \square

Note however that this property does not resist to the number of arms, as we can exhibit optimal policies with 3 arms that never visit the same state twice. Let $\mathcal{A} = \{a^{(1)}, a^{(2)}, a^{(3)}\}$, and consider the sequence of actions built as follows.

For $m = 1, 2, \dots$ play

- $a^{(1)}$ m times consecutively,
- $a^{(2)}$ m times consecutively,
- $a^{(3)}$ m times consecutively.

Assuming for simplicity that the states τ_i are initialized to 0 and not to 1 (this only impacts the states in the first block), it can be checked that, in block m , the state (τ_1, τ_2, τ_3) is given by

$$\begin{cases} (-t, m-1+t, t) & \text{after the } t^{\text{th}} \text{ pull of } a^{(1)} \\ (t, -t, m+t) & \text{after the } t^{\text{th}} \text{ pull of } a^{(2)} \\ (m+t, t, -t) & \text{after the } t^{\text{th}} \text{ pull of } a^{(3)} \end{cases}$$

From the above formula, it is immediate to check that no state is visited twice by the sequence. Now it is easy to choose μ_1, μ_2, μ_3 such that this policy is optimal. For example, $\mu_i \equiv 1$ for $i = 1, 2, 3$.

A.3 Proof of Proposition 3

Recall that Proposition 3 is proven in the main body of the paper for the special case where the μ_i are constant on \mathbb{Z}^- . The proof in the general case is very similar, but uses the second claim of Lemma 5, and more specifically the way B' is constructed from B .

Similarly to the constant case, we know that there exists a block B^* of length d with average reward (at least) greater than the average reward of the complete optimal sequence. Recalling the notation from the proof in the constant case, we have $B^* = [a_{t_0}^* \dots a_{t_0+d-1}^*]$ and

$$\frac{r(B^* | \boldsymbol{\tau}^*(t_0))}{d} \geq \frac{1}{T} \sum_{t=1}^T r_t^*.$$

Previously, we could simply repeat B^* , which is not possible anymore due to the possible effects showed in (24). Let B' , derived from B as in the proof of Lemma 5. Namely, if $a^{(1)}$ and $a^{(2)}$ are the first two different actions played in B , we have $B' = [a_{t_0}^* a^{(2)} a_{t_0+1}^* \dots a_{t_0+d-1}^*]$. Let $\boldsymbol{\tau}_{B'}$ be the state reached after a play of B' from initial state $\mathbf{1}$. Using Lemma 5, the expected rewards $r_\pi(t)$ obtained by policy π which plays cyclically B' satisfy

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T r_\pi(t) &= \frac{r(B' | \mathbf{1})}{T} + \frac{T-d}{T} \frac{r(B' | \boldsymbol{\tau}_{B'})}{d} \\ &= \frac{d}{T} \frac{r(B' | \mathbf{1})}{d} + \left(1 - \frac{d}{T}\right) \frac{r(B' | \boldsymbol{\tau}_{B'})}{d} \\ &\geq \frac{d}{T} \frac{r(B^* | \boldsymbol{\tau}^*(t_0)) - (K+2)}{d} + \left(1 - \frac{d}{T}\right) \frac{r(B^* | \boldsymbol{\tau}^*(t_0)) - (K+2)}{d} \\ &= \frac{r(B^* | \boldsymbol{\tau}^*(t_0)) - (K+2)}{d} \\ &\geq \frac{1}{T} \sum_{t=1}^T r_t^* - \frac{K+2}{d}. \end{aligned}$$

\square

A.4 Proof of Proposition 4

Similarly to Proposition 1, we only need to focus on the case where the reward functions are constant on \mathbb{Z}^- , and actually prove here a stronger result than Proposition 4. Namely, we show that for some d , computing

$$B^* = \operatorname{argmax}_{|B|=d} r(B \mid \tau_B) \quad (15)$$

where τ_B is the state reached by the system after a play of block B , is NP-hard, even when the reward functions can be totally ordered. The optimal cyclic policy (with cycle length d) is obtained by repeating indefinitely the solution to (15), and the NP-hardness of the latter problem then yields Proposition 4.

This proof is also based on a reduction of the Pinwheel Scheduling Problem (PSP), see Appendix A.1. Given a PSP dense instance $(d_i)_{i=1}^K$, we construct an instance of our problem as follows. For every arm $i = 1, \dots, K$, we set:

$$\mu_i(\tau) = \sqrt{\tau/d_i}. \quad (16)$$

Note that setting (16) is fundamentally different from setting (14) as we have a total ordering of the arms, i.e., for all τ we have $\mu_1(\tau) \geq \mu_2(\tau) \geq \dots \geq \mu_K(\tau)$, with the convention $d_1 \leq d_2 \leq \dots \leq d_K$. We also highlight that we do not introduce an additional null arm here, as opposed to (16). Finally, although the μ_i are unbounded for the moment, thus breaking Definition 1, we see at the end of the proof how to consider bounded μ_i without altering the subsequent analysis.

We now show that solving (15) with reward functions (16) allows us to determine if the PSP instance is a YES or a NO instance. Let n_i be the number of times action $i = 1, \dots, K$ is played in the block, and τ_{ij} , for $j = 1, \dots, n_i$ be the different states in which arm i is pulled. Problem (15) is equivalent to:

$$\max_{(n_i)_{i=1}^K} \max_{(\tau_{ij})_{i=1, j=1}^{K, n_i}} \sum_{i=1}^K \sum_{j=1}^{n_i} \sqrt{\tau_{ij}/d_i} \quad \text{subject to} \quad \begin{cases} \sum_{i=1}^K n_i = d \\ \sum_{j=1}^{n_i} \tau_{ij} = d \quad i = 1, \dots, K \end{cases} \quad (17)$$

We start by maximizing with respect to the $\tau = (\tau_{ij})_{i=1, j=1}^{K, n_i}$. The Lagrangian writes

$$\mathcal{L}(\tau, \lambda) = \sum_{i=1}^K \sum_{j=1}^{n_i} \sqrt{\tau_{ij}/d_i} + \sum_{i=1}^K \lambda_i \left(\sum_{j=1}^{n_i} \tau_{ij} - d \right).$$

The KKT conditions (gradient of the Lagrangian and primal feasibility) write

$$\frac{\partial \mathcal{L}(\tau, \lambda)}{\partial \tau_{ij}} = \frac{1}{2\sqrt{d_i \tau_{ij}}} + \lambda_i = 0 \quad i = 1, \dots, K, j = 1, \dots, d \quad (18)$$

$$\sum_{j=1}^{n_i} \tau_{ij} = d \quad i = 1, \dots, K. \quad (19)$$

Solving (18) for τ_{ij} we obtain a quantity independent of j . Then (19) implies that $\tau_{ij} = d/n_i$ for each $i = 1, \dots, K$. Replacing $\tau_{ij} = d/n_i$ into (17), we can now maximize with respect to the n_i . The Lagrangian writes

$$\mathcal{L}(\mathbf{n}, \lambda) = \sum_{i=1}^K \sqrt{dn_i/d_i} + \lambda \left(\sum_{i=1}^K n_i - d \right)$$

and the KKT conditions (gradient of the Lagrangian and primal feasibility) are

$$\frac{\partial \mathcal{L}(\mathbf{n}, \lambda)}{\partial n_i} = \sqrt{\frac{d}{4n_i d_i}} + \lambda = 0 \quad i = 1, \dots, K \quad (20)$$

$$\sum_{i=1}^K n_i = d \quad (21)$$

such that replacing (20) into (21), we obtain $n_i = d/d_i$, which implies $\tau_{ij} = d_i$.

Assume now that d can be divided by all the d_i , such that $n_i = d/d_i \in \mathbb{N}$. From the values of n_i and τ_{ij} obtained, we can see that the optimal block for (15) corresponds to a Pinwheel schedule. It yields an average reward equal to 1, and is achievable if and only if the Pinwheel instance is a YES instance. Therefore, if we can solve (15), we can tell if the average reward is equal to 1 or smaller, and thus decide whether the instance is a YES or NO instance. We have reduced PSP on dense instances to (15), which is therefore shown to be NP-hard. To our knowledge, this is the first hardness result for decomposable reward functions of the form $\mu_i(\tau) = \mu_i^0 \cdot \gamma(\tau)$.

Now, let $d_{\max} = \max_{i=1, \dots, K} d_i$. Note that replacing (16) with the bounded functions

$$\mu_i(\tau) = \begin{cases} \sqrt{\tau/d_i} & \text{if } \tau \leq d_{\max} \\ \sqrt{d_{\max}/d_i} & \text{otherwise} \end{cases}$$

does not change the optimal schedule (arm i is played every $d_i \leq d_{\max}$ time steps) and the analysis is unchanged. \square

A.5 Proof of Lemma 5

Recall that Lemma 5 is proven in the main body of the paper for the special case where the μ_i are constant on \mathbb{Z}^- . We focus here on the general version only. For any block B of size d , we want to find B' of length d such that for any initial states $\tau_{\text{init}}, \tau'_{\text{init}} \in \mathbb{Z}^K$, we have

$$r(B' \mid \tau'_{\text{init}}) \geq r(B \mid \tau_{\text{init}}) - (K + 2). \quad (22)$$

First, we analyze what happens if we keep the same B , instead of B' , in the left-hand side of (22). Similarly to the constant case, it is important to note that an action is impacted by a change of initial state only the first time it is played in the block (possibly with several consecutive pulls). Indeed, when this sequence (which might be of length 1 if a switch follows the first pull) is interrupted, the arm goes to state 1, no matter the state it was before, and similar subsequent actions then yield similar states/rewards. A second point to notice is that only the first action a_1 in the block can be pulled with a negative state at the beginning of its sequence of pulls. Indeed, let $\tau_{\text{new}} = \delta(\tau_{\text{init}}, a_1)$, where δ is a componentwise version of δ_a , and a_1 is the first action played in B . It is easy to check that for all $a \neq a_1$, we have $\tau_{\text{new}, a} \geq 1$ (either the action was in a negative state and was set to 1 as $a_1 \neq a$ has been played, or it was in a positive state and incremented by 1). Combining these two remarks, we know that for any action a different from $a^{(1)}$ (we can assume that $a_1 = a^{(1)}$ without loss of generality), the loss incurred by a change of initial state due to the pulls of a is bounded by 1. Indeed, the expected rewards obtained during the first play (possibly with consecutive pulls) with initial states τ_{init} and τ'_{init} are respectively

$$\mu_a(\tau_a), \mu_a(-1), \mu_a(-2), \dots \quad \text{and} \quad \mu_a(\tau'_a), \mu_a(-1), \mu_a(-2), \dots \quad (23)$$

where τ_a and τ'_a are two generic **positive** values (thanks to the remark we made earlier) that depend on $\tau_{\text{init}, a}$, $\tau'_{\text{init}, a}$, and the place of the first pull of a in the block. Making these values explicit is not important here, as the boundedness of μ_a ensures anyway that the difference of expected rewards obtained is smaller than 1.

Now, what happens for the first pulls of $a^{(1)}$? We assume that it is played n_1 times consecutively at the beginning of B (again, we might have $n_1 = 1$). Assuming that $\tau_{\text{init}, 1}$ is positive and $\tau'_{\text{init}, 1}$ is negative, the expected rewards collected are respectively

$$\mu_1(\tau_{\text{init}, 1}), \mu_1(-1) \dots \mu_1(-n_1 + 1) \quad \text{and} \quad \mu_1(\tau'_{\text{init}, 1}), \mu_1(\tau'_{\text{init}, 1} - 1) \dots \mu_1(\tau'_{\text{init}, 1} - n_1 + 1). \quad (24)$$

Unlike in the previous case, the difference between these two sequences cannot be bounded by 1 (it might even be equal to n_1 if $\tau'_{\text{init}, 1} \leq -n_1$). This is why the same B cannot be used on both sides of (22). To break this sequence of pulls, we define B' as follows. Without loss of generality, let $a^{(2)}$ be the second different action played in B (if B is only composed of pulls of $a^{(1)}$, we can set $a^{(2)}$ to be any action of \mathcal{A} different from $a^{(1)}$). The block B' is equal to B , except that the second pull of B' is necessarily $a^{(2)}$. We may now face 3 different cases.

- $n_1 = 1$. Note that here $B' = B$. Since $n = 1$, the difference between the two sequences of rewards due to the pulls of $a^{(1)}$ is at most 1. For all other actions, we can use the analysis we developed at the beginning, and the total difference is at most K .

- $n_1 = 2$. Then, denoting by n_2 the number of times $a^{(2)}$ is played consecutively after $a^{(1)}$, the expected rewards obtained by B with initial state τ_{init} and by B' with initial state τ'_{init} are respectively

$$\begin{aligned} & \mu_1(\tau_{\text{init},1}), \quad \mu_1(-1) \text{ or } \mu_1(\tau_{\text{init},1} - 1), \quad \mu_2(\tau_2), \quad \mu_2(-1) \quad \dots \quad \mu_2(-n_2 + 1) \\ \text{and} \quad & \mu_1(\tau'_{\text{init},1}), \quad \mu_2(\tau'_2), \quad \mu_2(-1), \quad \mu_2(-2) \quad \dots \quad \mu_2(-n_2) \end{aligned}$$

where the *or* comes from the fact that we don't know if $\tau_{\text{init},1}$ is positive (then the next reward is $\mu_1(-1)$) or negative (then next reward is $\mu_1(\tau_{\text{init},1} - 1)$). Here, τ_2 and τ'_2 are two generic positive numbers, whose values are not important as the difference between the two sequences is contained in the red rewards, and thus bounded by 3 anyway. For all other $K - 2$ actions, we can apply the standard analysis, such that in total the difference cannot exceed $K + 1$.

- $n_1 \geq 3$. Using the same notation as above, we have now respectively for the first n_1 pulls

$$\begin{aligned} & \mu_1(\tau_{\text{init},1}), \quad \mu_1(-1) \text{ or } \mu_1(\tau_{\text{init},1} - 1), \quad \mu_1(-2) \text{ or } \mu_1(\tau_{\text{init},1} - 2), \quad \dots \quad \mu_1(-n_1 + 1) \text{ or } \mu_1(\tau_{\text{init},1} - n_1 + 1) \\ \text{and} \quad & \mu_1(\tau'_{\text{init},1}), \quad \mu_2(\tau'_2), \quad \mu_1(1), \quad \dots \quad \mu_1(-n_1 + 3) \end{aligned}$$

The red terms incur a loss of at most 3. Then, if $\tau_{\text{init},1} \geq 1$, the remaining rewards (non red) in the play of B are $\sum_{j=3}^{n_1-1} \mu_1(-j) \leq \sum_{j=1}^{n_1-3} \mu_1(-j)$ as μ_1 is nondecreasing on \mathbb{Z}^- , so the rewards obtained by B' are greater. If $\tau_{\text{init},1} < 0$, we have $\sum_{j=3}^{n_1-1} \mu_1(\tau_{\text{init},1} - j) \leq \sum_{j=1}^{n_1-3} \mu_1(-j)$ for the same reasons. So overall, the difference is bounded by 3.

As for the n_2 following pulls, recalling that τ_2 and τ'_2 are positive since the preceding pull is $a^{(1)}$, we have

$$\begin{aligned} & \mu_2(\tau_2), \quad \mu_2(-1), \quad \dots \quad \mu_2(-n_2 + 1) \\ \text{and} \quad & \mu_2(\tau'_2), \quad \mu_2(-1), \quad \dots \quad \mu_2(-n_2 + 1) \end{aligned}$$

with a difference bounded by 1. For the $K - 2$ other actions, we still have the bound of $K - 2$, so that in total the difference does not exceed $K + 2$. □

A.6 Proof of Proposition 6

Recall the notation B_{double}^* , and the additional notation B^* and $\tau^*(t_0)$ introduced in the proof of Proposition 3 in the main body of the paper. We have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T r_{\text{double}}(t) &= \frac{r(B_{\text{double}}^* | \mathbf{1})}{T} + \frac{T-d}{T} \frac{r(B_{\text{double}}^* | \tau_{B_{\text{double}}^*})}{d} \\ &\geq \left(1 - \frac{d}{T}\right) \frac{r(B_{\text{double}}^* | \tau_{B_{\text{double}}^*})}{d} \\ &\geq \left(1 - \frac{d}{T}\right) \frac{r(B^* | \tau_{B^*})}{d} \end{aligned} \tag{25}$$

$$\geq \left(1 - \frac{d}{T}\right) \frac{r(B^* | \tau^*(t_0)) - K}{d} \tag{26}$$

$$\geq \left(1 - \frac{d}{T}\right) \left(\frac{1}{T} \sum_{t=1}^T r^*(t) - \frac{K}{d} \right), \tag{27}$$

where (25) holds because B_{double}^* is a maximizer of $r(B | \tau_B)$, (26) holds due to Lemma 5, and (27) is a direct consequence of the definition of B^* . Using similar arguments, we also have

$$\frac{1}{T} \sum_{t=1}^T r_{\sigma}(t) \geq \frac{r(B_{\sigma}^* | \tau_{B_{\sigma}})}{d+K} \geq \frac{d}{d+K} \frac{r(B^* | \tau_{B_{\sigma}})}{d} \geq \frac{d}{d+K} \frac{r(B^* | \tau^*(t_0)) - K}{d} \geq \frac{d}{d+K} \left(\frac{1}{T} \sum_{t=1}^T r^*(t) \right) - \frac{K}{d+K},$$

and

$$\frac{1}{T} \sum_{t=1}^T \tilde{r}^*(t) \geq \frac{\tilde{r}(\tilde{B}^*)}{d} \geq \frac{\tilde{r}(B^*)}{d} \geq \frac{r(B^* | \tau^*(t_0)) - K}{d} \geq \frac{1}{T} \sum_{t=1}^T r^*(t) - \frac{K}{d}.$$

□

B GENERAL STUDY WHEN μ_a IS NOT CONSTANT ON \mathbb{Z}^-

In this section, we detail the results in the general case where the μ_i are not constant on \mathbb{Z}^- , that were omitted in Section 3 for simplicity. If the definitions of B_{double}^* and B_σ^* remain unchanged, we need to adapt the definition of \tilde{B}^* . Indeed, the goal of \tilde{r} is to use first pulls as a calibration step, such that subsequent pulls are in a controlled state. However, assume that the system is in a state with unknown $\tau_a \leq 0$, and that ISI-CombUCB1 returns a block starting with several pulls of a . None of the rewards obtained by the sequence can be used to update our estimates, as they are obtained in states $\tau_a, \tau_a - 1, \tau_a - 2, \dots$ which are all unknown. This problem is avoided with constant μ_a on \mathbb{Z}^- , as the rewards would be obtained in states $\tau_a, -1, -1, \dots$ such that the first pull indeed plays its calibration role. Therefore, we now define \tilde{B}^* as follows

$$\tilde{B}^* = \underset{|B|=d, a_1 \neq a_2}{\operatorname{argmax}} \tilde{r}(B). \quad (28)$$

The fact that the first two actions in \tilde{B}^* are now different prevents the issues described above. On the other hand, note that B' in the second claim of Lemma 5 also possesses two different first actions, such that the extra constraint in (28) is not harmful in terms of approximation.

B.1 Equivalent of Proposition 6 and Theorem 7

Proposition 8 and Theorem 9 are the analog of Proposition 6 and Theorem 7 respectively.

Proposition 8. *Let $(\mu_i)_{i=1}^K$ be a LSD bandit with K arms. Let $T \geq 0$ be the horizon, and $d \geq 0$ that divides T . Let $r_{\text{double}}(t)$ be the expected rewards obtained at time step t by the policy playing cyclically B_{double}^* . We have*

$$\frac{1}{T} \sum_{t=1}^T r_{\text{double}}(t) \geq \left(1 - \frac{d}{T}\right) \left(\frac{1}{T} \sum_{t=1}^T r^*(t) - \frac{K+2}{d}\right).$$

Let $\sigma \in \mathfrak{S}_K$, and assume that $d + K$ divides T . Let $r_\sigma(t)$ be the expected rewards obtained at time step t by the policy playing cyclically $[B_\sigma, B_\sigma^*]$. We have

$$\frac{1}{T} \sum_{t=1}^T r_\sigma(t) \geq \frac{d}{d+K} \left(\frac{1}{T} \sum_{t=1}^T r^*(t)\right) - \frac{K+2}{d+K}.$$

Let $\tilde{r}^*(t)$ be the expected rewards obtained at time step t by the policy playing cyclically \tilde{B}^* . We have

$$\frac{1}{T} \sum_{t=1}^T \tilde{r}^*(t) \geq \frac{1}{T} \sum_{t=1}^T r^*(t) - \frac{K+2}{d}.$$

Proof. The proof is similar to that of Proposition 6, see Appendix A.6. The only difference is that we cannot use $r(B^* | \tau_{B^*}) \geq r(B^* | \tau^*(t_0)) - K$, as we did during the proof of the first claim of Proposition 6. Instead, we have to involve $(B^*)'$, as defined in Lemma 5. Then, we have

$$\begin{aligned} r(B_{\text{double}}^* | \tau_{B_{\text{double}}^*}) &\geq r((B^*)' | \tau_{(B^*)'}) \geq r(B^* | \tau^*(t_0)) - (K+2), \\ r(B_\sigma^* | \tau_{B_\sigma}) &\geq r((B^*)' | \tau_{B_\sigma}) \geq r(B^* | \tau^*(t_0)) - (K+2), \\ \tilde{r}(\tilde{B}^*) &\geq \tilde{r}((B^*)') \geq r(B^* | \tau^*(t_0)) - (K+2), \end{aligned}$$

which allow to complete the missing parts in the proofs. Note that in the last equation, the first inequality holds true thanks to the fact that $(B^*)'$ has also two first actions that are different, while the second inequality can be recovered from similar arguments as those used to prove Lemma 5. \square

Similarly to the constant case, we can use a CombUCB1 approach to produce a sequence of blocks with small regret against \tilde{B}^* . The only change is the size of the representation: it is now of dimension $2Kd^2$ since we have $2d$ possible states for the arms (ranging from $-d$ to d). Combining Theorem 6 in Kveton et al. (2015) with Proposition 8, we obtain Theorem 9.

Theorem 9. *Let $(\mu_i)_{i=1}^K$ be an LSD bandit with K arms. Let $T \geq 0$ be the horizon, and choose $d \geq 0$ that divides T . Then ISI-CombUCB1, run with block size d and exploration parameter $\alpha = 1.5$, has regret bounded by*

$$R_T \leq \frac{(K+2)T}{d} + 47d\sqrt{2KT \log \frac{T}{d}} + \left(\frac{\pi^2}{3} + 1\right) 2Kd^3.$$

Choosing $d \propto T^{1/4}$, we obtain $R_T = \tilde{O}(KT^{3/4})$, where \tilde{O} is neglecting logarithmic factors.

B.2 Integer Linear Program in the General Case

The last step we need to adapt is the Integer Linear Program which constitutes the optimization problem of ISI-CombUCB1. As explained at the beginning of the section, one difference is that we need to enforce the first two actions of the block to be different for calibration purposes. The second important difference is the size of the representation: we introduce Y^+ and Y^- , both of size Kd^2 , to encode pulls in positive and negative states respectively. The problem can be described as follows.

Let $F \in \{0, 1\}^{K \times d}$, such that $F[i, t] = 1$ if and only if arm i is played for the first time in the block at time step t . Let $Y^+ \in \{0, 1\}^{K \times d \times d}$, such that $Y^+[i, j, t] = 1$ if and only if arm i is played at time step t in state j . Let $Y^- \in \{0, 1\}^{K \times d \times d}$, such that $Y^-[i, j, t] = 1$ if and only if arm i is played at time step t in state $-j$. In comparison to the previous Integer Linear Program, the objective function, i.e., the sum of the current UCBs for the second actions present in the block, and the conditions (AC), (UFP), and (FPF) remain unchanged. Their formula are recalled for completeness. As for the novelties, we introduce the Change Action (CA) constraint, i.e., the second pull in the block must differ from the first. Time Consistency (TC) is now divided into TC for positive states (TC⁺), i.e., an arm can be pulled in state $j \geq 1$ at time step t only if it was pulled at time step $t - j$, and not pulled since, and TC for negative states, (TC⁻), i.e., an arm can be pulled in state $-j \leq -1$ at time step t only if it has been pulled consecutively for the last j time steps. Note that it is easier to express TC⁻ using two conditions, depending on whether $-j = -1$ or $-j \leq -2$. Overall, we have (for index limits that make sense)

$$\sum_{i,j,s} \left(Y^+[i, j, s] + Y^-[i, j, s] \right) U_t(i, j) \tag{objective}$$

$$\forall t \quad \sum_i F[i, t] + \sum_{i,j} Y^+[i, j, t] + \sum_{i,j} Y^-[i, j, t] = 1 \tag{AC}$$

$$\forall i \quad \sum_t F[i, t] \leq 1 \tag{UFP}$$

$$\forall i, t \quad \sum_j Y^+[i, j, t] + \sum_j Y^-[i, j, t] \leq \sum_{s=1}^{t-1} F[i, s] \tag{FPF}$$

$$\sum_i F[i, 2] = 1 \tag{CA}$$

$$\begin{aligned} \forall i, j, t \quad Y^+[i, j, t] &\leq F[i, t - (j + 1)] + \sum_l Y^+[i, l, t - (j + 1)] - \sum_s Y^+[i, j - s, t - s] \\ &\quad + \sum_l Y^-[i, l, t - (j + 1)] - \sum_l Y^-[i, l, t - j] \end{aligned} \tag{TC^+}$$

$$\forall i, t \quad Y^-[i, 1, t] \leq \sum_l Y^+[i, l, t - 1] + F[i, t - 1] \tag{TC_1^-}$$

$$\forall i, t, j \geq 2 \quad Y^-[i, j, t] \leq Y^-[i, j - 1, t - 1] \tag{TC_2^-}$$

Similarly to Section 3, we can approximately solve the above Integer Linear Program by a Branch-and-Bound-like approach, that we detail in the next section.

B.3 Details about the Branch-and-Bound Heuristic

In this section, we provide more details about the heuristic we use to approximately solve the Integer Linear Program (ILP). It works as follows. For every $i \leq K$, we set the first action of the block (of total size d) to $a^{(i)}$. We then solve a relaxed version of the ILP, optimizing only for actions a_2, \dots, a_d (recall that a_1 is fixed to $a^{(i)}$), and allowing for continuous values in $[0, 1]$, instead of $\{0, 1\}$. This can be done efficiently as the relaxed problem is a standard Linear Program. We finally set a_1 to the $a^{(i)}$ which has given the highest reward according to the relaxed ILP. We reiterate, by testing values for the second action, and solving the relaxed version with respect to $a_3 \dots a_d \in [0, 1]^{d-2}$, and so on. Let LP be the function that takes as input the current UCBs and a fixed block of size $s < d$, and outputs the best continuous solution in $[0, 1]^{d-s}$ for actions $a_{s+1} \dots a_d$. Let reward be the function returning the objective value of any sequence (possibly partially continuous). Our heuristic is summarized in Algorithm 2.

Algorithm 2 Approximate ILP Solver

```

input : Current UCBs  $U_t = [U_t(i, j)] \in \mathbb{R}^{K \times 2d}$  for all  $i \leq K$  and  $-d \leq j \leq d$ 
init   :  $\text{block} = []$ 
for  $s = 1 \dots d$  do
    for  $i = 1 \dots K$  do
         $\text{block}_{\text{tmp}} = \text{block} + [a_i]$  // test  $a_i$  as next discrete action (current block size:  $s$ )
         $\text{block}_{\text{cont}} = \text{LP}(U_t, \text{block}_{\text{tmp}})$  // find the best continuous continuation (of size  $d - s$ )
         $r^i = \text{reward}(\text{block}_{\text{tmp}}, \text{block}_{\text{cont}})$  // compute the total relaxed reward of the half-discrete
        // half-continuous block (of overall size  $d$ )
     $i^* = \text{argmax}_{i \leq K} r^i$ 
     $\text{block} = \text{block} + [a_{i^*}]$  // keep the discrete action with highest relaxed reward
return  $\text{block}$ 
    
```

C ADDITIONAL EXPERIMENTS

In this section, we provide additional experiments to elaborate more on the performance of ISI-CombUCB1. In particular, we enrich the comparison made in Section 4 by benchmarking two new algorithms based on calibration sequences (CS). These approaches first calibrate the system by playing a permutation σ of all the arms, and then play the best block according to the state reached after σ , see (7). CS-based approaches are known to be suboptimal, as they calibrate more arms than necessary, here 5 ($= K$) instead of 2 (number of different arms in the optimal block). In addition, since the calibration phase is of size $K = 5 > 3$, it might prevent from seeing the spike of arm $a^{(1)}$ at $\tau = 3$. We therefore benchmark two permutations, $\text{CS-worst} = [a^{(1)}, a^{(2)}, a^{(3)}, a^{(4)}, a^{(5)}]$, and $\text{CS-best} = [a^{(5)}, a^{(4)}, a^{(3)}, a^{(2)}, a^{(1)}]$, named depending on whether they allow to see the spike or not. Again, Algorithm 2 is used as an inner subroutine to approximately compute the solution to the optimization problem (7). Results are shown in Figure 3.

One may argue that the example of Section 4 is unfair to OracleGreedy, because the algorithm is tricked by phenomena that occur beyond the block size d , see Figure 2 (top). We now present an example where it is not the case. Consider the 2-armed LSD bandit with Bernoulli rewards and

$$\mu_1(\tau) = \begin{cases} 0.06 & \text{if } \tau = 1 \\ 0.95 & \text{if } \tau \geq 2 \end{cases} \quad \text{and} \quad \mu_2(\tau) = 0.05 \quad \forall \tau.$$

The reward functions are represented in Figure 4 (left). Here, OracleGreedy constantly pulls arm $a^{(1)}$, for an average reward of 0.06. Instead, ISI-CombUCB1 and CombUCB1 alternate between arms $a^{(1)}$ and $a^{(2)}$, for an optimal average reward of 0.5. Note that ISI-CombUCB1, as opposed to CombUCB1, requires a calibration pull before playing this optimal block. The latter bringing little reward, it slightly degrades the average reward and explains the small difference in performance seen in the plots. However, besides the calibration, both algorithms converge towards the same block, which is optimal. Results for $d = 10$ are reported in Figure 4 (right).




Figure 3: Reward functions with respect to τ (left) and cumulative rewards in thousands (right).




Figure 4: Reward functions with respect to τ (left) and cumulative rewards in thousands (right).

In conclusion, calibration is an operation that might occasionally degrade the performance (in a controlled way), but that on the other hand guarantees to avoid risky decoys, such as the one discussed in Section 4. This behavior is in line with the fact that the regret of ISI-CombUCB1 is well understood theoretically, while CombUCB1 is hard to analyze due to the interferences.

The code is written in Python, and uses the CVXOPT package (Andersen et al., 2013) to solve the Linear Programs. It is publicly available at the following GitHub repository: [GiuliaClerici/LSD_bandits](https://github.com/GiuliaClerici/LSD_bandits).