

Optimal Learning for Structured Bandits

Bart P.G. Van Parys^{*1} and Negin Golrezaei ^{†1}

¹*Sloan School of Management, MIT*

Abstract

We study structured multi-armed bandits, which is the problem of online decision-making under uncertainty in the presence of structural information. In this problem, the decision-maker needs to discover the best course of action despite observing only uncertain rewards over time. The decision-maker is aware of certain structural information regarding the reward distributions and would like to minimize his regret by exploiting this information, where the regret is its performance difference against a benchmark policy which knows the best action ahead of time. In the absence of structural information, the classical upper confidence bound (UCB) and Thomson sampling algorithms are well known to suffer only minimal regret. As recently pointed out by Russo and Van Roy (2018) and Lattimore and Szepesvari (2017), neither algorithm is, however, capable of exploiting structural information which is commonly available in practice. We propose a novel learning algorithm which we call “DUSA” whose worst-case regret matches the information-theoretic regret lower bound up to a constant factor and can handle a wide range of structural information. Our algorithm DUSA solves a dual counterpart of regret lower bound at the empirical reward distribution and follows the suggestion made by the dual problem. Our proposed algorithm is the first computationally viable learning policy for structured bandit problems that suffers asymptotic minimal regret.

Keywords— Structured Bandits, Online Learning, Convex Duality.

^{*}vanparys@mit.edu

[†]golrezae@mit.edu

1 Introduction

In the multi-armed bandit framework, a decision-maker is repeatedly offered a set of arms (options) with unknown rewards to choose from. In order to perform well in this framework, one needs to strike a balance between exploration and exploitation. Many classical bandit algorithms, including the UCB algorithm (Garivier and Cappé 2011) and Thompson sampling (Thompson 1933, Kaufmann et al. 2012), have been designed specifically to optimally trade-off exploration and exploitation in the absence of structural information. While these algorithms perform well when the rewards of all arms are arbitrary and unstructured, they perform poorly when there is a correlation structure between the rewards of the arms (Russo and Van Roy 2018, Lattimore and Szepesvari 2017).

Structural information allows a decision-maker to use rewards observed concerning one arm to indirectly deduce knowledge concerning other arms as well. As an example, in revenue management problems, where we use posted price mechanisms to sell items, every arm corresponds to a price. The probability of receiving a non-zero reward ought to decrease with the price as suggested by standard economic theory and represents structural information. In healthcare, one expects to see a similar performance for drugs (arms) which are composed of similar active ingredients. Structural information reduces the need for exploration and hence directly reduces the suffered regret. However, structured bandit problems are far more complicated than their classical counterparts in which rewards observed in the context of one arm carry no information about any other arm. We present here a unified yet flexible framework to study structured multi-armed bandit problems. *Our key research question is how can one exploit structural information to design efficient optimal learning algorithms that help the decision-maker to obtain higher rewards?*

Mimicking the regret lower bound. We take a principled approach to design an effective learning algorithm for structured bandits. Our approach is based on the information-theoretic regret lower bound for structured bandits characterized first by Graves and Lai (1997). This regret lower bound quantifies the exact extent to which structural information may reduce the suffered regret of any uniformly good bandit policy, i.e., any policy with a sublinear worst-case regret bound. The regret lower bound is characterized as a semi-infinite optimization problem where the decision variable can be interpreted as the rate with which suboptimal arms need to be explored. The constraints ensure that uniformly good bandit policies can distinguish the true reward distribution from certain “deceitful” distributions using an optimal hypothesis test; see Equation (3) for their formal definition.

Having this interpretation of the regret lower bound in mind, we design an algorithm that

aims at targeting the exploration rates suggested by the regret lower bound. However, targeting the suggested exploration rates exactly is not possible as obtaining these rates requires the actual reward distributions of the arms to be known. One way to overcome this challenge is to follow the exploration rates suggested by the regret lower bound, not for the actual unknown reward distribution, but for its empirical counterpart based on the past rewards instead. Unfortunately, this idea put forward by Combes et al. (2017) does not lead to a practical algorithm as obtaining the exploration rates suggested by the regret lower bound in every round involves solving a semi-infinite optimization problem which, except for a few exceptional cases, is computationally prohibitive. We address this issue in the following two ways:

- i) **Dual perspective.** First, we design a dual-based algorithm that instead of solving the regret lower bound problem directly solves its dual counterpart. The dual counterpart, which is a finite convex optimization problem, can be solved effectively and as a result, our dual-based algorithm is computationally efficient for a large class of structured bandits. By taking advantage of the dual counterpart, we offer a unified approach to incorporate convex structural information. To the best of our knowledge, this is the first use of convex duality to directly derive asymptotically optimal and efficiently computable online algorithms for a general class of structured bandit problems.
- ii) **Sufficient information test.** Second, via constructing an easy-to-solve *sufficient information* test, our policy, which is called DUal Structure-based Algorithm (DUSA), solves the dual counterpart of the regret lower bound problem in merely a logarithmic number of rounds, i.e., $O(\log(T))$, instead of solving it in every round, where T is the number of rounds. Although solving the dual problem is computationally efficient in principle, it is still desirable to remove the necessity to solve this problem in every round. In order to do so, our sufficient information test, which is a simple one-dimensional optimization problem, verifies in any particular round whether sufficient exploration has been done in the past and if so, the dual problem does not need to be resolved.

Minimal regret. We show that DUSA can optimally exploit any structural information which can be represented by an arbitrary convex constraint on the reward distribution. This novel design leads to an optimal learning algorithm whose regret matches the theoretical regret lower bound. Put differently, DUSA optimally utilizes the underlying structure by solving a sequence of dual counterparts of regret lower bound problem and this allows DUSA to offer a novel duality principle for structured bandit problems.

To recapitulate, we present (i) a unified approach to a large selection of structured bandit

problems which have been studied extensively in the literature, while at the same time (ii) present a novel approach which is flexible enough to generalize to the even larger class of bandit problems with arbitrary convex structural information. In particular, we can handle any arbitrary combination of the structural information studied prior to this work. Finally, (iii) we believe that the presented work is the first to put forward a computationally practical approach to this quite general class of structured bandit problems.

Outline. The rest of this manuscript is organized as follows. In Section 2, we review related work. Section 3 presents our structured bandit model. In Section 4, we provide the regret lower bound for structured multi-armed bandit problems. Section 5 presents a high level idea of our algorithm and characterizes the dual counterpart of the regret lower bound problem. In Section 6, we provide our learning algorithm DUSA. Section 7 presents the proof of the regret bound of our algorithm. We conclude in Section 8. The proof of the technical results and discussion on computational complexity of DUSA are provided in appendices.

2 Related Work

Our work contributes to the literature on the stochastic multi-armed bandit problem. As stated earlier, the multi-armed bandit framework has been widely applied in many different domains; see, for example, Besbes and Zeevi (2009), Keskin and Zeevi (2014), and Golrezaei et al. (2019). Most of the papers in the literature assume that the rewards of arms are statistically independent of each other; see Bubeck and Cesa-Bianchi (2012) for a survey on multi-armed bandits. While this independence assumption simplifies the problem of designing a learning algorithm, when arms are correlated, it can lead to a suboptimal regret. Considering this, many papers aim at designing learning algorithms that exploit structural information. However, most of these papers focus on very restricted types of structural information.

Several papers study the design of learning algorithms in the presence of a linear reward structure (Dani et al. 2008, Rusmevichientong and Tsitsiklis 2010, Mersereau et al. 2009, Lattimore and Szepesvari 2017). Some other papers focus on a Lipschitz reward structure that roughly speaking, enforces the reward of similar arms to be close to each other (Magureanu et al. 2014a, Mao et al. 2018). Other structures that have been studied include imposing upper bounds on the average rewards of arms (Gupta et al. 2019) and imposing lower and upper bounds on the realized rewards of the arms (Bubeck et al. 2019). There are also papers that study structural information in (i) contextual multi-armed bandit problems (e.g., Slivkins (2011) and Balseiro et al. (2019)) and (ii) combinatorial decision-making (e.g.,

Streeter and Golovin (2008), Zhang et al. (2019), and Niazadeh et al. (2020)).

In this work, we present a general and unified framework to exploit a wide range of structures including some of the structures we discussed above. The novel class of asymptotically optimal online algorithms we propose here is inspired by those discussed in Combes et al. (2017). They similarly consider online algorithms that imitate the information-theoretic lower bound by solving semi-infinite regret lower bound optimization problems over time. At a high level, there are three main differences between our algorithm and theirs. First and foremost, our algorithm does not require to solve a semi-infinite regret lower bound optimization problem. It instead solves its dual convex counterpart, which is computationally more tractable than the primal characterization of the regret lower bound problem. Second, while the online algorithm in Combes et al. (2017) needs to solve the regret lower bound optimization problem in every round, our algorithm only solves its dual counterpart in a logarithmic number of rounds. As stated earlier, by taking advantage of the convex duality theorem, we design a simple sufficient information test that allows our algorithm not to solve the dual problem in every round. Hence, by relying on the duality perspective of the lower regret bound, we not only obtain a tractable reformulation, we also significantly decrease the number of times that the dual counterpart of the regret lower problem needs to be solved. Finally, Combes et al. (2017) assume that the reward distribution of each arm is uniquely parameterized by its mean. We make no such limiting assumption and allow explicitly for arms to have almost any arbitrary reward distribution. We will discuss our model assumptions more closely in Section 3.

By mimicking the regret lower bound that embeds the structural information, our algorithm may play suboptimal arms in order to obtain information about other arms. This is in contrast with UCB and Thompson Sampling. These algorithms stop pulling suboptimal arms once they verified their suboptimality. This prevents these algorithms to exploit structural information to the fullest extent as observed by Lattimore and Szepesvari (2017) and Russo and Van Roy (2018). Pulling suboptimal arms in the presence of structural information has been shown to be an effective technique to reduce regret (Russo and Van Roy 2018). In Russo and Van Roy (2018), the authors design a novel algorithm, which they name Information Directed Sampling (IDS), that aims to capture the structural information by balancing the regret of pulling an arm with its information gain. Our work is different from that of Russo and Van Roy (2018) in three aspects. First, while we take a frequentist approach, IDS is a Bayesian bandit algorithm. Second, we present a systematic way to incorporate structural information which allows us to design an algorithm whose regret matches the regret lower bound whereas IDS does not necessarily obtain minimal regret. Third, while IDS may have

to update its decision in every round, our algorithm updates its strategy only in a logarithm number of rounds.

Our work is also related to the literature on robust optimization. In particular, the regret lower problem by Graves and Lai (1997), which our algorithm mimics, has semi-infinite constraints that can be interpreted as a linear robust constraint with respect to all neighboring distributions as measured by their information distance; see also Cover and Thomas (2012). Robust constraints which are to hold for all distributions in such a given *ambiguity set* have seen a surge of interest in the recent distributionally robust optimization literature building on the pioneering work of Delage and Ye (2010). For instance, distributionally robust constraints with respect to a myriad of statistical distances such as the Prokhorov metric (Erdoğan and Iyengar 2006), the Wasserstein distance (Pflug and Wozabal 2007, Mohajerin Esfahani and Kuhn 2018, Zhao and Guan 2018) and general ϕ -divergences (Bayraksan and Love 2015) have been studied extensively. In the distributionally robust optimization literature, our information distance function, which may be regarded as a particular ϕ -divergence¹, plays a predominant role. The tractability of ϕ -divergence based distributionally robust constraints was considered first by Wang et al. (2016) who show that they can be represented efficiently using standard convex cone constraints. From a computational point of view, this implies that their dual counterpart can be readily solved via modern interior point optimization algorithms when the structural information can be represented efficiently. Finally, Van Parys et al. (2017) show that in a stochastic optimization setting with independent sample data, a distributionally robust optimization formulation with respect to the same information distance function as discussed in the current work yields information-theoretically optimal decisions.

3 Structured Multi-armed Bandit Problems

In stochastic multi-armed bandit problems, a decision-maker needs to select an arm from a finite set X per round over the course of T rounds. When arm x is pulled in round t , the learner earns a nonnegative random reward $R_t(x) = r \in \mathfrak{R}$ with probability $P(r, x)$.² The rewards $R_t(x)$ are assumed to be independent across all arms $x \in X$ and rounds $t \in [T]$. The learner’s goal is to maximize the total expected reward over the course of T rounds

¹The information distance function which we formally state in Equation (4), otherwise known as relative entropy, is in fact a ϕ -divergence for $\phi : t \mapsto -\log(t)$.

²To ease the exposition, we assume that the set of all potential rewards \mathfrak{R} is common to all arms and is discrete. Extending the results in this paper to a setting in which the set of rewards $\mathfrak{R}(x)$ may depend on the arm x is trivial but requires more burdensome notation.

by pulling the arms. The reward distributions are unknown to the learner beyond the fact that $P = (P(r, x))_{r \in \mathfrak{R}, x \in X} \in \mathcal{P}$, where \mathcal{P} is a closed convex subset with non-empty interior of the set of all reward distributions.³ The compact set \mathcal{P} captures the structural information concerning the unknown reward distribution and allows us to have a unified representation of the structural information discussed previously. For instance, in the context of revenue management problem where each arm x corresponds to a price, the learner knows ahead of time that $P(r = 0, x) \geq P(r = 0, x')$ for any $x \geq x'$; that is, $\mathcal{P} = \{Q \in \mathcal{P}_\Omega : Q(r = 0, x) \geq Q(r = 0, x') \text{ for any } x \geq x'\}$, where \mathcal{P}_Ω is the set of all possible reward distributions:

$$\mathcal{P}_\Omega := \left\{ Q \in \mathbb{R}_+^{|\mathfrak{R}| \times |X|} : \sum_{r \in \mathfrak{R}} Q(r, x) = 1 \quad \forall x \in X \right\}. \quad (1)$$

As another example, if the learner is aware of a lower bound ℓ_x and an upper bound u_x on the expected reward of arm $x \in X$, then the compact set \mathcal{P} can be written as $\{Q \in \mathcal{P}_\Omega : \sum_{r \in \mathfrak{R}} rQ(r, x) \in [\ell_x, u_x] \quad \forall x \in X\}$. We present several illustrating examples at the end of this section in Example 1.

Consider a policy π used by the learner. Let x_t^π be the arm selected by the policy π in round t . This selection is based on previously selected arms and observations; that is, the random variable x_t^π is \mathcal{F}_t^π -measurable, where \mathcal{F}_t^π is the σ -algebra generated by $(x_1^\pi, R_1(x_1^\pi), \dots, x_{t-1}^\pi, R_{t-1}(x_{t-1}^\pi))$. Let Π be the set of all policies whose arm selection rules in any round $t \in [T]$ is \mathcal{F}_t^π -measurable. The performance of the policy $\pi \in \Pi$ over T rounds is quantified as its *regret*, defined below, which is the gap between the expected reward of policy π and the maximum expected reward of an omniscient learner:

$$\text{REG}_\pi(T, P) := T \arg \max_{x \in X} \sum_{r \in \mathfrak{R}} rP(r, x) - \sum_{t=1}^T \mathbb{E} [\sum_{r \in \mathfrak{R}} rP(r, x_t^\pi)]. \quad (2)$$

Here, the expectation is with respect to any potential randomness in the bandit policy π . For any reward distribution $P \in \mathcal{P}$, let $x^*(P) \in \arg \max_{x \in X} \sum_{r \in \mathfrak{R}} rP(r, x)$ denote the set of optimal arms. The per-round expected reward of an optimal arm is $\text{REW}^*(P) := \max_{x \in X} \sum_{r \in \mathfrak{R}} rP(r, x)$. Then, the first term in the regret of policy π can be written as $T \cdot \text{REW}^*(P)$. Note that pulling any suboptimal arm in set $\tilde{X}(P) := X \setminus x^*(P)$ offers less expected reward than pulling arm $x^*(P)$. We will denote this suboptimality gap of arm x with $\Delta(x, P) := \text{REW}^*(P) - \sum_{r \in \mathfrak{R}} rP(r, x)$. If a policy keeps pulling suboptimal arms in every round, it will suffer a large regret which is linear in T . A sensible bandit policy π must thus manage to keep a small regret uniformly over all \mathcal{P} . Definition 1 formalizes this notion,

³Note that our assumption that the set of all potential reward distributions \mathcal{P} is a compact convex set with a non-empty interior is rather mild. Indeed, it is satisfied for all examples we discuss in Section A.

which was first introduced by Lai and Robbins (1985).

Definition 1 (Uniformly Good Policies). *A policy $\pi \in \Pi$ is uniformly good if for all $\alpha > 0$ and for any reward distribution $P \in \mathcal{P}$, we have $\limsup_{T \rightarrow \infty} \text{REG}_\pi(T, P)/T^\alpha = 0$ when the number of rounds T tends to infinity.*

It is generally impossible for uniformly good policy $\pi \in \Pi$ to obtain a zero regret. Indeed, in most cases, an online policy must balance exploiting the arms that is optimal given the current information and exploring seemingly suboptimal arms. Fortunately, the total regret caused by exploring suboptimal arms can be kept relatively small. As we discuss in Section 4, the regret lower bounds for bandit problems indicate that one should expect the regret to grow logarithmically in T .

We finish this section with singling out three multi-armed bandit problems that each enjoys a varying amount of structural information. We want to remark that these examples are chosen not to illustrate the generality of our approach but rather serve as illustrative examples to the main concepts introduced in our paper. Other structured bandits that fit our framework include linear, polynomial, Kernel, monotone, and convex bandit problems; see Section A for more details. We would like to point out that the main strength of our proposed approach is not in solving the aforementioned multi-armed bandit problems better than the policies which have been tailored to one particular bandit multi-armed bandit problem. Rather, our algorithm is universal in the sense that it is applicable, suffering only minimal regret, to any structured bandit problem which can be efficiently characterized by a convex set \mathcal{P} . Hence, it is particularly well suited to practical problems which enjoy an idiosyncratic information structure which is not well described by structural bandit problems which have been studied in the literature.

Example 1 (Structural information: feasible set \mathcal{P}).

1. *Generic bandits.* Assume that the reward of any arms x is drawn from an arbitrary distribution with the discrete support of \mathfrak{R} . In this case, the set \mathcal{P} coincides with \mathcal{P}_Ω , defined in Equation (1).
2. *Separable bandits.* Assume that the reward of arm x can be any distribution in \mathcal{P}_x where \mathcal{P}_x , $x \in X$, is a closed feasible set for the reward distribution of arm x . Then, the set \mathcal{P} coincides with $\prod_{x \in X} \mathcal{P}_x$. For example, if \mathcal{P}_x is $\{Q \in \mathbb{R}_+^{|\mathfrak{R}|} : \sum_{r \in \mathfrak{R}} Q(r) = 1\}$ for all arms x , we have $\mathcal{P} = \mathcal{P}_\Omega$.
3. *Lipschitz Bernoulli bandits.* Assume that the reward distribution for all the arms is Bernoulli, i.e., $\mathfrak{R} = \{0, 1\}$, and the expected reward for each arm $x \in X \subseteq \mathfrak{R}$ is a Lipschitz

function. Then, the set $\mathcal{P} = \mathcal{P}_{\text{Lips}}$ is given as $\{Q \in \mathcal{P}_{\Omega} : Q(1, x) - Q(1, x') \leq L \cdot D(x, x') \quad \forall x \neq x' \in X\}$, where $D(x, x')$ is a symmetric distance function between arms x and x' and $L \geq 0$ is the Lipschitz constant.

4 Regret Lower Bound

The regret for bandit problems on which some structural information is imposed should get smaller as more information is provided on the unknown reward distribution P . Nevertheless, as we will show, the regret is still expected to scale logarithmically with the number of rounds T . The precise logarithmic rate with which the regret needs to grow can be quantified via an optimization problem. To do so, we focus on the set of “deceitful” reward distributions which makes learning the optimal arm challenging. Consider a multi-armed bandit problem with a given reward distribution P . For any arm x' , we define

$$\text{DECEIT}(x', P) := \left\{ Q \in \mathcal{P} : \begin{array}{l} Q(x^*(P)) = P(x^*(P)), \\ \sum_{r \in \mathfrak{R}} rQ(r, x') > \sum_{r \in \mathfrak{R}} rQ(r, x^*(P)) \end{array} \right\} \quad (3)$$

as sets of deceitful reward distributions relative to the actual but unknown reward distribution P . Deceitful reward distributions are precisely those reward distributions in \mathcal{P} that behave identical to the true distribution P when the unknown optimal arm $x^*(P)$ is played, i.e., see the first constraint in (3), but deceptively have a better arm x' to play as per the last constraint in (3). Put differently, the reward distribution for the actual optimal arm $x^*(P)$ is the same in P as for any deceitful reward distribution $Q \in \text{DECEIT}(x', P)$. For the deceitful reward distribution Q , however, the expected reward of arm x' is greater than that of arm $x^*(P)$.

To obtain a low regret, a policy should be able to distinguish the actual reward distribution P from its deceitful reward distributions based on the past observed rewards. Later, we discuss the extent to which this can be quantified in terms of the information distance between reward distribution P and its deceitful distributions $\text{DECEIT}(x', P)$. Specifically, the distribution P can be distinguished from its deceitful reward distributions if the information distance between the distribution P and its deceitful reward distributions is not too “small.” We will make this statement formal shortly hereafter. Before that, we recall that the information distance between any two positive measures M and M' on a set \mathfrak{R} is characterized as

$$I(M', M) := \sum_{r \in \mathfrak{R}} M'(r) \log(M'(r)/M(r)) - M'(r) + M(r) \quad (4)$$

if $M' \ll M$, and $+\infty$ otherwise. Here, $M' \ll M$ denotes the implication $M(r) = 0 \implies M'(r) = 0$ for all $r \in \mathfrak{X}$. Finally, we use the shorthand $M > 0$ for the implication $M(r) > 0$ for all $r \in \mathfrak{X}$. For those unfamiliar, all properties relevant to this paper concerning the information distance are conveniently presented in Section B. We are now ready to present the well-known result stating a lower regret bound as an optimization problem in terms of information distances between P and its deceitful distributions.

Proposition 1 (Regret Lower Bound). *Let π be any uniformly good policy. For any reward distribution $0 < P \in \mathcal{P}$, we have that $\liminf_{T \rightarrow \infty} \text{REG}_\pi(T, P)/\log(T) \geq C(P)$, where the regret lower bound function is characterized as*

$$\begin{aligned} C(P) &:= \inf_{\eta \geq 0} \sum_{x \in \tilde{X}(P)} \eta(x) \Delta(x, P) \\ \text{s.t. } &1 \leq \text{DIST}(\eta, x, P) \quad \forall x \in \tilde{X}(P). \end{aligned} \tag{5}$$

Here, if the set $\text{DECEIT}(x', P)$ is non-empty, the distance function DIST is defined as follows

$$\begin{aligned} \text{DIST}(\eta, x', P) &:= \inf \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) \\ \text{s.t. } &Q \in \text{DECEIT}(x', P). \end{aligned} \tag{6}$$

Otherwise, we take $\text{DIST}(\eta, x', P) = \infty$.

Proposition 1 is a straightforward corollary (Combes et al. 2017, Theorem 1) of a much more general result proven by Graves and Lai (1997, Theorem 1) in a Markovian setting. Thus, we do not provide its proof here. A special case of Proposition 1 for discrete bandit model sets \mathcal{P} was proven much earlier by Agrawal et al. (1989). As indicated in Proposition 1 by the infima, we make no assumptions whether or not either minimization problems (5) or (6) are attained. This will be a source of technical difficulty which needs to be addressed.

The decision variable η in the lower bound (5) can be interpreted as a proxy for the logarithmic rate $N_{T+1}(x) \approx \eta(x) \cdot \log(T)$ with which any uniformly good policy must pull each suboptimal arm $x \in \tilde{X}(P)$ during the first T rounds in order to have a shot at suffering only logarithmic loss. Here, $N_{T+1}(x)$ denotes the number of rounds in the first T rounds that arm x is pulled. The logarithmic rate with which the regret is accumulated is given as the objective function of the lower bound (5). Recall that $\Delta(x, P)$ is the suboptimality gap of arm x under reward distribution P .

The constraint in the lower bound (5) can be regarded as a condition on the amount of information collected on the suboptimal arms for any policy to be uniformly good. In fact, the constraint $\text{DIST}(\eta, x', P) \geq 1$ for $x' \in \tilde{X}(P)$ admits an interpretation in terms of optimal

hypothesis testing. In order to have a uniformly small regret, any policy must indeed identify the optimal arm with high probability. To do so, Graves and Lai (1997, Theorem 1) argue that any uniformly good policy must be able to distinguish the actual distribution P from its deceitful reward distributions $\text{DECEIT}(x, P)$ for all $x \in \tilde{X}(P)$, using some optimal hypothesis test.

The lower bound stated in Proposition 1 and its properties will be of central importance in this paper. Thus, in the next section, we will have a closer look at the distance function on which the lower bound $C(P)$ is critically based. In particular, we partition the suboptimal arms $\tilde{X}(P)$ into two groups: non-deceitful arms $\tilde{X}_n(P)$ and deceitful arms $\tilde{X}_d(P)$. For a suboptimal arm x in group $\tilde{X}_n(P)$, the set $\text{DECEIT}(x, P)$ is empty and as a result, such an arm is mainly pulled to rule out the suboptimality of other arms $x' \in \tilde{X}_d(P)$ for which the set $\text{DECEIT}(x', P)$ is not empty. We finish this section by presenting our main result informally.

Theorem 1 (Main Result (Informal)). *For any $0 < \epsilon < 1/|X|$, there exists a computationally efficient bandit policy π whose expected regret bound is essentially optimal; that is, $\limsup_{T \rightarrow \infty} \mathbb{E}[\text{REG}_\pi(T, P)]/\log(T) \leq (1+\epsilon)C(P)+O(\epsilon)$ where the regret is computed against the benchmark that knows the best arm in advance and the expectation is with respect to any randomness in policy π .*

4.1 Partitioning of Suboptimal Arms

The constraints in the regret lower bound problem, presented in Proposition 1, require that for any suboptimal arm $x \in \tilde{X}(P)$, the information distance between P and any of its deceitful distributions $\text{DECEIT}(x, P)$ to be at least one. However, for some suboptimal arms, the set of deceitful distributions may be empty. Let

$$\begin{aligned} \text{REW}_{\max}(x, P) &:= \max_{\sum_{r \in \mathfrak{R}} rQ(r, x)} \\ \text{s.t. } & Q \in \mathcal{P}, Q(x^*(P)) = P(x^*(P)) \end{aligned} \tag{7}$$

be the maximum reward an arm x can yield given that the first deceitful condition $Q(x^*(P)) = P(x^*(P))$ hold. Obviously, when one has $\text{REW}_{\max}(x, P) \leq \text{REW}^*(P)$, the second requirement, i.e., $\sum_{r \in \mathfrak{R}} rQ(r, x) > \sum_{r \in \mathfrak{R}} rQ(r, x^*(P)) = \text{REW}^*(P)$, cannot be satisfied for any deceitful reward distribution Q . Based on this observation, we partition the suboptimal arms into the following two groups.

- *Non-deceitful arms $\tilde{X}_n(P)$.* A suboptimal arm x is non-deceitful, i.e., $x \in \tilde{X}_n(P)$, if $\text{REW}_{\max}(x, P) \leq \text{REW}^*(P)$. For any non-deceitful arm $x \in \tilde{X}_n(P)$, the set of associated

deceitful distributions $\text{DECEIT}(x, P)$ is empty. Consequently, non-deceitful arms do not require any conditions on the rates with which suboptimal arms need to be played in the regret lower bound (5); we have indeed $\text{DIST}(\eta, x, P) = +\infty$ for any logarithmic exploration rate η . That is not to say that these arms are not to be explored. It may be beneficial to play non-deceitful arms to gain sufficient information on the set of arms we describe next.

- *Deceitful arms* $\tilde{X}_d(P)$. Suboptimal arm $x \in \tilde{X}_d(P)$ if $\text{REW}_{\max}(x, P) > \text{REW}^*(P)$. For these arms, the set of associated deceitful distributions $\text{DECEIT}(x, P)$ is non-empty and the distance $\text{DIST}(\eta, x, P)$ between reward distribution P and these deceitful distributions is finite. As a result, the learner has to gather enough information by exploring suboptimal arms to effectively reject their associated deceitful distributions.

In the following, we present an example to further clarify our notion of deceitful and non-deceitful arms. Consider a separable Bernoulli bandit problem with potential rewards $\mathfrak{R} = \{0, 1\}$ and only two arms $X = \{a, b\}$. Assume that the reward distribution of arm a is such that $P(0, a) \geq \frac{2}{5}$ and the reward distribution of arm b can be arbitrary. That is, we have a separable bandit problem with reward distribution in the set

$$\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2 := \{Q \in \mathbb{R}_+^2 : \sum_{r \in \{0,1\}} Q(r) = 1, Q(0) \geq \frac{2}{5}\} \times \{Q \in \mathbb{R}_+^2 : \sum_{r \in \{0,1\}} Q(r) = 1\}. \quad (8)$$

Clearly, the maximum rewards for each arm are given as $\text{REW}_{\max}(a, P) = 1 - \frac{2}{5} = \frac{3}{5}$ and $\text{REW}_{\max}(b, P) = 1$ and are independent of P . To see why $\text{REW}_{\max}(a, P) = \frac{3}{5}$, recall that $P(0, a) \geq \frac{2}{5}$. Now, let the actual reward distributions be $P(a) = [1/2, 1/2]$ and $P(b) = [1/5, 4/5]$. That is, the probability of receiving reward zero under arm a (respectively arm b) is $1/2$ (respectively $1/5$). Then, it is easy to see that the optimal arm is b , i.e., $x^*(P) = b$, and $\text{REW}^*(P) = 4/5$. Observe that $\text{REW}_{\max}(a, P) < \text{REW}^*(P)$ and as a result, arm a is not deceitful. To make the concept of non-deceitful arms more tangible, assume that the learner keeps playing the optimal arm b . Eventually, by just playing the optimal arm, he will learn that the non-deceitful arm a is not optimal. This is so because his empirical average reward of the optimal arm converges to $4/5$, which is greater than the maximum reward that one expects to obtain from suboptimal arm a . (Recall that $\text{REW}_{\max}(a, P) = \frac{3}{5}$ for any $P \in \mathcal{P}$.) This implies that for this example, the lower bound, $C(P)$, presented in Proposition 1, is zero.

The following proposition characterizes the distance function DIST for deceitful arms.

Proposition 2 (Distance Function for Deceitful Arms). *Consider any reward distribution*

$P \in \mathcal{P}$, deceitful arm $x' \in \tilde{X}_d(P)$, and positive exploration rate $\eta \geq 0$. Then,

$$\begin{aligned} \text{DIST}(\eta, x', P) &= \min_{Q \in \mathcal{P}} \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) \\ \text{s.t. } & Q(x^*(P)) = P(x^*(P)), \sum_{r \in \mathfrak{R}} rQ(r, x') \geq \text{REW}^*(P). \end{aligned} \quad (9)$$

The proof of this result is presented in Appendix G. The main benefit of the characterization of the distance function for a deceitful arm x' stated in Proposition 2 is that the minimum in the optimization problem (9) is in fact achieved and is finite. Hence, the partitioning of arms into two types alleviates the technical difficulty caused by potential non-attainment of the infimum (6). We finish this section by revisiting our running examples.

Example 2 (Regret Lower Bound). *Using Proposition 2, we characterize the lower bound and distance functions (C and DIST) for our running examples: generic, separable, and Lipschitz Bernoulli bandits.*

1–2. *Generic and separable bandits. For generic, and more generally separable bandits, the lower bound presented in Proposition 1 coincides precisely with the results first proven in the seminal work of Lai and Robbins (1985). Recall that for separable bandits, the set \mathcal{P} decomposes as the Cartesian product $\prod_{x \in X} \mathcal{P}_x$. This implies that for any reward distribution $P \in \mathcal{P}$ and suboptimal arm $x' \in \tilde{X}(P)$, we get $\text{REW}_{\max}(x', P) = \max \{ \sum_{r \in \mathfrak{R}} rQ(r) : Q \in \mathcal{P}_{x'} \}$. For any deceitful arm x' , we define a distribution $Q_{x'}^*(x) = P(x)$ for all $x \in X \setminus \{x'\}$ and let*

$$\begin{aligned} Q_{x'}^*(x') &\in \arg \min_Q I(P(x'), Q) \\ \text{s.t. } & Q \in \mathcal{P}_{x'}, \sum_{r \in \mathfrak{R}} rQ(r) \geq \text{REW}^*(P). \end{aligned}$$

The constructed reward distribution $Q_{x'}^$ can be interpreted as the worst deceitful reward distribution for deceitful arm x' . That is, this distribution is the optimal solution to problem (9). To see why, first of all note that $Q_{x'}^*$ is a feasible solution to problem (9). Second of all, the objective function of this problem at $Q_{x'}^*$ is given by $\eta(x') \cdot I(P(x'), Q_{x'}^*(x'))$, which by definition of $Q_{x'}^*(x')$ is the minimum value it can take. Next, we argue that the objective value at the worst distribution is bounded away from zero. This is so because x' is a deceitful arm and as a result, there exists a deceitful distribution Q such that $\sum_{r \in \mathfrak{R}} rQ(r, x') > \text{REW}^*(P)$. This observation guarantees that $I(P(x'), Q_{x'}^*(x')) > 0$. This is so because per Lemma 11, presented in Appendix B, $I(P(x'), Q_{x'}^*(x')) = 0$ would imply that $P(x') = Q_{x'}^*(x')$, which contradicts the fact that x' is a suboptimal arm. By Proposition 1 and the fact that $I(P(x'), Q_{x'}^*(x')) > 0$, the regret lower bound is hence finite and reduces to $C(P) = \sum_{x \in \tilde{X}_d(P)} \Delta(x, P) / I(P(x), Q_x^*(x)) < \infty$, where each term in*

the sum quantifies the minimum regret necessary to denounce each deceitful arm as sub-optimal. To see why note that the constraint $\text{DIST}(\eta, x', P) \geq 1$ in the regret lower bound problem demands that $\eta(x') \geq 1/I(P(x'), Q_{x'}^*(x'))$ for each deceitful arm. This leads to $\eta^*(x', P) = 1/I(P(x'), Q_{x'}^*(x'))$ for $x' \in \tilde{X}_d(P)$ and $\eta^*(x', P) = 0$ for $x' \in \tilde{X}_n(P)$ is the minimal solution in the minimization problem (5) defining the lower regret bound for all suboptimal arms $x' \in \tilde{X}(P)$.

3. *Lipschitz Bernoulli bandits.* For Lipschitz Bernoulli bandits, the lower bound presented in Proposition 1 coincides precisely with the results first proven in Magureanu et al. (2014b). It can be readily verified that here $\text{REW}_{\max}(x', P) = \min(\max(\mathfrak{R}), \text{REW}^*(P) + L \cdot D(x', x^*(P)))$. Hence, all suboptimal arms are deceitful when $\text{REW}^*(P) < \max(\mathfrak{R})$; otherwise if $\text{REW}^*(P) = \max(\mathfrak{R})$, then trivially all suboptimal arms are non-deceitful and consequently $C(P) = 0$. To see why note that arm x' is deceitful if $\text{REW}_{\max}(x', P) > \text{REW}^*(P)$. Then, when $\text{REW}^*(P) < \max(\mathfrak{R})$, we certainly have $\text{REW}_{\max}(x', P) = \min(\max(\mathfrak{R}), \text{REW}^*(P) + L \cdot D(x', x^*(P))) > \text{REW}^*(P)$, which shows that any suboptimal arm x' is deceitful. On the other hand, when $\text{REW}^*(P) = \max(\mathfrak{R})$, $\text{REW}_{\max}(x', P) = \min(\max(\mathfrak{R}), \text{REW}^*(P) + L \cdot D(x', x^*(P))) = \max(\mathfrak{R}) \leq \text{REW}^*(P)$, which shows that any suboptimal arm x' is not deceitful. We hence focus on the first more interesting case. We then define the following reward distribution for any suboptimal arm x'

$$Q_{x'}^*(1, x) = \max(P(1, x), \text{REW}^*(P) - L \cdot D(x, x')), \quad Q_{x'}^*(0, x) = 1 - Q_{x'}^*(1, x),$$

for all $x \in X$; see also Figure 1. As you can see in this figure, $Q_{x'}^*(1, x)$ (red dots) matches the true distribution $P(1, x)$ (blue dots) everywhere except when x is close to x' . In other words, $Q_{x'}^*(1, x)$ is intentionally distorted around x' . Note that in this figure, $D(x, x') = |x - x'|$. The constructed reward distribution $Q_{x'}^*$ can again be interpreted as the worst deceitful reward distribution for arm x' as we now have that $\text{DIST}(\eta, x', P) = \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q_{x'}^*(x))$. By Proposition 1, therefore, the regret lower bound reduces in this particular context to an ordinary linear optimization problem

$$\begin{aligned} C(P) &= \min_{\eta \geq 0} \sum_{x \in \tilde{X}(P)} \eta(x) \Delta(x, P) \\ \text{s.t.} \quad & 1 \leq \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q_{x'}^*(x)) \quad \forall x' \in \tilde{X}(P), \end{aligned} \tag{10}$$

which can be solved using a standard linear optimization method such as the simplex algorithm, c.f., Dantzig (1998), but in general does not admit a closed form solution. ■

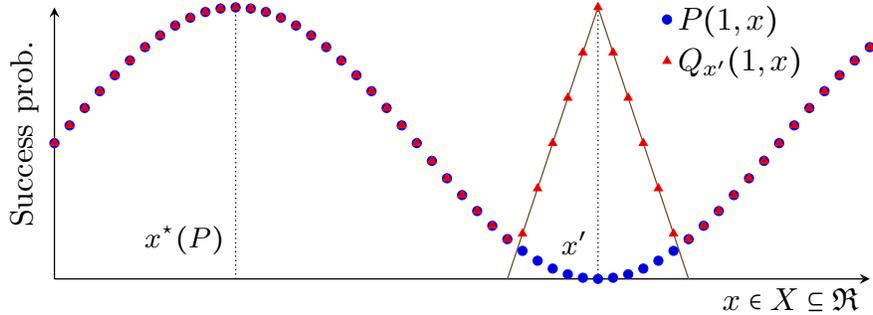


Figure 1: Worst-case deceitful reward distribution $Q_{x'}$ for a deceitful arm x' in a Lipschitz Bernoulli bandit problems with reward distribution P and optimal arm $x^*(P)$. Here, the worst-case deceitful reward distribution is independent of the logarithmic rate η with which we explore suboptimal arms. This observation, which is particular to the Lipschitz Bernoulli bandit problem and does not hold generally, guarantees that its regret lower bound minimization problem (5) reduces to an ordinary linear optimization problem.

5 High-Level Ideas of Our Policy

Our policy, which is called DUal Structure-based Algorithm (DUSA), is inspired by the lower bound formulation discussed in the previous section. The lower bound problem (5) suggests that to obtain minimal regret, each suboptimal arm x in $\tilde{X}(P)$ should be explored at least $N_{t+1}(x) = \eta_\epsilon(x, P) \cdot \log(t)$ times. The logarithmic exploration rate $\eta_\epsilon(P)$ is an arbitrary $0 < \epsilon$ -suboptimal solution, i.e., $\eta_\epsilon(P)$ is a feasible solution to problem (5) and achieves an objective value at most ϵ worse than $C(P)$, instead of an actual minimizer. We consider near optimal solutions to circumvent the technical difficulty that the infimum in (5) may not be attained. We will discuss later that considering such near optimal exploration rates $\eta_\epsilon(P)$ is desirable even if actual minimizers do exist.

Following the previously discussed strategy is evidently not possible without knowing the true reward distribution P . Our idea is to keep track of the reward distribution of each arm with the help of the empirical counterpart P_t constructed based on the historical rewards observed before round t . The policy then ensures that the number of times each arm x is pulled during the first t rounds, which is denoted by $N_{t+1}(x)$, is close to $\eta_\epsilon(x, P_t) \cdot \log(t)$ where the logarithmic rate $\eta_\epsilon(P_t)$ is an $0 < \epsilon$ -suboptimal solution of the following optimization problem

$$\begin{aligned}
 C(P_t) = \inf_{\eta} \quad & \sum_{x \in \tilde{X}(P_t)} \eta(x) \Delta(x, P_t) \\
 \text{s.t.} \quad & \eta \geq 0, \quad 1 \leq \text{DIST}(\eta, x, P_t) \quad \forall x \in \tilde{X}_d(P_t).
 \end{aligned} \tag{11}$$

In the following, we first discuss why and when mimicking the regret lower bound leads to a good learning policy. This discussion leads to a mild assumption on the true reward distribution P . Afterward, we argue that in general mimicking the regret lower bound, i.e., solving the semi-infinite problem (11), is computationally expensive. Motivated by this, we present a computationally-efficient convex dual-based approach instead.

5.1 When/Why Mimicking the Lower Bound Works

The proposed strategy aims to play suboptimal arms at the logarithmic target rate $\eta_\epsilon(P_t)$ in any round t . Such a strategy can attain asymptotic optimality when by converging the empirical distribution P_t to the true reward distribution P , (i) the empirical counterpart of the regret lower bound, i.e., $C(P_t)$, approaches the actual regret lower bound, i.e., $C(P)$, and (ii) the empirical (suboptimal) logarithmic target rate $\eta_\epsilon(P_t)$ converges to $\eta_\epsilon(P)$. To meet these requirements, at the reward distribution P , the regret lower bound function $C(Q)$ should be continuous and the ϵ -suboptimal exploration rates $\eta_\epsilon(Q)$ should admit a continuous selection at P . We note that $\eta_\epsilon(Q)$ may be a set-valued mapping⁴ and as a result, it is necessary to design a selection rule that chooses one of the suboptimal solutions while ensuring that the designed rule is continuous in the reward distribution.

The non-uniqueness of ϵ -suboptimal solutions, although at first glance perhaps perceived as an undesirable complication, is in fact a blessing in disguise. It indeed enables the existence of a continuous selection and echos our previous point that considering ϵ -suboptimal solutions is beneficial as the exact minimizers $\eta^*(Q)$ (should they exist) may fail to admit such continuous selection as our running example illustrates.

Example 3 (Continued). *We discuss the lack of continuity of the minimizer $\eta^*(Q)$ in the context of Lipschitz Bernoulli bandits.*

3. *Lipschitz Bernoulli bandits. We have previously determined $\eta^*(Q)$ in the context of Lipschitz Bernoulli Bandits as the minimizers in the parametric linear optimization problems stated in Equation (10). However, it is well known (Bank et al. 1982, Section 4.3) that the minimizers in such parametric linear optimization problems are upper semi-continuous but not necessarily lower semi-continuous mappings. In Figure 2, we attempt to illustrate the lack of continuous exploration rate selections from such upper semi-continuous mappings $\eta^*(Q)$. In this figure, at the reward distribution P , the optimal exploration rate η is not*

⁴We indeed expect multiple ϵ -suboptimal solutions to exist in problem (11).

unique and can be any point in a vertical line.⁵ The illustrated optimal mapping $\eta^*(Q)$ is upper semi-continuous, but not lower semi-continuous. For this mapping, one cannot choose one of the optimal solutions at P that leads to a continuous selection. However, this figure shows that this problem is alleviated by considering an ϵ -suboptimal mapping instead. Here, the green area illustrates all the ϵ -suboptimal solutions ($\eta_\epsilon(Q)$) and the red curve ($\eta'_\epsilon(Q)$) presents a particular continuous selection for the ϵ -suboptimal mapping.

The following mild assumption on the reward distribution P is necessary to guarantee the necessary continuity properties.

Assumption 1. *The reward distribution P (i) has a unique optimal arm, (ii) satisfies $\text{REW}_{\max}(x, P) \neq \text{REW}^*(P)$ for any $x \in X$, and (iii) is in the interior of \mathcal{P} . Here, $\text{REW}_{\max}(x, P)$ is defined in Equation (7) and $\text{REW}^*(P) = \sum_{r \in \mathfrak{R}} rP(r, x^*(P))$. Let $\mathcal{P}' \subseteq \mathcal{P}$ denote the set of all such distributions.*

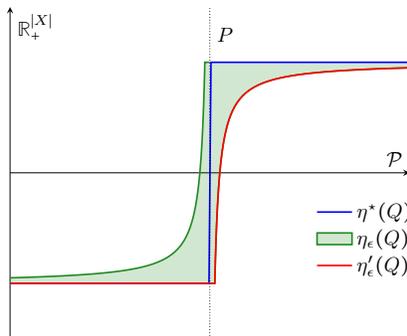


Figure 2: The mapping $\eta^*(Q)$ between a reward distribution Q and the set of all optimal exploration rates is typically merely upper semi-continuous but not necessarily lower semi-continuous. Even though the optimal exploration rates may exist, there exists no continuous selection at the reward distribution P where the optimal exploration rates are not uniquely defined. Considering ϵ -suboptimal solutions for $\epsilon > 0$ alleviates this problem and generalizes to situations where no optimal exploration rates may exist at all.

The first condition in Assumption 1 requires the optimal arm under the true reward distribution to be unique. The second condition, i.e., $\text{REW}_{\max}(x, P) \neq \text{REW}^*(P)$ for any $x \in X$, implies that there exists a neighborhood around the true reward distribution P such that for any distribution Q within this neighborhood, the sets of deceitful arms $\tilde{X}_d(Q)$,

⁵Having multiple optimal solutions to the regret lower problem does not necessarily imply that we have multiple optimal arms. To see why, consider an example with one optimal arm and two suboptimal arms whose expected rewards are identical. Suppose that the learner is aware of the fact that the expected reward of two arms is the same. In this case, the optimal solution to the regret lower problem may not be unique. However, the optimal arm is unique.

non-deceitful arms $\tilde{X}_n(Q)$, and optimal arm $x^*(Q)$ are the same. Recall that the suboptimal arm $x \in \tilde{X}_d(P)$ if $\text{REW}_{\max}(x, P) > \text{REW}^*(P)$ and the suboptimal arm $x \in \tilde{X}_d(P)$ if $\text{REW}_{\max}(x, P) \leq \text{REW}^*(P)$. As we show at the end of this section via counter example, both of these conditions are essential for the regret lower bound $C(P)$ as well its (suboptimal) solutions $\eta_\epsilon(P)$ to be continuous. The last assumption that $P \in \text{int}(\mathcal{P})$ ensures that our estimator P_t eventually realizes in the set of all potential reward distributions \mathcal{P} . That is, we have asymptotically $\lim_{t \rightarrow \infty} \text{Prob}[P_t \in \mathcal{P}] = 1$. We believe that this last assumption is not critical in that is tied to the naive empirical estimator P_t that does not exploit the structural information. As a more sophisticated estimator would make the regret analysis of our policy far more complicated, we do not explore this approach in this work and leave it as a future direction.

Proposition 3 (Continuity of Regret Lower Bound). *Consider the regret lower bound problem (5) characterizing the regret lower bound function. Then, (i) the regret lower bound function $C(Q)$ is continuous on \mathcal{P}' , and (ii) there exists an $0 < \epsilon$ -suboptimal rate selection $\eta_\epsilon^c(Q)$ which is continuous on \mathcal{P}' . Recall that $\mathcal{P}' \subseteq \mathcal{P}$ is the set of all distributions that satisfy Assumption 1.*

The proof of this result, along with other continuity results, is presented in Appendix G. We now argue that the first two conditions in Assumption 1 are necessary by considering a simple example. Consider a separable Bernoulli bandit problem with only two arms $X = \{a, b\}$. The reward distribution of the first arm a is such that $P(0, a) \geq \frac{2}{5}$ and the reward distribution of the second arm b can be arbitrary. That is, \mathcal{P} is given in Equation (8). For any $\lambda \in [0, 1]$, we define parameterized reward distributions $P_\lambda(a) = [\frac{1}{2}, \frac{1}{2}]$ and $P_\lambda(b) = [1 - \lambda, \lambda]$. That is, the probability of receiving reward zero under arm a (respectively arm b) is $\frac{1}{2}$ (respectively $1 - \lambda$). Then, it is easy to see that the optimal arms and reward are given by

$$x^*(P_\lambda) = \begin{cases} \{b\} & \text{if } \lambda \in (\frac{1}{2}, 1], \\ \{a, b\} & \text{if } \lambda = \frac{1}{2}, \\ \{a\} & \text{if } \lambda \in [0, \frac{1}{2}) \end{cases} \quad \text{and} \quad \text{REW}^*(P_\lambda) = \begin{cases} \lambda & \text{if } \lambda \in (\frac{1}{2}, 1], \\ \frac{1}{2} & \text{if } \lambda \in [0, \frac{1}{2}]. \end{cases}$$

Observe that at $\lambda = \frac{1}{2}$, the optimal arm is not unique. That is, the first condition of Assumption 1 is violated. Furthermore, at $\lambda = \frac{3}{5}$, we have $\text{REW}_{\max}(a, P) = \text{REW}^*(P)$. That is, the second condition of Assumption 1 is violated. Next, we show that these violations result in discontinuity in the regret lower bound.

For the parameter values $\lambda \in [\frac{3}{5}, 1]$, we have that $\text{REW}^*(P_\lambda) \geq \text{REW}_{\max}(a, P_\lambda) = \frac{3}{5}$ and hence $C(P_\lambda) = 0$. Likewise, when $\lambda = \frac{1}{2}$, both arms are optimal, and once again we have that

$C(P_\lambda) = 0$. It can be verified that

$$C(P_\lambda) = \begin{cases} 0 & \text{if } \lambda \in [\frac{3}{5}, 1] \cup \{\frac{1}{2}\}, \\ \frac{\lambda - \frac{1}{2}}{\frac{1}{2} \log(\frac{1/2}{\lambda}) + \frac{1}{2} \log(\frac{1/2}{1-\lambda})} & \text{if } \lambda \in (1/2, \frac{3}{5}), \\ \frac{1/2 - \lambda}{\lambda \log(\frac{\lambda}{1/2}) + (1-\lambda) \log(\frac{1-\lambda}{1/2})} & \text{if } \lambda \in [0, 1/2), \end{cases}$$

which we have visualized in Figure 3. The function $C(P_\lambda)$ is evidently discontinuous at parameter value λ of $\frac{1}{2}$ and $\frac{3}{5}$, as claimed.

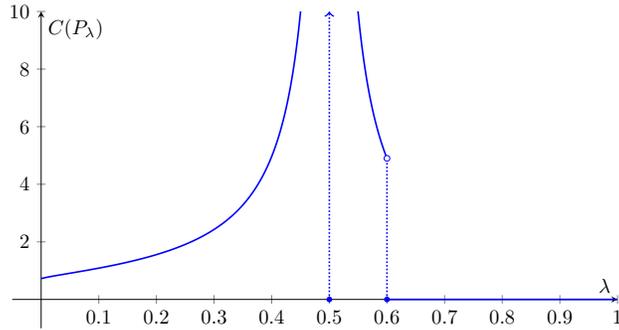


Figure 3: The regret lower bound $C(P_\lambda)$ for the separable bandit problem discussed in Section 5.1. The regret lower bound is discontinuous at $\lambda = \frac{1}{2}$ and $\lambda = \frac{3}{5}$.

5.2 Dual-based Representation of Regret Lower Bound

To follow the proposed strategy, we may potentially need to solve the regret lower bound problem, i.e., problem (11), in each round. Observe that this optimization problem is a semi-infinite minimization problem, which cannot in general be solved directly. Even verifying the feasibility of a fixed exploration rate η requires the solution to the auxiliary information minimization problem (6) defining our information distance function $\text{DIST}(\eta, x, P_t)$. We will overcome this challenge by using an equivalent dual representation of this distance function instead. This novel equivalent dual representation enables us to reformulate the regret lower bound problem (5) and its empirical counterpart problem (11) as standard finite convex minimization problems, which are computationally much more attractive compared to solving their semi-infinite representations directly.

We start with stating an equivalent formulation of the distance function with the help of conic hull of the feasible set \mathcal{P} . Let us denote the conic hull of \mathcal{P} as the set $\mathcal{K} = \text{cone}(\mathcal{P}) := \{\theta \cdot Q : \theta \geq 0, Q \in \mathcal{P}\}$. Observe that this convex cone represents the same information as the set of reward distributions \mathcal{P} as we have the following equivalence: $\mathcal{P} = \{Q : \sum_{x \in X, r \in \mathfrak{R}} Q(r, x) = |X|\} \cap \mathcal{K}$.

Working with the cone \mathcal{K} instead of the set \mathcal{P} helps us provide a unified dual formulation later on. This conic representation, along with Proposition 2, allows us to rewrite the information distance function for any deceitful arm $x' \in \tilde{X}_d(P)$ as the following minimization problem

$$\begin{aligned} \text{DIST}(\eta, x', P) = \min_Q \quad & \sum_{x \in \tilde{X}(P)} I(P(x), Q(x)) + \chi_\infty(P(x^*(P)) = Q(x^*(P))) \\ \text{s.t.} \quad & Q \in \mathcal{K}, & [\text{dual variable: } \lambda \in \mathcal{K}^*] \\ & \sum_{x \in X, r \in \mathfrak{R}} Q(r, x) = |X|, & [\text{dual variable: } \beta \in \mathbb{R}] \\ & \sum_{r \in \mathfrak{R}} rQ(r, x') \geq \text{REW}^*(P), & [\text{dual variable: } \alpha \in \mathbb{R}_+] \end{aligned} \quad (12)$$

where $\chi_\infty(A) = 0$ if event A occurs; $+\infty$ otherwise. Note that in problem (12), all the nonlinear constraints in the minimization problem (6) are captured by the conic constraint $Q \in \mathcal{K}$. We denote with λ the dual variable of the cone constraint, $\beta \in \mathbb{R}$. Further, $\alpha \in \mathbb{R}_+$ are the dual variables of the second and third constraints, respectively. Let the dual cone \mathcal{K}^* be defined as the unique cone that satisfies $Q \in \mathcal{K}$ if and only if $\langle \lambda, Q \rangle \geq 0$, $\forall \lambda \in \mathcal{K}^*$, where we take as inner product $\langle \lambda, Q \rangle = \sum_{x \in X, r \in \mathfrak{R}} \lambda(r, x)Q(r, x)$. Just as the cone \mathcal{K} captures the same information as the set of reward distributions \mathcal{P} , so does its dual cone \mathcal{K}^* . This is so because for any closed convex cones, we have $\mathcal{K}^{**} = \mathcal{K}$; see Boyd and Vandenberghe (2004, Section 2.6.1).

For any reward distribution $P \in \mathcal{P}$ and suboptimal arm $x' \in \tilde{X}(P)$, we now define our concave dual function:

$$\begin{aligned} \text{DUAL}(\eta, x', P ; \mu) := & \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \eta(x) \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha r \mathbb{1}(x = x')}{\eta(x)} \right) P(r, x) - \sum_{x \in x^*(P), r \in \mathfrak{R}} \lambda(r, x) P(r, x) \\ & + \alpha \text{REW}^*(P) + \beta |\tilde{X}(P)| + \chi_{-\infty}(\eta(x) - \lambda(r, x) - \beta - \alpha r \mathbb{1}(x = x') \geq 0). \end{aligned} \quad (13)$$

For notational convenience, we collect all our dual variables in the dual vector $\mu := (\alpha, \beta, \lambda) \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*$. Here, the concave characteristic function $\chi_{-\infty}(A)$ takes on the value zero when event A happens; $-\infty$ otherwise. The following lemma characterizes the dual of the distance function DIST using the DUAL function defined above.

Lemma 1 (Dual Formulation of the Distance Function). *For any reward distribution $P \in \mathcal{P}_\Omega$ and arm $x' \in X$, the following weak duality inequality $\text{DUAL}(\eta, x', P ; \mu) \leq \text{DIST}(\eta, x', P)$ holds, where $\mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*$ is any feasible dual vector to problem (12), the dual function DUAL and distance function DIST are defined in Equations (13) and (6), respectively. Furthermore, for any reward distribution $P \in \mathcal{P}$ and suboptimal deceitful arm $x' \in \tilde{X}_d(P)$, the following*

strong duality equality holds

$$\text{DIST}(\eta, x', P) = \sup_{\mu} \{ \text{DUAL}(\eta, x', P ; \mu) : \mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^* \}. \quad (14)$$

In the following theorem, we characterize the regret lower bound function using the dual formulation in Lemma 1. The proof of this lemma is presented in Appendix C.1. Note that all the duality results are presented in Appendix C.

Theorem 2 (Regret Lower Bound – Dual Formulation). *Consider any reward distribution $P \in \mathcal{P}$. The regret lower bound presented in Proposition 1 is characterized equivalently as*

$$\begin{aligned} C(P) &= \inf_{\eta \geq 0, \mu} \sum_{x \in \tilde{X}(P)} \eta(x) \Delta(x, P) \\ \text{s.t. } & \mu(x) \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^* & x \in X, \\ & 1 \leq \text{DUAL}(\eta, x, P ; \mu(x)) & x \in \tilde{X}_d(P). \end{aligned} \quad (15)$$

Observe that the dual characterization of the regret lower bound is convex and finite dimensional. We further highlight that the dual problem (15), unlike its primal counterpart (5), is no longer a nested optimization. The complexity and computational issues of solving the dual lower bound representation are further discussed in Appendix H.

6 Dual Structure Algorithm (DUSA)

In this section, we present our policy, called DUSA; see Algorithm 1. We shortly discuss each of its parts and give additional context and commentary. The DUSA algorithm is parameterized by a positive number $\epsilon > 0$ which we shall refer to as the accuracy parameter. Our optimal algorithm is best understood by considering its asymptotic limit $\epsilon \rightarrow 0$. Hence, most of our commentary regarding the intuition behind the algorithm will focus on this regime.

Initialization phase. DUSA starts with an initialization phase of $|X|$ rounds. During these first $|X|$ rounds, it pulls each arm $x \in X$ once. Then, based on the observed rewards in these rounds, it initializes its estimate for the reward distribution. As before, the estimate P_t is the empirical distribution of the observed rewards before round t . At the end of the initialization phase, i.e., round $t = |X|$, DUSA initializes the dual variables for each arm x , i.e., $\alpha_{t+1}(x)$, $\beta_{t+1}(x)$, and $\lambda_{t+1}(x)$, and the *reference logarithmic rate*, denoted by $\eta'_{t+1}(x)$.

Sufficient information and resolving test. After the initialization phase, DUSA aims

Algorithm 1 DUal Structure Algorithm (DUSA)

Input. Accuracy parameter $0 < \epsilon$.

• **Initialization.**

- During the first $|X|$ rounds, pull each arm $x \in X$ once and set $P_{|X|+1}(R(x), x) = 1$, where $R(x)$ is the observed reward of arm x in its corresponding round.
- Let $s_{|X|+1} = 0$ and $N_{|X|+1}(x) = 1$ for all $x \in X$.
- For all arms $x \in X$, initialize the dual variables $\lambda_{|X|+1}(x) \in \mathcal{K}^*$, $\alpha_{|X|+1}(x) \in \mathbb{R}_+$, and $\beta_{|X|+1}(x) \in \mathbb{R}$ and reference logarithmic rate $\eta'_{|X|+1}(x) > 0$.

• For $t \in |X| + 1, \dots$

- **Sufficient information condition.** For every deceitful arm $x \in \tilde{X}_d(P_t)$, solve the following convex univariate minimization problem

$$\overline{\text{DUAL}}_t(x) := \overline{\text{DUAL}}(N_t/\log(t), x, P_t; \mu_t(x)) := \max_{\rho \geq 0} \text{DUAL}(N_t/\log(t), x, P_t; \rho\mu_t(x)).$$

- **Exploitation.** If $\overline{\text{DUAL}}_t(x) \geq 1 + \epsilon$ for all deceitful arms $x \in \tilde{X}_d(P_t)$, then $s_{t+1} = s_t$, $\mu_{t+1} = \mu_t$, $\eta'_{t+1} = \eta'_t$, and pull a least played empirically optimal arm

$$x_t \in \arg \min \{N_t(x') : x' \in \arg \max_{x \in X} \sum_{r \in \mathfrak{R}} r P_t(r, x)\}.$$

- **Exploration.** Else, set $s_{t+1} = s_t + 1$. Furthermore,
 - * Update exploration target rates $\eta_t = \text{SU}(P_t, \eta'_t; \mu_t, \epsilon)$ using shallow update Algorithm 2.
 - * If $\min_{x \in X} N_t(x) \leq \epsilon s_t$, pull the least pulled arm

$$x_t = \underline{x}_t \in \arg \min_{x \in X} N_t(x).$$

- * Else, pull

$$x_t = \bar{x}_t \in \arg \min_{x \in \tilde{X}(P_t)} N_t(x)/\eta_t(x).$$

- * Update the dual variables $(\eta'_{t+1}, \mu_{t+1}) = \text{DU}(P_t; \epsilon)$ using deep update Algorithm 3.

– **Updating variables.**

- * Observe reward of the pulled arm $R_t(x_t)$. For any $x \in X$, $x \neq x_t$, and $r \in \mathfrak{R}$, set $P_{t+1}(r, x) = P_t(r, x)$. Further, for any $r \neq R_t(x_t)$, $P_{t+1}(r, x_t) = N_t(x_t)P_t(r, x_t)/(N_t(x_t) + 1)$, and $P_{t+1}(R_t(x_t), x_t) = (1 + N_t(x_t)P_t(R_t(x_t), x_t))/(N_t(x_t) + 1)$.
 - * Set $N_{t+1}(x_t) = N_t(x_t) + 1$ and for any $x \neq x_t$, $N_{t+1}(x) = N_t(x)$.
-

to strike a balance between exploration and exploitation. To do so, it checks if for any empirically deceitful arm $x \in \tilde{X}_d(P_t)$, the following sufficient information condition holds:

$$1 + \epsilon \leq \overline{\text{DUAL}}_t(x) := \overline{\text{DUAL}}\left(\frac{N_t}{\log(t)}, x, P_t; \mu_t(x)\right) := \max_{\rho \geq 0} \text{DUAL}\left(\frac{N_t}{\log(t)}, x, P_t; \rho \cdot \mu_t(x)\right). \quad (16)$$

Here, $\mu_t(x) = (\alpha_t(x), \beta_t(x), \lambda_t(x))$ are the dual variables for arm x employed in round t . Remark that verifying the sufficient information condition merely requires the resolution of a univariate convex optimization problem. We refer to previously defined function $\overline{\text{DUAL}}$ as the *dual-test* function. Roughly speaking, our condition verifies whether sufficient information has been collected to distinguish the empirical P_t from its deceitful distributions; that is, whether or not $\text{DIST}(N_t/\log(t), x', P_t) \geq \overline{\text{DUAL}}(N_t/\log(t), x', P_t; \mu_t) \geq 1 + \epsilon$ for any $x' \in \tilde{X}(P_t)$. In Section 6.1, we provide a further clarifying interpretation of this sufficient information test as the information distance between P_t and a half-space containing all deceitful distributions $\text{DECEIT}(x', P)$. As we will show later, our sufficient information condition helps DUSA solve the lower bound problem (15) only in $O(\log(T))$ rounds.

Exploitation. If the sufficient information condition holds for all deceitful arms $x \in \tilde{X}_d(P_t)$, DUSA enters an exploitation phase and it selects a least played empirically optimal arm, i.e., $x_t \in \arg \min \{N_t(x') : x' \in \arg \max_{x \in X} \sum_{r \in \mathfrak{R}} r P_t(r, x)\}$.

Exploration. If the sufficient information test fails, DUSA enters an exploration phase as insufficient information has been collected on the empirically suboptimal arms $\tilde{X}(P_t)$ in the previous rounds. Let s_t be the number of exploration rounds in the first $t - 1$ rounds. Then, if $\min_{x \in X} N_t(x) \leq \epsilon s_t$, DUSA pulls the least pulled arm \underline{x}_t . Otherwise, it considers a new 2ϵ -suboptimal logarithmic target rate η_t in

$$\begin{aligned} \overline{C}(P_t; \mu_t) &:= \inf_{\eta \geq 0} \sum_{x \in \tilde{X}(P_t)} \eta(x) \Delta(x, P_t) \\ \text{s.t. } &1 \leq \overline{\text{DUAL}}(\eta, x, P_t; \mu_t(x)) \quad \forall x \in \tilde{X}_d(P_t). \end{aligned} \quad (17)$$

This infimum can be computed efficiently because it is characterized as the minimum of a linear objective over an ordinary convex optimization constraint set. However, as 2ϵ -suboptimal solutions are not unique, a particular selection must be made. To do so, the shallow update (SU) Algorithm 2 selects the exploration rate closest to some reference logarithmic rate η'_t . As the optimization characterization (18) is strictly convex, its minimizer when feasible must be unique. Our selection procedure thus circumvents the potential pitfall that the infimum (17) may not in fact admit a minimizer and as outlined in Proposition 5 enjoys much better continuity properties even if this infimum is in fact attained. Let $\eta_t = \text{SU}(P_t, \eta'_t; \mu_t, \epsilon)$ be the output of the shallow update Algorithm. Then, DUSA pulls a

suboptimal arm $\bar{x}_t \in \arg \min_{x \in \tilde{X}(P_t)} N_t(x)/\eta_t(x)$ whose empirical logarithmic exploration rate deviates from its optimal target rate to the largest relative extent.

Algorithm 2 Shallow Update $SU(Q, \eta' ; \mu, \epsilon)$

Input. Reward distribution Q , reference exploration target rate η' , dual variable $\mu = (\alpha, \beta, \lambda)$, and accuracy parameter $\epsilon > 0$. Compute

$$\begin{aligned} \eta_\epsilon(Q) := \arg \min_{\eta \geq 0} \quad & \sum_{x \in X} (\eta(x) - \eta'(x))^2 \\ \text{s.t.} \quad & 1 \leq \overline{\text{DUAL}}(\eta, x, Q ; \mu(x)) \quad \forall x \in \tilde{X}_d(Q), \\ & \sum_{x \in \tilde{X}(Q)} \eta(x) \Delta(x, Q) \leq \overline{C}(Q ; \mu) + 2\epsilon. \end{aligned} \tag{18}$$

If no feasible solution exists, we simply set $\eta_\epsilon(Q) := +\infty$.

Return. $\eta_\epsilon(Q)$.

Finally, DUSA updates the reference rate η'_t and dual variables $(\eta'_t, \mu_t) = \text{DU}(P_t, \epsilon)$ using the deep update (DU) Algorithm 3. The deep update algorithm returns a ϵ -suboptimal solution of the dual counterpart of the lower bound problem (15) computed at the empirical reward distribution P_t . As the dual counterpart of the lower bound problem (15) is an convex optimization problem, such ϵ -suboptimal solutions can be readily determined. However, for reasons discussed previously, these solutions are not unique and a particular selection must be made. Our deep update algorithm selects a certain minimal norm solution over a slightly restricted feasible set. As the optimization formulation (19) in the deep update algorithm is strictly convex, its minimizer is always unique. Note that the first and second sets of constraints in problem (19) are the same as those in problem (15). This ensures that any feasible solution to problem (19) is a feasible solution to problem (15). The third set of constraints of problem (19) guarantees any feasible solution to this problem to be an ϵ -suboptimal solution to problem (15). Note that the optimal value $C(Q)$, which is appeared in the third constraint, can be obtained simply by solving an ordinary convex optimization problem. The fourth set of constraints of problem (19) ensures η to be bounded away from zero and helps us establish continuity of the optimal solution to problem (19). The last set of constraints of this problem guarantees the dual function DUAL to be finite.⁶ We would like to stress that, thanks to our sufficient information test, we do not carry out either the deep or shallow update in every round.

As discussed earlier, it is critical to have a continuous selection rule as the ϵ -suboptimal solutions are not unique. The subsequent proposition points out that our deep update selection rule is indeed well-defined and continuous on \mathcal{P}' as required.

⁶The dual function, which is defined in Equation (13), is written as a function of the logarithm of variable $\omega = \eta(x) - \lambda(r, x) - \beta(x) - \alpha(x)r\mathbb{1}(x = x')/\eta(x)$. The last set of constraints of problem (19) insures that the ω 's are bounded away from zero and consequently the dual functions are finite.

Algorithm 3 Deep Update DU(Q ; ϵ)

Input. Reward distribution Q and accuracy parameter $\epsilon > 0$. Let

$$\begin{aligned} \eta'_\epsilon(Q), \mu_\epsilon(Q) &:= \arg \min_{\eta, \mu} \sum_{x \in X} \eta(x)^2 + \sum_{x \in X} \|\mu(x)\|_2^2 \\ \text{s.t. } \eta(x) &\in \mathbb{R}_+, \mu(x) = (\alpha(x), \beta(x), \lambda(x)) \in \mathcal{K}^* \times \mathbb{R}_+ \times \mathbb{R} \quad \forall x \in X, \\ 1 &\leq \text{DUAL}(\eta, x, Q ; \mu(x)) \quad \forall x \in \tilde{X}_d(Q), \\ \sum_{x \in \tilde{X}(Q)} \eta(x) \Delta(x, Q) &\leq C(Q) + \epsilon, \\ \eta(x) &\geq \epsilon / (2 \sum_{x \in \tilde{X}(Q)} \Delta(x, Q)) \quad \forall x \in \tilde{X}(Q), \\ \eta(x) &\geq \lambda(r, x)(x') + \beta(x') + \alpha(x') \cdot r \mathbb{1}(x = x') \\ &\quad + \epsilon / (2 \sum_{x \in \tilde{X}(Q)} \Delta(x, Q)) \quad \forall x' \in \tilde{X}_d(Q), x \in \tilde{X}(Q). \end{aligned} \tag{19}$$

Return. $(\eta'_\epsilon(Q), \mu_\epsilon(Q))$.

Proposition 4 (Continuous Selection Rule for the Deep Update Algorithm). *Our deep update selection rule $(\eta'_\epsilon(Q), \mu_\epsilon(Q))$, presented in Equation (19), exists and is continuous on \mathcal{P}' , where \mathcal{P}' consists of all the distributions in \mathcal{P} that satisfy Assumption 1.*

The proof of this technical result is provided in Appendix G.3. To establish our main result stated in Theorem 3, we also require that the shallow update selection rule to be stable. The shallow update is stable in the sense that $\text{SU}(P_t, \eta'_\epsilon(P_{t'}); \mu_\epsilon(P_{t'}), \epsilon)$ converges to the desired reference rate $\eta'_\epsilon(P)$ when both empirical distributions P_t and $P_{t'}$ approach the true reward distribution P . This notion of stability for the shallow update is formalized in the following proposition and will be crucial to establish the optimality of our DUSA policy. The proof of this technical result is referred to Appendix G.4.

Proposition 5 (Stability of the Shallow Update Selection Rule). *For any $P \in \mathcal{P}'$, we have*

$$\begin{aligned} \lim_{\kappa \rightarrow 0} \max_{Q_1, Q_2} \|\text{SU}(Q_1, \eta'_\epsilon(Q_2); \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(P)\|_2 &= 0. \\ \text{s.t. } \|Q_1 - P\|_\infty \leq \kappa, \|Q_2 - P\|_\infty \leq \kappa \end{aligned}$$

We now present our main result: regret bound of DUSA.

Theorem 3 (Regret Bound of DUSA). *Let the true reward distribution $P \in \mathcal{P}'$ and consider any accuracy parameter $0 < \epsilon < 1/|X|$. Then, DUSA suffers the following asymptotic regret:*

$$\limsup_{T \rightarrow \infty} \frac{\text{REG}_\pi(T, P)}{\log(T)} \leq C(P)(1 + \epsilon) + \epsilon \left(1 + (1 + \epsilon)|X|^2 (\|\eta'_\epsilon(P)\|_\infty + \epsilon) + \sum_{x \in X} \Delta(x, P) \right), \tag{20}$$

where the regret is computed against the benchmark that knows the best arm in advance. Furthermore, the expected number of rounds in which DUSA enters the exploration phase is bounded by

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[s_T]}{\log(T)} \leq (1 + \epsilon) |X|^2 (\|\eta'_\epsilon(P)\|_\infty + \epsilon), \quad (21)$$

where the expectation is taken with respect to the randomness in the observed rewards and choices made by DUSA.

The maximum target exploration rate $\|\eta'_\epsilon(P)\|_\infty = \max_{x \in \tilde{X}(P)} \eta'(x, P)$ appearing in both bounds in Theorem 3 is uniformly bounded for all accuracy parameter values $\epsilon > 0$. This is so because (i) we have from ϵ -suboptimality that $\sum_{x \in \tilde{X}(P)} \eta'_\epsilon(x, P) \Delta(x, P) \leq C(P) + \epsilon < \infty$, and (ii) $\Delta(x, P) > 0$ for all suboptimal arms. Then, comparing the regret lower bound in Proposition 1 with the regret upper bound found in Equation (20), we can hence conclude that DUSA enjoys the minimal asymptotic logarithmic regret as the accuracy parameter ϵ goes to zero. Furthermore, the number of times we have to update the target exploration rates and dual variables using either the shallow or deep update Algorithms 2 and 3 grows merely logarithmically in the number of rounds T as indicated by Equation (21). Note that the number of updates is equal to the number of the exploration rounds. We note that Theorem 3 does not bound the regret of DUSA for any finite number of rounds but instead is completely asymptotic in nature. Nevertheless, in the proof of Theorem 3, which we will present in Section 7, non-asymptotic regret bounds are also derived.

6.1 More on the Sufficient Information Condition

As stated earlier, the dual-test function provides a convenient procedure to verify if enough information has been collected already to distinguish the empirical P_t from its deceitful models; that is, whether or not $\text{DIST}(N_t/\log(t), x', P_t) \geq \overline{\text{DUAL}}(N_t/\log(t), x', P_t; \mu_t) \geq 1 + \epsilon$ for any deceitful arm $x' \in \tilde{X}_d(P_t)$. The strong duality results in Lemma 1, relates the information distance between P_t and $\text{DIST}(x', P_t)$ and the dual function as

$$\text{DIST}(N_t/\log(t), x', P_t) = \max_{\mu(x')} \{\text{DUAL}(N_t/\log(t), x', P_t; \mu(x')) : \mu(x') \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*\}.$$

Our sufficient information test however involves instead a seemingly unrelated restricted dual problem

$$\begin{aligned}
& \overline{\text{DUAL}}(N_t/\log(t), x', P_t ; \mu_t) \\
&= \max_{\rho \geq 0, \mu(x')} \{ \text{DUAL}(N_t/\log(t), x', P_t ; \mu(x')) : \rho \geq 0, \mu(x') = \rho \cdot \mu_t(x') \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^* \} \\
&= \max_{\rho \geq 0} \{ \text{DUAL}(N_t/\log(t), x', P_t ; \rho \cdot \mu_t(x')) : \rho \geq 0 \}.
\end{aligned}$$

Nevertheless, we will argue here that also this restricted dual problem admits a satisfying interpretation as the information distance condition between P_t and a certain set of deceitful distributions. Indeed, for any feasible dual variable $\mu = (\alpha, \beta, \lambda)$, arbitrary empirical reward distribution $P_t \in \mathcal{P}$, and deceitful arm $x' \in \tilde{X}_d(P_t)$, we define $\mathcal{H}(x', P_t ; \mu)$ as the following set of reward distributions:

$$\left\{ Q : \begin{array}{l} Q(x^*(P_t)) = P_t(x^*(P_t)), \\ \sum_{x \in X, r \in \mathfrak{R}} Q(r, x) (\lambda(r, x) + \beta + \alpha r \mathbb{1}(x = x')) \geq \beta |X| + \alpha \text{REW}^*(P_t) \end{array} \right\}. \quad (22)$$

We point out that the previously defined set contains all deceitful distributions whenever the dual variable μ is dual feasible; that is, $\mathcal{H}(x', P_t ; \mu) \supseteq \text{DECEIT}(x', P_t)$, where the set of deceitful distributions $\text{DECEIT}(x', P_t)$ is presented as the feasible set of the convex minimization problem (12). Obviously, any deceitful distribution will satisfy $Q(x^*(P_t)) = P_t(x^*(P_t))$ by construction. That is, the reward distributions of the empirically optimal arms of P_t are in full agreement with those of Q . We now show that any deceitful distribution will also satisfy the second constraint $\sum_{x \in X, r \in \mathfrak{R}} Q(r, x) (\lambda(r, x) + \beta + \alpha r \mathbb{1}(x = x')) \geq \beta |X| + \alpha \text{REW}^*(P_t)$ characterizing the set $\mathcal{H}(x', P_t ; \mu)$. The first term $\sum_{x \in X, r \in \mathfrak{R}} Q(r, x) \lambda(r, x)$ is nonnegative because $Q \in \mathcal{K}$ and dual feasible variable $\lambda \in \mathcal{K}^*$. The second term satisfies $\sum_{x \in X, r \in \mathfrak{R}} \beta Q(r, x) = \beta |X|$ irrespective of the value of the dual variable β ; see the second constraints in the minimization problem (12). For any dual feasible variable $\alpha \geq 0$, the final term satisfies the inequality $\sum_{x \in X, r \in \mathfrak{R}} Q(r, x) r \mathbb{1}(x = x') \geq \text{REW}^*(P_t)$; see the third constraints of minimization Problem (12). Thus, we have indeed the claimed inclusion $\mathcal{H}(x', P_t ; \mu) \supseteq \text{DECEIT}(x', P_t)$. The following lemma provides an interpretation of the sufficient information condition.

Lemma 2 (Dual-test Function). *Let $P \in \text{int}(\mathcal{P})$ and $\eta > 0$ and consider any deceitful arm $x' \in \tilde{X}_d(P)$. For any dual feasible variable $\mu = (\alpha, \beta, \lambda) \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*$, the dual-test function*

is equal to

$$\begin{aligned} \overline{\text{DUAL}}(\eta, x', P ; \mu) &= \min_Q \sum_{x \in \tilde{\mathcal{X}}(P)} \eta(x) I(P(x), Q(x)) \\ \text{s.t. } Q &\in \mathcal{H}(x', P ; \mu). \end{aligned} \tag{23}$$

Proof. See Appendix C.3. □

The optimization problem (23), which states an alternative characterization of our dual-test function, bears some resemblance to the minimization problem (6), which characterizes the distance function DIST . While the objective function of two optimization problems is the same, in the minimization Problem (6), we enforce the distribution Q to belong to the deceitful set $\text{DECEIT}(x', P)$, whereas in the minimization problem (23), we enforce the same distribution belong to merely belong to the superset $\mathcal{H}(x', P ; \mu)$. Thus, the condition $1 \leq \overline{\text{DUAL}}(\eta, x', P ; \mu)$ implies that sufficient information has been obtained to distinguish the null hypothesis P not only from its deceitful distributions $\text{DECEIT}(x', P)$ but in fact from the entire superset $\mathcal{H}(x', P ; \mu)$. The simple observation that distinguishing P from the larger set $\mathcal{H}(x', P ; \mu)$ must necessarily be harder from a statistical point of view than distinguishing P merely from the smaller set of deceitful distributions $\text{DECEIT}(x', P)$ explains the weak duality inequality $\overline{\text{DUAL}}(\eta, x', P ; \mu) \leq \text{DIST}(\eta, x', P)$ for any feasible dual variable μ . Strong duality can be interpreted as the existence of a feasible dual variable μ^* so that the distance between P and the sets $\mathcal{H}(x', P ; \mu^*)$ and $\text{DECEIT}(x', P)$ is the same. That is, $\overline{\text{DUAL}}(\eta, x', P ; \mu^*) = \text{DIST}(\eta, x', P) = \sum_{x \in \tilde{\mathcal{X}}(P)} \eta(x) I(P(x), Q^*(x))$ for some worst-case deceitful reward distribution $Q^* \in \text{DECEIT}(x', P)$. Under those circumstances, the hyperplane $\mathcal{H}(x', P ; \mu^*)$ separates the deceitful reward distributions $\text{DECEIT}(x', P)$ and the information ball $\{Q : Q(x^*(P)) = P(x^*(P)), \sum_{x \in \tilde{\mathcal{X}}(P)} \eta(x) I(P(x), Q(x)) \leq \text{DIST}(\eta, P, x')\}$ as implied by the necessary and sufficient KKT optimality conditions of minimization problem (6) and visually illustrated in Figure 4.

The previous discussion also offers additional insight into the deep and shallow update algorithms introduced before. The deep update Algorithm 3 minimizes the regret over all exploration rates η and dual variables so that the reward distribution P can be distinguished from its deceitful distributions $\text{DECEIT}(x', P)$ with sufficient statistical power; that is, the deep update ensures that $\text{DIST}(\eta, x', P) \geq 1$. The shallow update, on the other hand, merely minimizes the regret over the logarithmic rates η so that the reward distribution P can be distinguished from the superset $\mathcal{H}(x', P ; \mu)$ associated with the dual variable μ . That is, the shallow update ensures that $\overline{\text{DUAL}}(\eta, x', P ; \mu) \geq 1$.

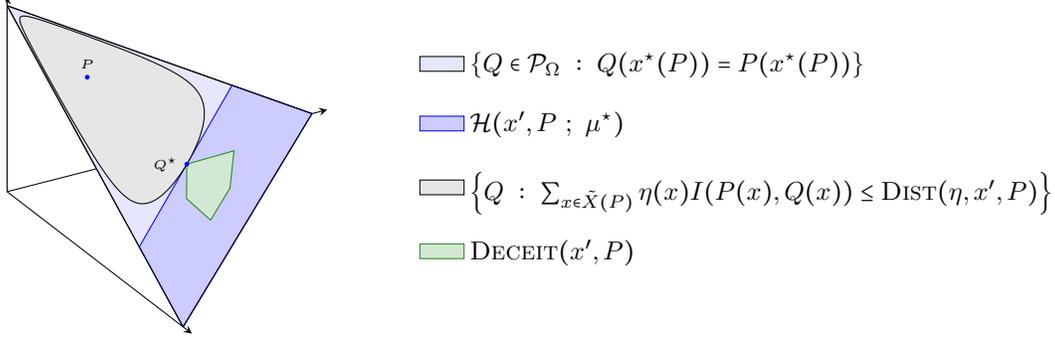


Figure 4: The sufficient information test $\overline{\text{DUAL}}(\eta, x', P; \mu) \leq \text{DIST}(\eta, x', P)$ for dual feasible variable μ quantifies the information distance between P and the half-space $\mathcal{H}(x', P; \mu')$ containing all deceitful distributions $\text{DECEIT}(x', P)$. When $\mathcal{H}(x', P; \mu^*)$ for some μ^* defines a separating half-space between $\text{DECEIT}(x', P)$ and the information ball $\{Q : Q(x^*(P)) = P(x^*(P)), \sum_{x \in \hat{X}(P)} \eta(x)I(P(x), Q(x)) \leq \text{DIST}(\eta, P, x)\}$, the information distance between P and either the set of deceitful distributions $\text{DECEIT}(x', P)$ or its superset $\mathcal{H}(x', P; \mu^*)$ is the same. In that case, we have the strong duality equalities $\overline{\text{DUAL}}(\eta, x', P; \mu^*) = \text{DIST}(\eta, x', P) = \sum_{x \in \hat{X}(P)} \eta(x)I(P(x), Q^*(x))$, where Q^* is the worst-case distribution, solving problem (23).

7 Proof of Theorem 3: Regret Bound of DUSA

We first present our concentrations bound and explain how the regret of DUSA can be decomposed.

7.1 Concentration Bounds and Regret Decomposition

We start with presenting our supporting concentration bounds. These bounds will be used multiple times throughout the proof and their proofs are presented in Appendix D.

Lemma 3 (Concentration Bound). *For any arm $x \in X$, let $t(s, x) \in [s, \dots, T + 1]$ be any stopping time such that either $N_{t(s, x)}(x) \geq \epsilon s/4$ or $t(s, x) = T + 1$. Then, for any $r \in \mathfrak{X}$, we have*

$$\text{Prob} \left[\left| P_{t(s, x)}(r, x) - P(r, x) \right| \geq \kappa, t(s, x) \leq T \right] \leq 2 \exp(-s\epsilon\kappa^2/2),$$

where for any $\tau \in [T]$, $N_\tau(x)$ is the number of rounds within the first τ rounds that arm x is played.

The proof of Lemma 3 is inspired by the proof of Lemma 4.3 in Combes and Proutiere (2014).

Lemma 4 (Concentration Bound for Information Distance). *Let $\delta \geq |X|(|\mathfrak{X}| - 1) + 1$. Then,*

for any $t \geq 1$,

$$\text{Prob} \left[\sum_{x \in X} N_t(x) I(P_t(x), P(x)) \geq \delta \right] \leq e \cdot \left(\frac{\delta [\log(t)\delta + 1] 2e}{|X|(|\mathfrak{X}| - 1)} \right)^{|X|(|\mathfrak{X}| - 1)} \cdot \exp(-\delta).$$

The proof of concentration bound in Lemma 4, which is one of our contributions, is quite involved and requires a careful decomposition of the reward distributions for which we construct a martingale sequence. Our concentration bound bears resemblance to the bound in Magureanu et al. (2014b), where they bound $\text{Prob} [\sum_{x \in X} N_t(x) I(P_t(x), P(x)) \geq \delta]$ for Binomial random variables. We generalize their bounds to arbitrary discrete distributions, which may of independent interest for non-parametric multi-armed bandit problems.

Regret decomposition. Having presented the concentration bounds, we now focus on upper bounding the regret of DUSA. The regret can be broken down into three parts: Regret during (i) the initialization phase, (ii) the exploitation phase and (iii) the exploration phase. That is, for any $T > |X|$, we have

$$\text{REG}_\pi(T, P) = \mathbb{E} \left[\sum_{1 \leq t \leq |X|} \Delta(x_t, P) \right] + \mathbb{E} \left[\sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t\} \Delta(x_t, P) \right] + \mathbb{E} \left[\sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{X}_t\} \Delta(x_t, P) \right],$$

where $\Delta(x, P)$ is the suboptimality gap of arm x under the actual reward distribution P . Here, we define events \mathcal{E}_t and \mathcal{X}_t as the event that DUSA is the exploitation and exploration phase, respectively, in a given round t . Evidently, we have that $\mathcal{E}_t \cup \mathcal{X}_t = \{t > |X|\}$. In the initialization phase, we select each arm $x \in X$ once. The regret caused in this phase is evidently finite and equal to $\mathbb{E} \left[\sum_{1 \leq t \leq |X|} \Delta(x_t, P) \right] = \sum_{x \in X} \Delta(x, P)$. Note that this regret is upper bounded by $|X|$ as the regret in each round is bounded by one. In the following sections, each of the other two terms will be bounded.

7.2 Exploitation Phase

To characterize the regret during the exploitation phase, we start with defining the following bad event $\mathcal{B}_t^1(x, \kappa) = \{\|P_t(x) - P(x)\|_\infty > \kappa\}$ for some positive threshold $\kappa > 0$ and arm $x \in X$. The bad event occurs when the empirical reward distribution $P_t(x)$ and the true reward distribution $P(x)$ of the considered arm x deviate significantly. In other words, the reward distribution of arm x is poorly estimated. The following lemma characterizes the regret accumulated during the exploitation phase due to a poorly estimated empirically optimal arm; i.e., when the events \mathcal{E}_t and $\bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)$ occur simultaneously. In the following lemma, we show that the regret accumulated during the first T rounds under these events is

negligible for any arbitrary small $\kappa > 0$. As a part of the proof, we invoke the concentration bound in Lemma 3.

Lemma 5 (Regret under the Bad Events in the Exploitation Phase). *The regret accumulated during the exploitation phase when the bad event $\bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)$ occurs remains finite; that is,*

$$\mathbb{E} \left[\sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t, \bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x_t, P) \right] \leq 8|X| + |X| |\mathfrak{R}| \sum_{1 \leq s \leq T} 2 \exp(-s\kappa^2/(2|X|)) \quad (24)$$

for any $\kappa > 0$.

Next, we assume that we are in the exploitation phase and our estimate of the reward distributions of the empirically optimal arms are accurate, i.e., $\bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)$ does not occur. Recall that we assumed that P has a unique optimal arm. By Lemma 7, presented at the end of this section, this implies that we can take κ small enough such that either $x^*(P) = x^*(P_t)$ or $x^*(P) \notin x^*(P_t)$. The regret under the former event is zero. Thus, we only bound the regret under the latter event:

$$\begin{aligned} & \mathbb{E} \left[\sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t, \bigcap_{x \in x^*(P_t)} \neg \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x, P) \right] \\ &= \mathbb{E} \left[\sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t, x^*(P) \notin x^*(P_t), \bigcap_{x \in x^*(P_t)} \neg \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x, P) \right] \\ &\leq \sum_{|X| < t \leq T} \text{Prob} \left[\mathcal{E}_t, x^*(P) \notin x^*(P_t), \|P_t(x) - P(x)\|_\infty \leq \kappa \quad \forall x \in x^*(P_t) \right], \end{aligned}$$

where the inequality holds because $\Delta(x, P) \leq 1$. Here, for any event E , $\neg E$ denotes its complement.

To further bound the regret, we take advantage of our sufficient information test. If the event \mathcal{E}_t occurs, we must be in the exploitation phase and thus the sufficient information test (16) should have been passed by every empirically deceitful arm in $\tilde{X}_d(P_t)$. By the weak duality inequality in Lemma 1, this implies that $1 + \epsilon \leq \overline{\text{DUAL}}_t(x') \leq \text{DIST}(N_t/\log(t), x', P_t)$ for any empirically deceitful arm x' , and hence by definition of the distance function DIST in (6), the event

$$\sum_{x \in \tilde{X}(P_t)} N_t(x) I(P_t(x), Q(x)) \geq (1 + \epsilon) \log(t) \quad \forall Q \in \text{DECEIT}(x', P_t), \quad \forall x' \in \tilde{X}_d(P_t) \quad (25)$$

must have occurred. By (25), if we had $P \in \text{DECEIT}(x^*(P), P_t)$ and $x^*(P) \in \tilde{X}_d(P_t)$, then

$$\sum_{x \in X} N_t(x) I(P_t(x), P(x)) \geq (1 + \epsilon) \log(t). \quad (26)$$

Under (26), Lemma 4 implies that the total regret caused by such events is finite uniformly in T . Unfortunately, the actual reward distribution P may not be deceitful, i.e., P may not belong to $\text{DECEIT}(x^*(P), P_t)$. Nonetheless, we prove that the event defined in (26) does occur as long as κ is chosen sufficiently small.

Lemma 6. *For small enough $\kappa > 0$, if the bad event i.e., $\bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)$, does not occur, then the event defined in Equation (26), must occur.*

The total regret

$$\sum_{1 \leq t \leq T} \text{Prob}[\mathcal{E}_t, x^*(P) \notin x^*(P_t), \|P_t(x) - P(x)\|_\infty \leq \kappa \quad \forall x \in x^*(P_t)]$$

accumulated here is upper bounded by $\sum_{1 \leq t \leq T} \text{Prob}[\sum_{x \in X} N_t(x) I(P_t(x), P(x)) \geq (1 + \epsilon) \log(t)]$. Let $t > t_0$ with $\log(t_0)(1 + \epsilon) > |X|(|\mathfrak{R}| - 1) + 1$. Then, by Lemma 4, we have

$$\begin{aligned} & \text{Prob}[\sum_{x \in X} N_t(x) I(P_t(x), P(x)) \geq (1 + \epsilon) \log(t)] \\ & \leq \exp(-(1 + \epsilon) \log(t)) \cdot e \cdot \left(\frac{\delta[\log(t)\delta + 1]2e}{|X|(|\mathfrak{R}| - 1)} \right)^{|X|(|\mathfrak{R}| - 1)} = \frac{e}{t^{1+\epsilon}} \cdot \left(\frac{\delta[\log(t)\delta + 1]2e}{|X|(|\mathfrak{R}| - 1)} \right)^{|X|(|\mathfrak{R}| - 1)}. \end{aligned}$$

Hence, the regret in the event of interest, i.e., when the bad event $\bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)$ does not happen, is bounded as

$$\mathbb{E} \left[\sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t, \bigcap_{x \in x^*(P_t)} \neg \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x_t, P) \right] \leq t_0 + \sum_{t_0 \leq t \leq T} \frac{e}{t^{1+\epsilon}} \cdot \left(\frac{\delta[\log(t)\delta + 1]2e}{|X|(|\mathfrak{R}| - 1)} \right)^{|X|(|\mathfrak{R}| - 1)}.$$

Thus, the total regret in the exploitation phase is upper bounded by

$$C_{\text{exploit}} := 8|X| + |X||\mathfrak{R}| \sum_{1 \leq s \leq T} 2 \exp(-s\kappa^2/(2|X|)) + t_0 + \sum_{t_0 \leq t \leq T} \frac{e}{t^{1+\epsilon}} \left(\frac{\delta[\log(t)\delta + 1]2e}{|X|(|\mathfrak{R}| - 1)} \right)^{|X|(|\mathfrak{R}| - 1)},$$

which remains finite for any T .

Lemma 7 (Stability of the Optimal Arm and Continuity of the Maximal Reward). *Let $\mathcal{P}_u(x_0^*)$ be the set of all reward distributions $P \in \mathcal{P}_\Omega$ with a unique optimal arm x_0^* . Then, (i) the set $\mathcal{P}_u(x_0^*)$ is an open set, and (ii) the maximum reward $\text{REW}_{\max}(x', Q)$ of any arm $x' \in X$ is a continuous function at $P \in \text{int}(\mathcal{P})$.*

7.3 Exploration Phase

To characterize the regret during the exploration phase, we first bound the probability that our estimate of reward distributions in two consecutive exploration rounds is bad and we then bound the regret when our estimate of reward distributions in two consecutive exploration rounds is indeed good. Let $\kappa > 0$ be defined as in the previous section. We say our estimate at the end of round t is bad, i.e., \mathcal{B}_t^2 occurs, when $\|P_t - P\|_\infty > \kappa$. Here, $\|P_t - P\|_\infty = \max_{r,x} |P_t(r,x) - P(r,x)|$. We a slight abuse of notation, we define the stopping time $\tau(s)$ as the first time in the first T rounds in which we are in the exploration phase, i.e., \mathcal{X}_t occurs, and $s_t = s$. If no such round exists, we set $\tau(s) = T + 1$.

Regret during bad exploration rounds. We now bound the regret caused by events in which either now or in the last time we were in the exploration phase our estimate of the reward distributions is bad.

Lemma 8. *If $\mathcal{X}_t \cap \{s_t > 2/\epsilon\}$ occurs, then we must have $\min_{x \in X} N_t(x) \geq \epsilon \lceil s_t/4 \rceil$.*

When event $\mathcal{X}_t \cap \{s_t > 2/(\epsilon)\} \cap \mathcal{B}_t^2$ happens, (i) we are in the exploration phase, (ii) $\|P_t - P\|_\infty \geq \kappa$, and (iii) $\min_{x \in X} N_t(x) \geq \epsilon \lceil s_t/4 \rceil$. Our regret under the aforementioned events can be written as:

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{1}\{\mathcal{B}_{\tau(1)}^2, \tau(1) \leq T\} \cdot \Delta(x_{\tau(1)}, P) + \sum_{2 \leq s \leq T} \mathbb{1}\{(\mathcal{B}_{\tau(s-1)}^2 \cup \mathcal{B}_{\tau(s)}^2), s \leq \tau(s) \leq T\} \cdot \Delta(x_{\tau(s)}, P) \right] \\
& \leq 1 + \sum_{2 \leq s \leq T} \text{Prob} \left[(\mathcal{B}_{\tau(s-1)}^2 \cup \mathcal{B}_{\tau(s)}^2), s \leq \tau(s) \leq T \right] \leq 1 + 2 \sum_{1 \leq s \leq T} \text{Prob} \left[\mathcal{B}_{\tau(s)}^2, s \leq \tau(s) \leq T \right] \\
& \leq 1 + 2 \left(\sum_{1 \leq s < T} \text{Prob} [\mathbb{1}\{s \leq 2/\epsilon\}] + \sum_{1 \leq s \leq T} \text{Prob} [\mathbb{1}\{s > 2/\epsilon\}, \|P_{\tau(s)} - P\|_\infty \geq \kappa, s \leq \tau(s) \leq T] \right) \\
& \leq 1 + \frac{4}{\epsilon} + 4|X| |\mathfrak{R}| \sum_{1 \leq s \leq T} \exp(-s\epsilon\kappa^2/2) \leq 1 + \frac{4}{\epsilon} + 8 \frac{|X| |\mathfrak{R}|}{\epsilon\kappa^2} := C_2,
\end{aligned}$$

where the third inequality follows from Lemmas 3 and 8.

Upper bound on the number of good exploration rounds. In a good exploration round, clearly, we are in the exploration phase and the event \mathcal{B}_t^2 does not occur now nor did occur the last time we were in the exploration phase. Recall that event \mathcal{B}_t^2 does not happen when at the end of round t , $\|P_t - P\|_\infty \leq \kappa$. From Lemma 7 and Assumption 1, we can consider a sufficiently small κ such that when the event \mathcal{B}_t^2 does not occur, we have $x^*(P_t) = x^*(P)$, $\tilde{X}_d(P_t) = \tilde{X}_d(P)$ and $\tilde{X}_n(P_t) = \tilde{X}_n(P)$. Thus, under those aforementioned fortunate circumstances, we correctly identified each arm as either optimal, deceitful or undeceitful both now and in the previous exploration round. We will denote this fortunate event with \mathcal{G}_t . We define $r(x) = \max\{|X| < t \leq T : x_t = x, \mathcal{G}_t\}$ as the most recent exploration

round before time T in which the good event \mathcal{G}_t occurred and arm x is played. We set $r(x) = 0$ if no such event occurred. Consider the number $\sum_{t=1}^T \mathbb{1}\{x_t = x, \mathcal{X}_t, \mathcal{G}_t\}$ of good exploration rounds in which some arm $x \in X$ is played. Evidently, as $r(x)$ is the most recent time before time T in which the good event \mathcal{G}_t occurred and arm x is played, we must have that $\sum_{t=1}^T \mathbb{1}\{x_t = x, \mathcal{X}_t, \mathcal{G}_t\} = \sum_{t=1}^{r(x)} \mathbb{1}\{x_t = x, \mathcal{X}_t, \mathcal{G}_t\} \leq \sum_{t=1}^{r(x)} \mathbb{1}\{x_t = x\} = N_{r(x)}(x)$. We bound this further using the following lemma.

Lemma 9. *During an exploration round t , we have $N_t(x_t) \leq (1 + \epsilon)\|\eta_t\|_\infty \log(T)$.*

By Lemma 9, we have the following bound on the number of good exploration rounds:

$$\begin{aligned} \sum_{t=1}^T \mathbb{1}\{\mathcal{X}_t, \mathcal{G}_t\} &= \sum_{t=1}^T \sum_{x \in X} \mathbb{1}\{x_t = x, \mathcal{X}_t, \mathcal{G}_t\} = \sum_{x \in X} \sum_{t=1}^T \mathbb{1}\{x_t = x, \mathcal{X}_t, \mathcal{G}_t\} \\ &\leq \sum_{x \in X} N_{r(x)}(x) \leq \sum_{x \in X} (1 + \epsilon)\|\eta_{r(x)}\|_\infty \log(T). \end{aligned}$$

Recall that in rounds $r(x)$ we have that the estimate $P_{r(x)}$ and the estimate of the exploration round preceding $r(x)$ are both close to P . Because of the stability result in Proposition 5, we can take $0 < \kappa$ sufficiently small so that the target rates $\eta_{r(x)}$ computed by the shallow update algorithm in such good rounds satisfies $\|\eta_{r(x)} - \eta'_\epsilon(P)\|_\infty \leq \epsilon$. That is, these rounds come with sufficiently accurate estimates to guarantee that the exploration rates set by the shallow update algorithm are very close to the desired exploration rate $\eta'_\epsilon(P)$. Consequently, we obtain the bound

$$\sum_{t=1}^T \mathbb{1}\{\mathcal{X}_t, \mathcal{G}_t\} \leq (1 + \epsilon)|X|(\|\eta'_\epsilon(P)\|_\infty + \epsilon) \log(T). \quad (27)$$

Note that the inequality holds even when $r(x) = 0$. When $r(x) = 0$, we have $\sum_{t=1}^T \mathbb{1}\{x_t = x, \mathcal{X}_t, \mathcal{G}_t\} = \sum_{t=1}^{r(x)} \mathbb{1}\{x_t = x, \mathcal{X}_t, \mathcal{G}_t\} = 0$.

Regret under good exploration rounds. We are now left with the regret caused by rounds in which we are in the exploration phase and in the reward model P was estimated accurately both in the current round and the previous time we were not in the exploitation phase. That is, we need to bound the regret caused when the good events \mathcal{G}_t occur. Let

$$W_T^1 = \sum_{t=1}^T \Delta(x, P) \mathbb{1}\{x_t = \underline{x}_t, \mathcal{X}_t, \mathcal{G}_t\} \quad \text{and} \quad W_T^2 = \sum_{t=1}^T \Delta(x, P) \mathbb{1}\{x_t = \bar{x}_t, \mathcal{X}_t, \mathcal{G}_t\}.$$

We first upper bound the expected regret $\mathbb{E}[W_T^1]$ when the good event \mathcal{G}_t happens and arm \underline{x}_t is played in the exploration phase. Consider the number of times $\sum_{t=1}^T \mathbb{1}\{x_t = \underline{x}_t = x, \mathcal{X}_t, \mathcal{G}_t\}$ we play arm x in such events. Recall that when $x_t = \underline{x}_t$ and \mathcal{G}_t occurs, we must have that

$N_t(\underline{x}_t) \leq \epsilon s(t)$. Let $\underline{r}(x)$ be the most recent time before round T that the event $x_t = \underline{x}_t = x$, \mathcal{G}_t occurred. If no such time exists than we set $\underline{r}(x) = 0$. Evidently, we have that

$$\sum_{t=1}^T \mathbb{1}\{x_t = \underline{x}_t = x, \mathcal{X}_t, \mathcal{G}_t\} = \sum_{t=1}^{\underline{r}(x)} \mathbb{1}\{x_t = \underline{x}_t = x, \mathcal{X}_t, \mathcal{G}_t\} \leq \sum_{t=1}^{\underline{r}(x)} \mathbb{1}\{x_t = \underline{x}_t = x\} \leq N_{\underline{r}(x)}(x) \leq \epsilon s_{\underline{r}(x)} \leq \epsilon s_T.$$

We then have

$$\begin{aligned} W_T^1 &= \sum_{x \in X} \sum_{t=1}^T \Delta(x, P) \mathbb{1}\{x_t = \underline{x}_t = x, \mathcal{X}_t, \mathcal{G}_t\} \leq \sum_{x \in X} \sum_{t=1}^T \mathbb{1}\{x_t = \underline{x}_t = x, \mathcal{X}_t, \mathcal{G}_t\} \leq \epsilon |X| s_T \\ &\leq |X| \epsilon (\sum_{\tau=1}^T \mathbb{1}\{\mathcal{X}_\tau, \mathcal{G}_\tau\} + \sum_{\tau=1}^T \mathbb{1}\{\mathcal{X}_\tau, -\mathcal{G}_\tau\}) \\ &\leq |X| \epsilon ((1 + \epsilon) |X| (\|\eta'_\epsilon(P)\|_\infty + \epsilon) \log(T) + \sum_{\tau=1}^T \mathbb{1}\{\mathcal{X}_\tau, -\mathcal{G}_\tau\}), \end{aligned}$$

where the last inequality follows from Equation (27). Taking expectations gives us

$$\mathbb{E} [W_T^1] \leq \epsilon |X| ((1 + \epsilon) |X| (\|\eta'_\epsilon(P)\|_\infty + \epsilon) \log(T) + C_2),$$

where we use the previous result that the expected number of exploration bounds associated with bad estimates, i.e., $\mathbb{E} [\sum_{\tau=1}^T \mathbb{1}\{\mathcal{X}_\tau, -\mathcal{G}_\tau\}]$, is bounded in expectation by the constant C_2 defined previously.

We now upper bound the expected regret $\mathbb{E} [W_T^2]$. Consider the number of times $\sum_{t=1}^T \mathbb{1}\{x_t = \bar{x}_t = x, \mathcal{X}_t, \mathcal{G}_t\}$ we play arm x in such events. As we argued before, since we are in the exploration phase and $x_t = \bar{x}_t$, we must have that $N_t(x_t) \leq (1 + \epsilon) \eta_t(x_t) \log(T)$ where $\eta_t = \text{SU}(P_t, \eta'_t; \mu_t, \epsilon)$. (Recall that when we are in the exploration phase, there exists an empirically suboptimal arm $x' \in \tilde{X}(P_t)$ such that $N_t(x')/\eta_t(x') \leq (1 + \epsilon) \log(t)$). This implies that $N_t(x_t)/\eta_t(x_t) \leq (1 + \epsilon) \log(t)$ when $x_t = \bar{x}_t = \arg \min_{x \in \tilde{X}(P_t)} N_t(x)/\eta_t(x)$. Let $\bar{r}(x)$ be the most recent time before round T that the event $x_t = \bar{x}_t = x$, \mathcal{G}_t occurred. If no such time exists than we set $\bar{r}(x) = 0$. Evidently, we have that

$$\begin{aligned} \sum_{t=1}^T \mathbb{1}\{x_t = \bar{x}_t = x, \mathcal{X}_t, \mathcal{G}_t\} &\leq \sum_{t=1}^{\bar{r}(x)} \mathbb{1}\{x_t = \bar{x}_t = x, \mathcal{X}_t, \mathcal{G}_t\} \\ &\leq \sum_{t=1}^{\bar{r}(x)} \mathbb{1}\{x_t = \underline{x}_t = x\} = N_{\bar{r}(x)}(x) \leq (1 + \epsilon) \eta_{\bar{r}(x)}(x) \log(T). \end{aligned}$$

Thus,

$$\mathbb{E} [W_T^2] = \mathbb{E} \left[\sum_{x \in X} \Delta(x, P) \sum_{t=1}^T \mathbb{1}\{x_t = \bar{x}_t = x, \mathcal{G}_t\} \right] \leq \mathbb{E} \left[(1 + \epsilon) \sum_{x \in X} \Delta(x, P) \eta_{\bar{r}(x)}(x) \log(T) \right].$$

Because of the stability result in Proposition 5, we can take $0 < \kappa$ sufficiently small so that

the target rates $\eta_{r(x)}$ computed by the shallow update algorithm in such good rounds satisfies $\|\eta_{r(x)} - \eta'_\epsilon(P)\|_\infty \leq \epsilon$. Hence,

$$\mathbb{E} [W_T^2] \leq (1 + \epsilon) \sum_{x \in X} \Delta(x, P) (\eta'_\epsilon(x, P) + \epsilon) \log(T) \leq (1 + \epsilon) (C(P) + \epsilon + \epsilon \sum_{x \in X} \Delta(x, P)) \log(T),$$

where we exploit the fact that $\eta'_\epsilon(P)$ is an ϵ -suboptimal exploration rate. As $\kappa > 0$ can be chosen arbitrary small, and ignoring the regret constant terms, we obtain

$$\limsup_{T \rightarrow \infty} \frac{\text{REG}_\pi(T)}{\log(T)} \leq (1 + \epsilon) C(P) + \epsilon (1 + (1 + \epsilon) |X|^2 (\|\eta_\epsilon(P)\|_\infty + \epsilon) + \sum_{x \in X} \Delta(x, P)). \quad \square$$

8 Concluding Remarks

In this paper, we have presented the first bandit policy that optimally exploits convex structural information in a computationally feasible fashion. Our policy, DUSA, automatically exploits any known structural reward information with the help of convex constraints on the unknown reward distribution. Hence, our powerful policy can be used in many practical problems where structural information is often available but may be peculiar to a very specific problem. Rather than developing bandit policies for some idiosyncratic class of bandit problems, our paper presents a universally optimal algorithm. We believe that our DUSA algorithm presents a useful blueprint which can be extended in several directions.

References

- R. Agrawal, D. Teneketzis, and V. Anantharam. Asymptotically efficient adaptive allocation schemes for controlled iid processes: Finite parameter space. *IEEE Transactions on Automatic Control*, 34(3), 1989.
- J.-P. Aubin and H. Frankowska. *Set-Valued Analysis*. Springer Science & Business Media, 2009.
- S. Balseiro, N. Golrezaei, M. Mahdian, V. Mirrokni, and J. Schneider. Contextual bandits with cross-learning. In *Advances in Neural Information Processing Systems*, pages 9676–9685, 2019.
- B. Bank, J. Guddat, D. Klatte, B. Kummer, and K. Tammer. *Non-Linear Parametric Optimization*. Springer, 1982.
- A. Barvinok. *A Course in Convexity*, volume 54. American Mathematical Society, 2002.
- G. Bayraksan and D.K. Love. Data-driven stochastic programming using phi-divergences. In *The Operations Research Revolution*, pages 1–19. INFORMS, 2015.

- C. Berge. *Topological Spaces: Including a Treatment of Multi-Valued Functions, Vector Spaces, and Convexity*. Courier Corporation, 1997.
- D.P. Bertsekas. *Convex Optimization Theory*. Athena Scientific Belmont, 2009.
- O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, 2009.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends[®] in Machine Learning*, 5(1):1–122, 2012.
- S. Bubeck, N. Devanur, Z. Huang, and R. Niazadeh. Multi-scale online learning: Theory and applications to online auctions and pricing. *Journal of Machine Learning Research*, 2019.
- S. Burer. On the copositive representation of binary and continuous nonconvex quadratic programs. *Mathematical Programming*, 120(2):479–495, 2009.
- V. Chandrasekaran and P. Shah. Relative entropy optimization and its applications. *Mathematical Programming*, 161(1-2):1–32, 2017.
- R. Combes and A. Proutiere. Unimodal bandits: Regret lower bounds and optimal algorithms. In *International Conference on Machine Learning*, pages 521–529, 2014.
- R. Combes, S. Magureanu, and A. Proutiere. Minimal exploration in structured stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 1763–1771, 2017.
- P.L. Combettes. Perspective functions: Properties, constructions, and examples. *Set-Valued and Variational Analysis*, 26(2):247–264, 2018.
- T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.
- V. Dani, T.P. Hayes, and S.M. Kakade. Stochastic linear optimization under bandit feedback. In *International Conference on Algorithmic Learning Theory*, pages 355–366, 2008.
- G.B. Dantzig. *Linear programming and extensions*. Princeton University Press, 1998.
- E. Delage and Y. Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612, 2010.
- E. Erdoğan and G. Iyengar. Ambiguous chance constrained problems and robust optimization. *Mathematical Programming*, 107(1-2):37–61, 2006.
- A. Garivier and O. Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th annual conference on learning theory*, pages 359–376, 2011.
- N. Golrezaei, A. Javanmard, and V. Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. In *Advances in Neural Information Processing Systems*, pages 9756–9766, 2019.
- T.L. Graves and T.L. Lai. Asymptotically efficient adaptive choice of control laws in controlled Markov chains. *SIAM Journal on Control and Optimization*, 35(3):715–743, 1997.

- S. Gupta, S. Chaudhari, G. Joshi, and O. Yağın. Multi-armed bandits with correlated arms. *arXiv preprint arXiv:1911.03959*, 2019.
- E. Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- W. Hoeffding. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pages 409–426. Springer, 1994.
- E. Kaufmann, N. Korda, and R. Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *International Conference on Algorithmic Learning Theory*, pages 199–213. Springer, 2012.
- N.B. Keskin and A. Zeevi. Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. *Operations Research*, 62(5):1142–1167, 2014.
- T.L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- T. Lattimore and C. Szepesvári. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 728–737, Fort Lauderdale, FL, USA, 4 2017. PMLR.
- S. Magureanu, R. Combes, and A. Proutiere. Lipschitz bandits: Regret lower bounds and optimal algorithms. *arXiv preprint arXiv:1405.4758*, 2014a.
- S. Magureanu, R. Combes, and A. Proutiere. Lipschitz bandits: Regret lower bound and optimal algorithms. In Maria Florina Balcan, Vitaly Feldman, and Csaba Szepesvári, editors, *Proceedings of The 27th Conference on Learning Theory*, volume 35 of *Proceedings of Machine Learning Research*, pages 975–999, Barcelona, Spain, 2014b. PMLR.
- J. Mao, R. Leme, and J. Schneider. Contextual pricing for Lipschitz buyers. In *Advances in Neural Information Processing Systems*, pages 5643–5651, 2018.
- A.J. Mersereau, P. Rusmevichientong, and J.N. Tsitsiklis. A structured multiarmed bandit problem and the greedy policy. *IEEE Transactions on Automatic Control*, 54(12):2787–2802, 2009.
- P. Mohajerin Esfahani and D. Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1-2):115–166, 2018.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer Academic Publishers, 2004.
- Rad Niazadeh, Negin Golrezaei, Joshua Wang, Fransisca Susan, and Ashwinkumar Badanidiyuru. Online learning via offline greedy: Applications in market design and optimization. *Available at SSRN 3613756*, 2020.

- G. Pflug and D. Wozabal. Ambiguity in portfolio selection. *Quantitative Finance*, 7(4):435–442, 2007.
- P. Rusmevichientong and J. N. Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.
- D. Russo and B. Van Roy. Learning to optimize via information-directed sampling. *Operations Research*, 66(1):230–252, 2018.
- B. Schölkopf and A.J. Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- A. Slivkins. Contextual bandits with similarity information. In *Proceedings of the 24th annual Conference On Learning Theory*, pages 679–702, 2011.
- Matthew Streeter and Daniel Golovin. An online algorithm for maximizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 1577–1584, 2008.
- J.A.K. Suykens and J. Vandewalle. Least squares support vector machine classifiers. *Neural processing letters*, 9(3):293–300, 1999.
- W.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.
- B.P.G. Van Parys, P.M. Esfahani, and D. Kuhn. From data to decisions: Distributionally robust optimization is optimal. *arXiv preprint arXiv:1704.04118*, 2017.
- Z. Wang, P.W. Glynn, and Y. Ye. Likelihood robust optimization for data-driven problems. *Computational Management Science*, 13(2):241–261, 2016.
- Mingrui Zhang, Zebang Shen, Aryan Mokhtari, Hamed Hassani, and Amin Karbasi. One sample stochastic frank-wolfe. *arXiv preprint arXiv:1910.04322*, 2019.
- C. Zhao and Y. Guan. Data-driven risk-averse stochastic optimization with wasserstein metric. *Operations Research Letters*, 46(2):262–267, 2018.

A Further Examples

The following (non-exhaustive) list of structured bandit problems are subsumed by our flexible structured bandit framework.

Linear bandits. A rather popular class of structured bandit problems, c.f, (Dani et al. 2008, Rusmevichientong and Tsitsiklis 2010, Mersereau et al. 2009, Lattimore and Szepesvari 2017), assumes that the expected rewards of the arms $x \in X \subseteq \mathbb{R}^n$ are linearly related. That is, there exists an unobserved vector $a \in \mathbb{R}^n$ so that $\sum_{x \in X} (\mathbb{E}[R_t(x)] - a^\top x)^2/2 + \|a\|_2^2/(2\gamma) \leq \delta$ where $\gamma > 0$ and $\delta > 0$ allows for a small deviation from linearity. We can capture this structure with the help of the set

$$\mathcal{P}_{\text{lin}}(\delta) := \{Q \in \mathcal{P}_\Omega : a \in \mathbb{R}^n, \sum_{x \in X} (\sum_{r \in \mathcal{R}} rQ(r, x) - a^\top x)^2/2 + \|a\|_2^2/(2\gamma) \leq \delta\}.$$

and remark that it is closed for any regularization parameter $\gamma > 0$ as it is recognized as the projection of a bounded convex set; see also Bertsekas (2009, Proposition 1.3.6). Note that for the set \mathcal{P}_{lin} to have non-empty interior, it suffices to require the mild assumption that $\delta \in \text{int}(\{\delta' \geq 0 : \mathcal{P}_{\text{lin}}(\delta') \neq \emptyset\})$.

Polynomial bandits. One can extend the previous class of multi-armed bandit problems by allowing arms to be related through a polynomial of degree d . That is, there exists some unobserved vector $c \in \mathbb{R}^f$ of dimension $f := \binom{d+n}{n}$ so that $\sum_{x \in X} (\mathbb{E}[R_t(x)] - \sum_{i=1}^f c_i m_i(x))^2/2 + \|c\|_2^2/(2\gamma) \leq \delta$ with $m_i(x)$, $i \in [1, \dots, f]$ a basis of monomials of degree at most d , where again $\gamma > 0$. We can capture this structure with the help of the set

$$\mathcal{P}_{\text{poly}}(\delta) := \{Q \in \mathcal{P}_\Omega : c \in \mathbb{R}^f, \sum_{x \in X} (\sum_{r \in \mathcal{R}} rQ(r, x) - \sum_{i=1}^f c_i m_i(x))^2/2 + \|c\|_2^2/(2\gamma) \leq \delta\}.$$

This is a closed convex set with non-empty interior under the mild assumption that $\delta \in \text{int}(\{\delta' \geq 0 : \mathcal{P}_{\text{poly}}(\delta') \neq \emptyset\})$ via the same line of argument as in the linear case.

Kernel bandits. The first two classes of structured multi-armed bandit problems can ultimately be subsumed by leaving the rewards between arms to be related as an arbitrary sum of f basis functions. That is, there exists some unobserved vector $c \in \mathbb{R}^f$ of dimension f so that $\sum_{x \in X} (\mathbb{E}[R_t(x)] - \sum_{i=1}^f c_i \varphi_i(x))^2/2 + \|c\|^2/(2\gamma) \leq \delta$ with $\varphi_i(x)$, $i \in [1, \dots, f]$ linearly

Transformation φ	Kernel $K(x, x') =$
Linear functions	$\langle x, x' \rangle$
Polynomial functions (degree d)	$(\langle x, x' \rangle + 1)^d$
Normal basis functions (smoothness σ)	$\exp(-\ x - x'\ _2^2 / \sigma^2)$
Sigmoidal functions (intercept c)	$\tanh(\langle x, x' \rangle + c)$

Table 1: Function spaces and associated kernel functions. For additional examples and references see Schölkopf and Smola (2002).

independent functions. We can capture this structure with the help of the set

$$\mathcal{P}_{\text{kernel}}(\delta) := \left\{ Q \in \mathcal{P}_\Omega : c \in \mathbb{R}^f, \sum_{x \in X} (\sum_{r \in \mathfrak{R}} rQ(r, x) - \sum_{i=1}^f c_i \varphi_i(x))^2 + \|c\|^2 / (2\gamma) \leq \delta \right\}.$$

Clearly, even just verifying whether a particular multi-armed bandit belongs to this convex class is in general not trivial when f becomes exceedingly large. This seems to suggest that polynomial bandits may quickly prove computationally intractable as their degree parameter d is anything but a small number. However, if the inner product between the basis functions can be computed efficiently (see also Table A) with the help of a kernel matrix, i.e., $K_{ij} = \langle \varphi(x_i), \varphi(x_j) \rangle \forall x_i, x_j \in X$, then it is well known that the dependency on f can be avoided altogether as the subsequent lemma indicates.

Lemma 10 (Kernel Representation (Suykens and Vandewalle 1999)). *Suppose that $\forall x_i, x_j \in X$, we have $K_{ij} = \langle \varphi(x_i), \varphi(x_j) \rangle$. Then, we have that*

$$\mathcal{P}_{\text{kernel}}(\delta) = \left\{ Q \in \mathcal{P}_\Omega : R(x) = \sum_{r \in \mathfrak{R}} rQ(r, x) \quad \forall x \in X, 2\delta \geq R^\top (I + \gamma K)^{-1} R \right\}.$$

The previous lemma also shows that $\mathcal{P}_{\text{kernel}}(\delta)$ is always closed. Furthermore, the condition $\delta \in \text{int}(\{\delta' \geq 0 : \mathcal{P}_{\text{kernel}}(\delta') \neq \emptyset\})$ again suffices for our set of reward distributions $\mathcal{P}_{\text{kernel}}(\delta)$ to have non-empty interior.

Monotone bandits. As discussed in the context of revenue management in the Section 3, monotonicity is observed in many practical problems. Here, we assume that the set of arms X comes equipped with some partial order “ \geq ” and the expected rewards of the arms satisfy the inequality $\mathbb{E}[R_t(x_2)] \geq \mathbb{E}[R_t(x_1)] - \delta$ for some $\delta \geq 0$ whenever $x_2 \geq x_1$. (Note that we may do not have access to a complete order over any pair of arms. Otherwise the problem would be trivial if $\delta = 0$.) We can capture this structure with the help of the set

$$\mathcal{P}_{\text{mon}}(\delta) := \left\{ Q \in \mathcal{P}_\Omega : \sum_{r \in \mathfrak{R}} rQ(r, x_2) \geq \sum_{r \in \mathfrak{R}} rQ(r, x_1) - \delta \quad \forall x_1, x_2 \in X \text{ with } x_2 \geq x_1 \right\},$$

which is clearly convex and compact. Again, the set $\mathcal{P}_{\text{mon}}(\delta)$ has non-empty interior under the mild assumption that $\delta \in \text{int}(\{\delta' \geq 0 : \mathcal{P}_{\text{mon}}(\delta') \neq \emptyset\})$.

Convex bandits. Adversarial bandits in which the rewards enjoy a convex structure have received considerable attention in the literature; see (Hazan 2016) and references therein. A stochastic version with finitely many arms $X \subseteq \mathbb{R}^n$ can be modeled simply by considering the set

$$\mathcal{P}_{\text{conv}}(\delta) := \left\{ Q \in \mathcal{P}_{\Omega} : \sum_{r \in \mathfrak{X}} r Q(r, x_2) \geq \sum_{r \in \mathfrak{X}} r Q(r, x_1) + g(x)^\top (x_2 - x_1) - \delta, g(x_1) \in \mathbb{R}^n \forall x_1, x_2 \in X \right\}.$$

The previous set captures a convex structure between the arms by forcing (modulo a constant $\delta \geq 0$) a subgradient inequality $\mathbb{E}[R_t(y)] \geq \mathbb{E}[R_t(x)] + g(x)^\top (y - x)$ to hold between any two arms x_1 and x_2 with $g(x_1) \in \mathbb{R}^n$ the subgradient at arm x_1 ; see also (Boyd and Vandenberghe 2004, Section 6.5.5). Clearly, the set $\mathcal{P}_{\text{conv}}(\delta)$ is a closed convex polytopic set. Once more, our set of reward distributions $\mathcal{P}_{\text{conv}}(\delta)$ has non-empty interior under the mild assumption that $\delta \in \text{int}(\{\delta' \geq 0 : \mathcal{P}_{\text{conv}}(\delta') \neq \emptyset\})$.

B Topological Properties of the Information Distance

With a slight abuse of notation, we denote the set of all distributions supported on a set of outcomes \mathfrak{X} as

$$\mathcal{P}_{\mathfrak{X}} = \{Q \in \mathbb{R}_+^{|\mathfrak{X}|} : \sum_{r \in \mathfrak{X}} Q(r) = 1\}. \quad (28)$$

In this paper, we equip the probability simplex $\mathcal{P}_{\mathfrak{X}}$ with the standard topology relative to its affine hull. The open sets

$$\{Q : \sum_{r \in \mathfrak{X}} Q(r) = 1, \|Q - P\|_{\infty} := \|Q - P\|_{\infty} = \max_{r \in \mathfrak{X}} |Q(r) - P(r)| < \epsilon\}$$

for all $P \in \mathcal{P}_{\mathfrak{X}}$ and $\epsilon > 0$ are a base for this topology. The set of all interior points of $\mathcal{P}_{\mathfrak{X}}$ in this topology is then identified with $\{Q \in \mathbb{R}_{++}^{|\mathfrak{X}|} : \sum_{r \in \mathfrak{X}} Q(r) = 1\}$. The information distance is a natural statistical notion of distance on the probability simplex $\mathcal{P}_{\mathfrak{X}}$ and satisfies several well-known key properties which we use in this paper.

Lemma 11 (Properties of Information Distance). *Let \mathfrak{X} be a discrete set and $\mathcal{P}_{\mathfrak{X}}$ be the set defined in (28). The information distance enjoys the following key properties:*

1. **Information inequality.** $I(P, P') \geq 0$ for all $P, P' \in \mathcal{P}_{\mathfrak{R}}$, and $I(P, P') = 0$ if and only if $P = P'$.

2. **Convexity.** For all pairs $(P_1, P'_1), (P_2, P'_2) \in \mathcal{P}_{\mathfrak{R}} \times \mathcal{P}_{\mathfrak{R}}$ and $\lambda \in [0, 1]$ we have

$$I((1 - \lambda)P_1 + \lambda P_2, (1 - \lambda)P'_1 + \lambda P'_2) \leq (1 - \lambda)I(P_1, P'_1) + \lambda I(P_2, P'_2).$$

3. **Continuity in P' .** $I(P, P')$ is continuous in P' on $\{P' \in \mathcal{P}_{\mathfrak{R}} : P \ll P'\}$ for any fixed $P \in \mathcal{P}_{\mathfrak{R}}$.

4. **Lower semi-continuity.** I is lower semicontinuous on $\mathcal{P}_{\mathfrak{R}} \times \mathcal{P}_{\mathfrak{R}}$. That is, its sub-level sets defined as $\{(P, P') \in \mathcal{P}_{\mathfrak{R}} \times \mathcal{P}_{\mathfrak{R}} : I(P, P') \leq \alpha\}$ are closed (and bounded) for any $\alpha \geq 0$.

Proof of Lemma 11. The information inequality and convexity properties can be found in most textbooks on information theory. For instance, see Theorem 2.6.3 and Theorem 2.7.2 in Cover and Thomas (2012), respectively. For showing the lower semi-continuity result, let us write $I(P, P')$ as

$$I(P, P') = \sum_{r \in \mathfrak{R}} P(r) \log(P(r)/P'(r))$$

for P and P' in $\mathcal{P}_{\mathfrak{R}}$, where we use the convention $0 \log(0/a) = 0$ and $b \log(b/0) = \infty$ for all $a \geq 0$ and $b > 0$. The lower semi-continuity property follows from the fact that functions of the type $(u, v) \mapsto u \log(u/v)$ are lower semi-continuous and summation preserves lower semi-continuity; see (Bertsekas 2009, Proposition 1.1.5).

For continuity in P' , let $\text{supp}(P) = \{r \in \mathfrak{R} : P(r) > 0\}$ denote the support of distribution P . We can write

$$I(P, P') = \sum_{r \in \text{supp}(P)} P(r) \log(P(r)/P'(r)).$$

Continuity of $I(P, P')$ in P' when $P \ll P'$ follows from the immediate observation that the logarithm function $\log(P(r)/P'(r))$ is indeed finite and continuous at any $P'(r) > 0$ for $r \in \text{supp}(P)$. \square

C Duality Results

To show the duality results in this section, we make use of Section I.1, where supplementary materials for duality is presented.

C.1 Proof of Lemma 1

At a high level, we will show that the maximization problem (14) is the dual problem associated with the primal minimization problem (12) as discussed in Section I. However, the primal minimization problem (12) does not satisfy the constraint qualification conditions required by Proposition 8 and hence establishing strong duality directly is problematic. To overcome this challenge, we will show strong duality for a slightly perturbed distance function. By letting the perturbation vanish asymptotically, we will then obtain the stated strong duality result.

We introduce the following perturbed distance functions

$$\begin{aligned} \text{DIST}_\gamma(\eta, x', P) = \min_{Q \in \mathcal{K}} & \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) + \gamma I(P(x^*(P)), Q(x^*(P))) \\ \text{s.t.} & Q \in \mathcal{K}, \\ & \sum_{x \in X, r \in \mathfrak{R}} Q(r, x) = |X|, \\ & \sum_{r \in \mathfrak{R}} r Q(r, x') \geq \text{REW}^*(P) \end{aligned} \quad (29)$$

where the minimum is achieved, as by Lemma 11, its objective function is lower semi-continuous in Q while its feasible domain is compact. The perturbed distance function can indeed be regarded as a penalty formulation of the conic minimization problem (12) with positive penalty term $I(P(x^*(P)), Q(x^*(P))) \geq 0$. That is, we have $I(P(x^*(P)), Q(x^*(P))) = 0$ for any $P(x^*(P)) = Q(x^*(P))$ and $I(P(x^*(P)), Q(x^*(P))) > 0$ otherwise. It is hence evident by comparing problems (29) and (12) that

$$\begin{aligned} \text{DIST}_\gamma(\eta, x', P) \leq \text{DIST}(\eta, x', P) = \min_Q & \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) \\ \text{s.t.} & Q \in \mathcal{K}, \\ & \sum_{x \in X, r \in \mathfrak{R}} Q(r, x) = |X|, \\ & \sum_{r \in \mathfrak{R}} r Q(r, x') \geq \text{REW}^*(P), \\ & Q(x^*(P)) = P(x^*(P)) \end{aligned}$$

for any $\gamma \geq 0$ and deceitful arm $x' \in \tilde{X}_d(P)$. From Nesterov (2004, Theorem 1.3.1) it also follows that $\text{DIST}(\eta, x', P) = \lim_{\gamma \rightarrow \infty} \text{DIST}_\gamma(\eta, x', P)$ as the sublevel sets of the objective function in the perturbed minimization problem (29) are bounded for any $\gamma > 0$.

We will identify a Lagrangian dual formulation for the perturbed minimization problem (29)

with the help of the following perturbed dual function

$$\begin{aligned} \text{DUAL}_\gamma(\eta, x', P; \mu) &:= \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) + \alpha \text{REW}^*(P) + |X| \beta + \\ &\quad \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \eta(x) \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')}{\eta(x)} \right) P(r, x) + \\ &\quad \sum_{x \in x^*(P), r \in \mathfrak{R}} \gamma \log \left(\frac{\gamma - \lambda(r, x) - \beta}{\gamma} \right) P(r, x) + \sum_{x \in x^*(P), r \in \mathfrak{R}} \chi_{-\infty}(\gamma \geq \lambda(r, x) + \beta). \end{aligned}$$

Indeed, using this perturbed dual function we will show the following strong duality result.

Lemma 12 (Dual Formulation of the Perturbed Distance Function). *For any reward distribution $P \in \mathcal{P}$ and suboptimal deceitful arm $x' \in \tilde{X}_d(P)$, the following strong duality equality holds.*

$$\begin{aligned} \text{DIST}_\gamma(\eta, x', P) &= \sup_{\mu} \text{DUAL}_\gamma(\eta, x', P; \mu) \\ \text{s.t. } \mu &\in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*. \end{aligned} \tag{30}$$

We indicate first that the strong duality result stated in Lemma 12 between the perturbed distance function and perturbed dual function implies the claimed result as an asymptotic case by letting $\gamma \rightarrow \infty$. Indeed, as the functions DIST_γ and DUAL_γ are both non-decreasing in γ , and our strong duality result, we can write

$$\begin{aligned} \text{DIST}(\eta, x', P) &= \lim_{\gamma \rightarrow \infty} \text{DIST}_\gamma(\eta, x', P) = \sup_{\gamma \geq 0} \text{DIST}_\gamma(\eta, x', P) \\ &= \sup_{\gamma \geq 0} \sup_{\mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*} \text{DUAL}_\gamma(\eta, x', P; \mu) \\ &= \sup_{\mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*, \gamma \geq 0} \text{DUAL}_\gamma(\eta, x', P; \mu) \\ &= \sup_{\mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*} \lim_{\gamma \rightarrow \infty} \text{DUAL}_\gamma(\eta, x', P; \mu) \end{aligned}$$

where the final limit is identified with

$$\begin{aligned} \text{DUAL}(\eta, x', P; \mu) &= \lim_{\gamma \rightarrow \infty} \text{DUAL}_\gamma(\eta, x', P; \mu) \\ &= \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \eta(x) \log \left(\frac{(\eta(x) - \lambda(x, r) - \beta - \alpha \cdot r \mathbb{1}(x = x'))}{\eta(x)} \right) P(r, x) \\ &\quad - \sum_{x \in x^*(P), r \in \mathfrak{R}} \lambda(r, x) P(r, x) + \alpha \text{REW}^*(P) + \beta |\tilde{X}(P)| \\ &\quad + \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \end{aligned}$$

from which the claimed result follows. □

C.2 Proof of Lemma 12

We will prove Lemma 12 as a special case of the strong duality result stated in Proposition 8 stated in Section I. We first derive the associated dual function for the perturbed problem (29). Afterwards, we verify that the perturbed problem (29) satisfies the required constraint qualification conditions necessary to invoke Proposition 8.

Perturbed dual function. The Lagrangian function associated with the perturbed minimization problem (29) which we introduce in Section I is

$$\begin{aligned}
L_\gamma(\eta, x', P, Q ; \mu) = & \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \left[\eta(x)P(r, x) \log \left(\frac{P(r, x)}{Q(r, x)} \right) - \eta(x)P(r, x) + \eta(x)Q(r, x) \right] \\
& + \sum_{x \in x^*(P), r \in \mathfrak{R}} \left[\gamma P(r, x) \log \left(\frac{P(r, x)}{Q(r, x)} \right) - \gamma P(r, x) + \gamma Q(r, x) \right] \\
& - \sum_{x \in X, r \in \mathfrak{R}} Q(r, x) \lambda(r, x) + \beta \left(|X| - \sum_{x \in X, r \in \mathfrak{R}} Q(r, x) \right) \\
& + \alpha \left(\text{REW}^*(P) - \sum_{r \in \mathfrak{R}} r Q(r, x') \right)
\end{aligned}$$

where $\mu = (\alpha, \beta, \lambda) \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*$. The Lagrangian function can be simplified significantly by grouping terms to

$$\begin{aligned}
L_\gamma(\eta, x', P, Q ; \mu) = & \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \left[Q(r, x) (\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')) + \eta(x)P(r, x) \log \left(\frac{P(r, x)}{Q(r, x)} \right) \right] \\
& + \sum_{x \in x^*(P), r \in \mathfrak{R}} \left[Q(r, x) (\gamma - \lambda(r, x) - \beta) + \gamma P(r, x) \log \left(\frac{P(r, x)}{Q(r, x)} \right) \right] \\
& + \alpha \text{REW}^*(P) + |X| \beta - \sum_{x \in \tilde{X}(P)} \eta(x) - |x^*(P)| \gamma
\end{aligned}$$

where we used the fact that we only are interested in deceitful arms which by definition can not be optimal, i.e., $x' \in \tilde{X}_d(P) \implies x' \notin x^*(P)$. The dual function associated with the perturbed problem (29) is $\text{DUAL}_\gamma(\eta, x', P ; \mu) := \min_{Q \geq 0} L_\gamma(\eta, x', P, Q ; \mu)$ and satisfies the weak duality inequality $\text{DUAL}_\gamma(\eta, x', P ; \mu) \leq \text{DIST}_\gamma(\eta, x', P)$. The dual function can be

expressed equivalently as

$$\begin{aligned}
& \text{DUAL}_\gamma(\eta, x', P; \mu) \\
&= \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \inf_{q \geq 0} \left[q(\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')) + \eta(x) P(r, x) \log \left(\frac{P(r, x)}{q} \right) \right] \\
&\quad + \sum_{x \in x^*(P), r \in \mathfrak{R}} \inf_{q \geq 0} \left[q(\gamma - \lambda(r, x) - \beta) + \gamma P(r, x) \log \left(\frac{P(r, x)}{q} \right) \right] \\
&\quad + \alpha \text{REW}^*(P) + |X| \beta - \sum_{x \in \tilde{X}(P)} \eta(x) - |x^*(P)| \gamma \\
&= \sum_{\{x \in \tilde{X}(P), r: \eta(x) P(r, x) = 0\}} \chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \\
&\quad + \sum_{\{x \in \tilde{x}^*(P), r: \gamma P(r, x) = 0\}} \chi_{-\infty}(\gamma \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \\
&\quad + \sum_{\{x \in \tilde{X}(P), r: \eta(x) P(r, x) > 0\}} P(r, x) \inf_{\theta \geq 0} \left[\theta(\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')) - \eta(x) \log(\theta) \right] \\
&\quad + \sum_{\{x \in x^*(P), r: \gamma P(r, x) > 0\}} P(r, x) \inf_{\theta \geq 0} \left[\theta(\gamma - \lambda(r, x) - \beta) - \gamma \log(\theta) \right] \\
&\quad + \alpha \text{REW}^*(P) + |X| \beta - \sum_{x \in \tilde{X}(P)} \eta(x) - |x^*(P)| \gamma,
\end{aligned}$$

where we obtain the last equality by a change of variables q to $P(r, x)\theta$. Recall that $\chi_{-\infty}(A)$ takes on the value 0 when event A happens and $-\infty$ otherwise. Finally, using the convex conjugacy relationship $\max_{\theta \geq 0} u\theta + \log(\theta) = -(1 + \log(-u)) + \chi_{+\infty}(u \leq 0)$, we can simplify the

perturbed dual function to

$$\begin{aligned}
& \text{DUAL}_\gamma(\eta, x', P ; \mu) \\
&= \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \\
&\quad + \sum_{x \in x^*(P), r \in \mathfrak{R}} \chi_{-\infty}(\gamma \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \\
&\quad + \sum_{\{x \in \tilde{X}(P), r: \eta(x) P(r, x) > 0\}} P(r, x) \eta(x) \left[1 + \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')}{\eta(x)} \right) \right] \\
&\quad + \sum_{\{x \in x^*(P), r: \gamma P(r, x) > 0\}} P(r, x) \gamma \left[1 + \log \left(\frac{\gamma - \lambda(r, x) - \beta}{\gamma} \right) \right] \\
&\quad + \alpha \text{REW}^*(P) + |X| \beta - \sum_{x \in \tilde{X}(P)} \eta(x) - |x^*(P)| \gamma \\
&= \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \\
&\quad + \sum_{x \in x^*(P), r \in \mathfrak{R}} \chi_{-\infty}(\gamma \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \\
&\quad + \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \eta(x) \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')}{\eta(x)} \right) P(r, x) \\
&\quad + \sum_{x \in x^*(P), r \in \mathfrak{R}} \gamma \log \left(\frac{\gamma - \lambda(r, x) - \beta}{\gamma} \right) P(r, x) \\
&\quad + \alpha \text{REW}^*(P) + |X| \beta.
\end{aligned}$$

Strong duality perturbed problem. We can obtain the strong duality result claimed in Lemma 12, i.e., the equality

$$\sup_{\mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*} \text{DUAL}_\gamma(\eta, P, x' ; \mu) = \text{DIST}_\gamma(\eta, x', P)$$

by verifying that the Slater's constraint qualification conditions required by Proposition 8 are satisfied. In the following, we find a Slater point for problem (29). Precisely, we identify a distribution \bar{Q} in the interior of \mathcal{P} which is strictly feasible, i.e., $\sum_{r \in \mathfrak{R}} r \bar{Q}(r, x') > \text{REW}^*(P)$. To do so, we consider any distribution U in the interior of \mathcal{P} ; that is, $U \in \text{int}(\mathcal{P})$. Finding such an interior point is possible as we take that $\text{int}(\mathcal{P}) \neq \emptyset$ as a standing assumption.

Let the distribution Q_{\max} be a distribution which satisfies $\sum_{r \in \mathfrak{R}} r Q_{\max}(r, x') = \text{REW}_{\max}(x', P) > \text{REW}^*(P)$ were the previous strict inequality is guaranteed by the fact that x' is a deceitful arm. Define $Q_\theta := Q_{\max} \cdot \theta + U \cdot (1 - \theta) \in \mathcal{P}$ for all $\theta \in [0, 1]$. We will show that for some values

of θ , the reward distribution Q_θ can serve as a Slater point. Let

$$\theta' = \frac{\text{REW}^*(P) - \sum_{r \in \mathfrak{R}} rU(r, x')}{\text{REW}_{\max}(x', P) - \sum_{r \in \mathfrak{R}} rU(r, x')} \in [-\infty, 1).$$

Observe that for any $\theta \in (\theta', 1] \cap [0, 1)$, we have that the distribution $Q_\theta \in \text{int}(\mathcal{P})$ satisfies $\sum_{r \in \mathfrak{R}} rQ_\theta(x', r) > \text{REW}^*(P)$. Hence, the particular distribution $\bar{Q} = Q_{\bar{\theta}}$ with $\bar{\theta} = (\max\{\theta', 0\} + 1)/2$ serves as a Slater point for problem (29). Note that as $\bar{Q} \in \text{int}(\mathcal{P})$ is also in the interior of \mathcal{P}_Ω and hence in the interior of the domain of the objective function of problem (29). \square

C.3 Proof of Lemma 2

The information minimization problem (23) can be written explicitly as

$$\begin{aligned} \min \quad & \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) \\ \text{s.t.} \quad & \sum_{x \in \tilde{X}, r \in \mathfrak{R}} Q(r, x) [\lambda(r, x) + \beta + \alpha r \mathbf{1}(x = x')] \geq \beta |X| + \alpha \text{REW}^*(P) \quad [\text{dual variable: } \rho] \quad (31) \\ & Q(x^*(P)) = P(x^*(P)) \quad [\text{not dualized}] \end{aligned}$$

where we use the explicit characterization of $\mathcal{H}(x', P; \mu)$ given in Equation (22). We will claim first that the minimum in problem (31) is indeed attained. The remainder of the claimed results will be proven as a special strong duality result of the type discussed in Section I.

Attainment of minimum. We first show that the minimum in (31) is in fact attained. The objective function of the information minimization problem has non-empty level sets

$$\{Q \in \mathcal{H}(x', P; \mu) : \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) \leq \gamma\}$$

for any level $\gamma \geq \text{DIST}(\eta, x', P) = \min_{Q \in \text{DECEIT}(x', P)} \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x))$ as indeed we have the set inclusion $\text{DECEIT}(x', P) \subseteq \mathcal{H}(x', P; \mu)$. These level sets are furthermore closed as the set $\mathcal{H}(x', P; \mu)$ is closed while the objective function is lower semi-continuous as per Lemma 11. Recall that we assume that the logarithmic rate $\eta > 0$ is positive. Hence, the level sets are also bounded as we have

$$\begin{aligned} & \{Q \in \mathcal{H}(x', P; \mu) : \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) \leq \gamma\} \\ & \subseteq \{Q : \eta(x) I(P(x), Q(x)) \leq \gamma \quad \forall x \in \tilde{X}(P), Q(x^*(P)) = P(x^*(P))\}. \end{aligned}$$

Bertsekas (2009, Prop. 3.2.1) guarantees that the set of minima in (31) is non-empty and compact.

Dual function. The associated Lagrangian to the information minimization problem (31) introduced in Section I is equal to

$$L(\eta, x', P, Q ; \rho) := \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) \\ + \rho \left(\sum_{x \in X, r \in \mathfrak{R}} Q(r, x) [-\lambda(r, x) - \beta - \alpha r \mathbb{1}(x = x')] \right) + \beta |X| + \alpha \text{REW}^*(P)$$

which can be simplified further by collecting terms to

$$L(\eta, x', P, Q ; \rho) = \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} Q(r, x) (-\rho \lambda(r, x) - \rho \beta - \rho \alpha r \mathbb{1}(x = x')) \\ + \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q(x)) \\ + \sum_{x \in x^*(P), r \in \mathfrak{R}} Q(r, x) (-\rho \lambda(r, x) - \rho \beta) + \rho \beta |X| + \rho \alpha \text{REW}^*(P).$$

The dual function associated with the perturbed problem (29) is as discussed in Section I identified with $g(\eta, x', P ; \rho) := \min_{Q(x^*(P))=P(x^*(P))} L(\eta, x', P, Q ; \rho)$. The dual function can be expressed equivalently as

$$g(\eta, x', P ; \mu) \\ = \min_{Q(x^*(P))=P(x^*(P))} \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} Q(r, x) (-\rho \lambda(r, x) - \rho \beta - \rho \alpha r \mathbb{1}(x = x')) \\ + \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \eta(x) \left[P(r, x) \log \left(\frac{P(r, x)}{Q(r, x)} \right) - P(r, x) + Q(r, x) \right] \\ + \sum_{x \in x^*(P), r \in \mathfrak{R}} Q(r, x) (-\rho \lambda(r, x) - \rho \beta) + \rho \beta |X| + \rho \alpha \text{REW}^*(P) \\ = \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \inf_{q \geq 0} \left[q (\eta(x) - \rho \lambda(r, x) - \rho \beta - \rho \alpha \cdot r \mathbb{1}(x = x')) + \eta(x) P(r, x) \log \left(\frac{P(r, x)}{q} \right) \right] \\ - \sum_{x \in x^*(P)} \rho \lambda(r, x) P(r, x) + \rho \alpha \text{REW}^*(P) + |\tilde{X}(P)| \beta \rho - \sum_{x \in \tilde{X}(P)} \eta(x) \\ = \sum_{\{x \in \tilde{X}(P), r: \eta(x) P(r, x) = 0\}} \chi_{-\infty}(\eta(x) \geq \rho \lambda(r, x) + \rho \beta + \rho \alpha \cdot r \mathbb{1}(x = x')) \\ + \sum_{\{x \in \tilde{X}(P), r: \eta(x) P(r, x) > 0\}} P(r, x) \inf_{\theta \geq 0} \left[\theta (\eta(x) - \rho \lambda(r, x) - \rho \beta - \rho \alpha \cdot r \mathbb{1}(x = x')) - \eta(x) \log(\theta) \right] \\ - \sum_{x \in x^*(P)} \rho \lambda(r, x) P(r, x) + \rho \alpha \text{REW}^*(P) + |\tilde{X}(P)| \beta \rho - \sum_{x \in \tilde{X}(P)} \eta(x)$$

where we obtain the last equality by a change of variables q to $P(r, x)\theta$. Recall that $\chi_{-\infty}(A)$ takes on the value 0 when event A happens and $-\infty$ otherwise. Finally, using the convex

conjugacy relationship $\max_{\theta \geq 0} u\theta + \log(\theta) = -(1 + \log(-u)) + \chi_{+\infty}(u \leq 0)$, we can simplify the perturbed dual function to

$$\begin{aligned}
& g(\eta, x', P; \rho) \\
&= \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \chi_{-\infty}(\eta(x) \geq \rho\lambda(r, x) + \rho\beta + \rho\alpha \cdot r\mathbb{1}(x = x')) \\
&\quad + \sum_{\{x \in \tilde{X}(P), r: \eta(x)P(r, x) > 0\}} P(r, x)\eta(x) \left[1 + \log\left(\frac{\eta(x) - \rho\lambda(r, x) - \rho\beta - \rho\alpha \cdot r\mathbb{1}(x = x')}{\eta(x)}\right) \right] \\
&\quad - \sum_{x \in x^*(P)} \rho\lambda(r, x)P(r, x) + \rho\alpha \text{REW}^*(P) + \rho|\tilde{X}(P)|\beta - \sum_{x \in \tilde{X}(P)} \eta(x) \\
&= \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \chi_{-\infty}(\eta(x) \geq \rho\lambda(r, x) + \rho\beta + \rho\alpha \cdot r\mathbb{1}(x = x')) \\
&\quad + \sum_{x \in \tilde{X}(P), r \in \mathfrak{R}} \eta(x) \log\left(\frac{\eta(x) - \rho\lambda(r, x) - \rho\beta - \rho\alpha \cdot r\mathbb{1}(x = x')}{\eta(x)}\right) P(r, x) \\
&\quad - \sum_{x \in x^*(P)} \rho\lambda(r, x)P(r, x) + \rho\alpha \text{REW}^*(P) + |\tilde{X}(P)|\rho\beta. \\
&= \text{DUAL}(\eta, x', P; \rho \cdot \mu).
\end{aligned}$$

Strong duality. We can now obtain the strong duality equality

$$(31) = \max_{\rho \geq 0} \text{DUAL}(\eta, x', P; \rho \cdot \mu)$$

via Proposition 8. Recall that the set $\text{DECEIT}(x', P) \subseteq \mathcal{H}(x', P, \mu)$ and hence in order to apply Proposition 8, it suffices to construct a distribution $\bar{Q} \in \text{int}(\mathcal{P}_\Omega)$ which satisfies

$$\sum_{r \in \mathfrak{R}} \bar{Q}(r, x') > \text{REW}^*(P) \quad \text{and} \quad \bar{Q}(x^*(P)) = P(x^*(P)).$$

Let $Q_\theta = \theta Q^* + (1 - \theta)P \in \mathcal{P}$ with Q^* a minimizer in problem (12). We will show there for some values of θ , the reward distribution Q_θ can serve as a Slater point. Let

$$\theta' = \frac{\text{REW}^*(P) - \sum_{r \in \mathfrak{R}} rP(r, x')}{\text{REW}_{\max}(x', P) - \sum_{r \in \mathfrak{R}} rP(r, x')} \in [-\infty, 1).$$

Observe that as $P \in \text{int}(\mathcal{P})$ for any $\theta \in (\theta', 1] \cap [0, 1)$, we have that the distribution $Q_\theta \in \text{int}(\mathcal{P}) \subseteq \text{int}(\mathcal{P}_\Omega)$ satisfies $\sum_{r \in \mathfrak{R}} rQ_\theta(x', r) > \text{REW}^*(P)$. Hence, the particular distribution $\bar{Q} = Q_{\bar{\theta}}$ with $\bar{\theta} = (\max\{\theta', 0\} + 1)/2$ serves as a Slater point for problem (31).

C.4 Dual Cones for Running Examples

Example 4 (Dual Cones). *In what follows, we discuss the cones \mathcal{K} and its dual \mathcal{K}^* in the context of our running examples. Detailed derivation of dual cones are deferred to Section C.4.1.*

1–2. *Generic and separable bandits. Recall that the feasible set \mathcal{P} of separable bandits is $\mathcal{P} = \prod_{x \in X} \mathcal{P}_x$, where \mathcal{P}_x , $x \in X$, is a closed feasible set for the reward distribution of arm x . The cone \mathcal{K} associated with feasible set \mathcal{P} can be written as*

$$\mathcal{K} = \left\{ Q \in \mathbb{R}_+^{|\mathfrak{R}| \times |X|} : \begin{array}{ll} \sum_{r \in \mathfrak{R}} Q(r, x) = \theta & \forall x \in X, \\ Q(x) \in \text{cone}(\mathcal{P}_x) & \forall x \in X. \end{array} \right\}, \quad (32)$$

The dual cone of the separable bandit cone \mathcal{K} in Equation (32) is given as

$$\mathcal{K}^* = \left\{ \lambda \in \mathbb{R}^{|\mathfrak{R}| \times |X|} : \begin{array}{ll} \exists \gamma(x) \in \mathbb{R} & \forall x \in X, \\ \lambda(x) - \mathbf{1} \cdot \gamma(x) \in \text{cone}(\mathcal{P}_x)^* & \forall x \in X, \\ \sum_{x \in X} \gamma(x) = 0 \end{array} \right\},$$

where $\mathbf{1}$ is the all ones vector with a size of $|\mathfrak{R}|$. The dual cone is here characterized by $|X|$ conic constraints with the help an auxiliary variables $\gamma \in \mathbb{R}^{|X|}$. This implies that when $\text{cone}(\mathcal{P}_x)^$ is efficiently representable for each arm, then \mathcal{K}^* admits an efficient barrier function. If for instance the structural information \mathcal{P}_x are polyhedral, than $\text{cone}(\mathcal{P}_x)$ is a polyhedral cone and so will be its dual cone $\text{cone}(\mathcal{P}_x)^*$ and hence also the cone \mathcal{K}^* , implying it is efficiently representable.*

3. *Lipschitz Bernoulli bandits. Recall that the feasible set of Lipschitz Bernoulli Bandits is given by $\mathcal{P}_{\text{Lips}} = \left\{ P \in \mathcal{P}_\Omega : P(1, x) - P(1, x') \leq L \cdot D(x, x') \quad \forall x, x' \in X \right\}$. Then, it is easy to observe that the cone \mathcal{K} associated with $\mathcal{P}_{\text{Lips}}$ is given by*

$$\mathcal{K} = \left\{ Q \in \mathbb{R}_+^{2 \times |X|} : \begin{array}{ll} \exists \theta \in \mathbb{R}, & \\ Q(0, x) + Q(1, x) = \theta & \forall x \in X, \\ Q(1, x) - Q(1, x') \leq \theta \cdot L \cdot D(x, x') & \forall x, x' \in X \end{array} \right\}. \quad (33)$$

Note that the last constraint enforces the Lipschitz property of the reward distributions. The

dual cone of the Lipschitz bandits cone \mathcal{K} in Equation (33) is given as

$$\mathcal{K}^* = \left\{ \lambda \in \mathbb{R}^{2 \times |X|} : \begin{array}{ll} \exists \Lambda(x, x') \in \mathbb{R}_+, \exists \gamma(x) \in \mathbb{R} & \forall x, x' \in X, \\ \lambda(1, x) + \sum_{x' \in X} (\Lambda(x, x') - \Lambda(x', x)) \geq \gamma(x) & \forall x \in X, \\ \lambda(0, x) \geq \gamma(x) & \forall x \in X, \\ \sum_{x \in X} \gamma(x) = L \sum_{x, x' \in X} D(x, x') \Lambda(x, x') \end{array} \right\}.$$

The cone \mathcal{K}^* is here a polyhedral set characterized by $2|X|$ conic constraints with the help of two auxiliary variables $\gamma \in \mathbb{R}^{|X|}$ and $\Lambda(x, x') \in \mathbb{R}_+^{|X| \times |X|}$, implying \mathcal{K}^* admits an efficient barrier function. \square

C.4.1 Detailed Derivation of Dual Cones

Dual cone of generic and separable bandits. From the definition of dual cones, it follows that $\lambda \in \mathcal{K}^* \iff 0 \leq \min_{Q \in \mathcal{K}} \langle \lambda, Q \rangle$, where for separable bandits, the cone \mathcal{K} is given in Equation (32). Here, we have that the previous primal minimization problem $\min_{Q \in \mathcal{K}} \langle \lambda, Q \rangle$ can be characterized as

$$\begin{aligned} \min_{Q, \theta} \quad & \langle \lambda, Q \rangle \\ \text{s.t.} \quad & \sum_{r \in \mathfrak{R}} Q(r, x) = \theta \quad \forall x \in X, & [\text{dual variable } \gamma \in \mathbb{R}^{|X|}] \\ & Q(x) \in \text{cone}(\mathcal{P}_x) \quad \forall x \in X, & [\text{dual variable } \Lambda(x) \in \text{cone}(\mathcal{P}_x)^*] \end{aligned} \quad (34)$$

where the constraint of the above optimization problem enforces Q to belong to cone \mathcal{K} ; see Equation (32). Please remark that this is not the only possible valid optimization characterization. Nevertheless, its primal characterization will be found to have a particularly simple dual. We remark that the special case of generic bandit problems is subsumed in the separable bandit problem and corresponds to the particular choice cone (\mathcal{P}_x) the positive orthant. The associated Lagrangian function as introduced in Section I is

$$L_\lambda(Q, \theta ; \gamma, \Lambda) = \langle \lambda, Q \rangle + \sum_{x \in X} \gamma(x) (\theta - \sum_{x \in X, r \in \mathfrak{R}} Q(r, x)) - \sum_{x \in X} \langle \Lambda(x), Q(x) \rangle,$$

which can be simplified by collecting terms to

$$L_\lambda(Q, \theta ; \gamma, \Lambda) = \sum_{x \in X, r \in \mathfrak{R}} (\lambda(r, x) - \gamma(x) - \Lambda(r, x)) \cdot Q(r, x) + \theta \sum_{x \in X} \gamma(x).$$

The associated dual function as introduced in Section I is hence

$$g_\lambda(\gamma, \Lambda) := \min_{Q, \theta} L_\lambda(Q, \theta ; \gamma, \Lambda) = \sum_{x \in X, r \in \mathfrak{R}} \chi_{-\infty}(\lambda(r, x) - \gamma(x) - \Lambda(r, x) = 0) + \chi_{-\infty}(\sum_{x \in X} \gamma(x) = 0).$$

The sets $\text{cone}(\mathcal{P}_x)$ have non-empty interior, as $\text{int}(\prod_{x \in X} \mathcal{P}_x) = \prod_{x \in X} \text{int}(\mathcal{P}_x)$ is non-empty. Thus, the primal optimization problem (34) satisfies Slater's constraint qualification and hence strong duality holds, per Proposition 8. Hence, we have the equivalence

$$0 \leq \min_{Q \in \mathcal{K}} \langle \lambda, Q \rangle = \max_{\gamma, \Lambda} g_\lambda(\gamma, \Lambda) \iff \left\{ \begin{array}{l} \exists \gamma(x) \in \mathbb{R} \quad \forall x \in X, \\ \lambda(x) - \mathbf{1} \cdot \gamma(x) \in \text{cone}(\mathcal{P}_x)^* \quad \forall x \in X, \\ \sum_{x \in X} \gamma(x) = 0. \end{array} \right\} \iff \lambda \in \mathcal{K}^*,$$

where $\mathbf{1}$ is the all ones vector of size of $|\mathfrak{X}|$.

Dual cone of Lipschitz bandits. Using the definition of dual cone, we again have $\lambda \in \mathcal{K}^* \iff 0 \leq \min_{Q \in \mathcal{K}} \langle \lambda, Q \rangle$, where for Lipschitz bandits, the cone \mathcal{K} is given in Equation (33). Here, we have that the previous primal minimization problem $\min_{Q \in \mathcal{K}} \langle \lambda, Q \rangle$ can be characterized as

$$\begin{array}{ll} \min_{Q, \theta} & \langle \lambda, Q \rangle \\ \text{s.t.} & Q(0, x) + Q(1, x) = \theta, \quad \forall x \in X \quad [\text{dual variable } \gamma \in \mathbb{R}] \\ & Q(1, x) - Q(1, x') \leq \theta \cdot L \cdot D(x, x') \quad \forall x, x' \in X, \quad [\text{dual variable } \Lambda \in \mathbb{R}_+^{|\mathfrak{X}| \times |\mathfrak{X}|}] \quad (35) \\ & Q \geq 0, \quad [\text{not dualized}] \\ & \theta \in \mathbb{R} \quad [\text{not dualized}] \end{array}$$

where the constraint of the above optimization problem enforces Q to belong to cone \mathcal{K} ; see Equation (33). Please remark that this is not the only possible valid characterization. Nevertheless, its primal characterization will be found to have a particularly simple dual. The associated Lagrangian function as introduced in Section I is

$$\begin{aligned} L_\lambda(Q, \theta; \gamma, \Lambda) &= \langle \lambda, Q \rangle + \sum_{x \in X} \gamma(x) (\theta - Q(0, x) - Q(1, x)) \\ &\quad + \sum_{x, x' \in X} \Lambda(x, x') (Q(1, x) - Q(1, x') - \theta \cdot L \cdot D(x, x')), \end{aligned}$$

which can be simplified by collecting terms to

$$\begin{aligned} &L_\lambda(Q, \theta; \gamma, \Lambda) \\ &= \sum_{x \in X} (\lambda(0, x) - \gamma(x)) \cdot Q(0, x) + \theta (\sum_{x \in X} \gamma(x) - \sum_{x, x' \in X} L \cdot D(x, x') \cdot \Lambda(x, x')) \\ &\quad + \sum_{x \in X} \left(\lambda(1, x) - \gamma(x) + \sum_{x' \in X} \Lambda(x, x') - \sum_{x' \in X} \Lambda(x', x) \right) \cdot Q(1, x). \end{aligned}$$

The associated dual function as introduced in Section I is hence

$$\begin{aligned}
g_\lambda(\gamma, \Lambda) &:= \min_{Q \geq 0, \theta} L_\lambda(Q, \theta; \gamma, \Lambda) \\
&= \sum_{x \in X} \chi_{-\infty}(\lambda(0, x) - \gamma(x) \geq 0) \\
&\quad + \sum_{x \in X} \chi_{-\infty}(\lambda(1, x) - \gamma(x) + \sum_{x' \in X} (\Lambda(x, x') - \Lambda(x', x)) \geq 0) \\
&\quad + \chi_{-\infty}(\sum_{x \in X} \gamma(x) \geq \sum_{x, x' \in X} L \cdot D(x, x') \cdot \Lambda(x, x')).
\end{aligned}$$

As the feasible set in (35) is a non-empty bounded polyhedron, strong linear duality holds. Hence, we have the equivalence

$$0 \leq \min_{Q \in \mathcal{K}} \langle \lambda, Q \rangle = \max_{\gamma, \Lambda} g_\lambda(\gamma, \Lambda) \iff \left\{ \begin{array}{ll} \exists \Lambda(x, x') \in \mathbb{R}_+, \gamma(x) \in \mathbb{R} & \forall x, x' \in X, \\ \lambda(0, x) \geq \gamma(x) & \forall x \in X, \\ \lambda(1, x) + \sum_{x' \in X} \Lambda(x, x') - \sum_{x' \in X} \Lambda(x', x) \geq \gamma(x) & \forall x \in X, \\ \sum_{x \in X} \gamma(x) = \sum_{x, x' \in X} L \cdot D(x, x') \cdot \Lambda(x, x') \end{array} \right\}.$$

D Proof of Concentration Bounds in Section 7.1

D.1 Proof of Lemma 3

The proof is inspired by the proof of Lemma 4.3 in Combes and Proutiere (2014). In this proof, to simplify notation, we denote the stopping time $t(s, x)$ by $t(s)$.

Let \mathcal{F}_t^π be the σ -algebra defined in Section 3 generated by the rewards observed and arms pulled before round t . Let us consider a fixed arm x and reward $r \in \mathfrak{R}$. We then define B_τ as $\mathbb{1}(x_\tau = x)$; that is, B_τ is one if arm x is played in round τ , and zero otherwise. Observe that B_τ is \mathcal{F}_τ measurable. Let

$$S_t := \sum_{\tau=1}^{t-1} B_\tau \left(\mathbb{1}(R_\tau(x) = r) - \mathbb{E}[\mathbb{1}(R_\tau(x) = r)] \right).$$

Note that S_t can be written as

$$S_t = \sum_{\tau=1}^{t-1} B_\tau \left(\mathbb{1}(R_\tau(x) = r) - P(r, x) \right) = N_t(x) \left(P_t(r, x) - P(r, x) \right),$$

where the last equation holds because our estimate of $P(r, x)$ in round n , i.e., $P_t(r, x)$, is

$\sum_{\tau=1}^{t-1} B_\tau \mathbb{1}(R_\tau(x) = r) / N_t(x)$, where $N_t(x) := \sum_{\tau=1}^{t-1} B_\tau$ is the number of times in before round t that arm x is played. For any $s > 1$, we would like to show that

$$\text{Prob} \left[|S_{t(s)}| \geq \kappa \cdot N_{t(s)}(x), t(s) \leq T \right] \leq \exp(-s\epsilon\kappa^2/2),$$

where $t(s) \in [s, \dots, T+1]$ is any stopping time such that either $N_{t(s)}(x) \geq \epsilon s/4$ or $t(s) = T+1$, and κ is a positive number. Showing the aforementioned inequality gives us the desired result.

With this goal in mind, define $G_t = \exp(4\kappa(S_t - \kappa N_t(x))) \mathbb{1}(t \leq T)$ for all $t \in [1, \dots, T+1]$. Then, we have

$$\begin{aligned} \text{Prob} \left[S_{t(s)} \geq \kappa N_{t(s)}, t(s) \leq T \right] &= \text{Prob} \left[\left(\exp(4\kappa(S_{t(s)} - \kappa N_{t(s)}(x))) \cdot \mathbb{1}(t(s) \leq T) \right) \geq 1 \right] \\ &= \text{Prob} \left[G_{t(s)} \geq 1 \right] \leq \mathbb{E} \left[G_{t(s)} \right] \end{aligned}$$

where the ultimate inequality is due to Markov as $G_{t(s)} \geq 0$. In the following, to upper bound $\text{Prob} \left[S_{t(s)} \geq \kappa N_{t(s)}, t(s) \leq T \right]$, we upper bound $\mathbb{E} \left[G_{t(s)} \right]$. To do so, we need a few definitions. Let

$$\tilde{G}_t = \exp \left(\sum_{\tau=1}^{t-1} Y_\tau \right) \mathbb{1}(t \leq T) \quad \text{with} \quad Y_\tau = B_\tau \left(4\kappa(\mathbb{1}(R_\tau(x) = r) - P(r, x)) - 2\kappa^2 \right).$$

We note that

$$G_t = \tilde{G}_t \exp(-N_t(x)(4\kappa^2 - 2\kappa^2)) = \tilde{G}_t \exp(-2N_t(x)\kappa^2).$$

We now use the fact that $N_{t(s)}(x) \geq s\epsilon/4$ if $t(s) \leq T$ to upper bound $G_{t(s)}$ by

$$G_{t(s)} = \tilde{G}_{t(s)} \exp(-2N_{t(s)}(x)\kappa^2) \leq \tilde{G}_{t(s)} \exp(-s\epsilon\kappa^2/2).$$

The above inequality holds even when $t(s) = T+1$ as $G_{T+1} = \tilde{G}_{T+1} = 0$. Therefore,

$$\mathbb{E} \left[G_{t(s)} \right] \leq \mathbb{E} \left[\tilde{G}_{t(s)} \right] \exp(-s\epsilon\kappa^2/2)$$

So far, we have upper bounded $\mathbb{E} \left[G_{t(s)} \right]$ as a function of $\mathbb{E} \left[\tilde{G}_{t(s)} \right]$. Next, we will show that $(\tilde{G}_n)_n$ is a super-martingale sequence. This allows us to upper bound $\mathbb{E} \left[\tilde{G}_{t(s)} \right]$. First observe that $\mathbb{E} \left[\tilde{G}_{T+1} | \mathcal{F}_T \right] = 0 \leq \tilde{G}_T$. For any $t \leq T-1$, since B_t is \mathcal{F}_t measurable,

$$\mathbb{E} \left[\tilde{G}_{t+1} | \mathcal{F}_t \right] = \tilde{G}_t (1 - B_t) + \tilde{G}_t B_t \mathbb{E} \left[\exp(Y_t) \right].$$

Recall that when $B_t = 1$, we have $Y_t = 4\kappa(\mathbb{1}(R_t(x) = r) - P(r, x)) - 2\kappa^2$. Then, by invoking Equation 4.16 in Hoeffding (1994) and considering the fact that $\mathbb{1}(R_t(x) = r) \leq 1$, we have

$$\mathbb{E} \left[\exp(4\kappa(\mathbb{1}(R_t(x) = r) - P(r, x))) \right] \leq \exp(2\kappa^2).$$

This implies that $\mathbb{E}[\exp(Y_t)] \leq 1$, and as a result, $(\tilde{G}_t)_t$ is a super-martingale sequence: $\mathbb{E}[\tilde{G}_{t+1}|\mathcal{F}_t] \leq \tilde{G}_t$. Since $t(s) \leq T + 1$, sequence $(\tilde{G}_t)_t$ is super-martingale, we can apply the Doob's optional stopping theorem, to get $\mathbb{E}[\tilde{G}_{t(s)}] \leq \mathbb{E}[\tilde{G}_1] = 1$. Putting everything together, we have

$$\text{Prob}[S_{t(s)} \geq \kappa N_{t(s)}, t(s) \leq T] \leq \mathbb{E}[G_{t(s)}] \leq \mathbb{E}[\tilde{G}_{t(s)}] \exp(-s\epsilon\kappa^2/2) \leq \exp(-s\epsilon\kappa^2/2).$$

Then, by symmetry, we can show that

$$\text{Prob}[|S_{t(s)}| \geq \kappa N_{t(s)}, t(s) \leq T] \leq 2 \exp(-s\epsilon\kappa^2/2),$$

which is the desired result. □

D.2 Proof of Lemma 4

Let us take $\mathfrak{R} = \{1, 2, 3, \dots, k\} = [1, \dots, k]$ without loss of generality. Define $\delta \geq |X|(|\mathfrak{R}|-1)+1$ and $\eta = 1/(\delta-1) > 0$. The core of the argument involves partitioning the event set of interest in geometrically sized slices. Define $D := \lceil \log(t)/\log(1+\eta) \rceil$ as the number of such slices. Let now $\mathcal{D} := \{1, \dots, D\}^{(|\mathfrak{R}|-1)|X|}$ and $\mathcal{Y} := \{-1, 1\}^{(|\mathfrak{R}|-1)|X|}$. Introduce the event

$$A = \{\sum_{x \in X} N_t(x) I(P_t(x), P(x)) \geq \delta\},$$

and partitions

$$B_d = \cap_{x \in X, r \in [k-1]} \{(1+\eta)^{d(r,x)-1} \leq N'_t(r, x) \leq (1+\eta)^{d(r,x)}\},$$

$$C_y = \cap_{x \in X, r \in [k-1]} \{P'_t(r, x)y(r, x) \leq P'(r, x)y(r, x)\}$$

indexed in $d \in \mathcal{D}$ and $y \in \mathcal{Y}$. Here,

$$N'_t(r, x) := \sum_{\tau=1}^{t-1} \mathbb{1}\{x_\tau = x, R_\tau(x) \geq r\} \quad r \in [1, \dots, k-1], x \in X$$

is the number of times before round t in which we play arm x and receive a reward of at least r . Furthermore,

$$P'(r, x) := \frac{P(r, x)}{\sum_{r' \geq r} P(r', x)}, \quad r \in [1, \dots, k-1], x \in X, \quad (36)$$

$$P'_t(r, x) := \frac{P_t(r, x)}{\sum_{r' \geq r} P_t(r', x)}, \quad r \in [1, \dots, k-1], x \in X. \quad (37)$$

We will also slightly abuse notation and extent B_d to all $d \in \mathcal{D}' = \{0, \dots, D\}^{(|\mathfrak{A}|-1)|X|}$ with the understanding that $d(r, x) = 0$ implies $N'_t(r, x) = 0$. The index $y(r, x)$ partitions the events into two cases: (1) $y(r, x) = 1$ under which $P'_t(r, x) \leq P'(r, x)$ and (2) $y(r, x) = -1$ under which $P'_t(r, x) \geq P'(r, x)$.

We have that $A = \cup_{d \in \mathcal{D}', y \in \mathcal{Y}} (A \cap B_d \cap C_y)$ and this allows us to use the union bound:

$$\text{Prob}[A] \leq \sum_{d \in \mathcal{D}'} \sum_{y \in \mathcal{Y}} \text{Prob}[A \cap B_d \cap C_y]. \quad (38)$$

Consider a particular $d \in \mathcal{D}'$ and $y \in \mathcal{Y}$. Applying Lemma 13 with $\bar{N}'(r, x) = (1 + \eta)^{d(r, x) - 1}$ and $\delta \geq (1 + \eta) |X| (|\mathfrak{A}| - 1)$, we obtain the inequality

$$\text{Prob}[A \cap B_d \cap C_y] \leq \left(\frac{\delta e}{|X| (|\mathfrak{A}| - 1)} \right)^{|X| (|\mathfrak{A}| - 1)} \cdot \exp(-\delta / (1 + \eta)).$$

We have $|\mathcal{D}' \times \mathcal{Y}| = (2(D + 1))^{|X| (|\mathfrak{A}| - 1)}$ and thus from Inequality (38), we obtain

$$\text{Prob}[A] \leq \left(\frac{2(D + 1)\delta e}{|X| (|\mathfrak{A}| - 1)} \right)^{|X| (|\mathfrak{A}| - 1)} \exp(-\delta / (1 + \eta)).$$

With the choice $\eta = 1/(\delta - 1)$ and using $\log(1 + \eta) = -\log(1/(1 + \eta)) \geq 1 - 1/(1 + \eta) = 1/\delta$, we obtain

$$\text{Prob}[A] \leq \left(\frac{\delta [\log(T)\delta + 1] 2e}{|X| (|\mathfrak{A}| - 1)} \right)^{|X| (|\mathfrak{A}| - 1)} \exp(-\delta / (1 + \eta)).$$

Furthermore, as $1/(\eta + 1) = (\delta - 1)/\delta$, we get the desired result

$$\text{Prob}[A] \leq \left(\frac{\delta [\log(T)\delta + 1] 2e}{|X| (|\mathfrak{A}| - 1)} \right)^{|X| (|\mathfrak{A}| - 1)} e \cdot \exp(-\delta).$$

□

Lemma 13. For any $r \in [1, \dots, k-1]$ and $x \in X$, let $\bar{N}'(r, x) \in [0, \dots, t]$, where $k = |\mathfrak{A}|$,

$\mathfrak{K} = \{1, 2, \dots, k\}$. For any $\eta > 0$, define the event

$$B := \cap_{x \in X} \cap_{r=1}^{k-1} \{\bar{N}'(r, x) \leq N'_t(r, x) \leq (1 + \eta)\bar{N}'(r, x)\}$$

$$C := \cap_{x \in X} \cap_{r=1}^{k-1} \{P'_t(r, x)y(r, x) \leq P'(r, x)y(r, x)\}$$

for $y(r, x) \in \{+1, -1\}$, $r \in [1, \dots, k-1]$ and $x \in X$. For $\delta \geq (1 + \eta)(|\mathfrak{K}| - 1) \cdot |X|$, we have the inequality

$$\text{Prob} \left[B \cap C \cap \left(\sum_{x \in X} N_t(x) I(P_t(x), P(x)) \geq \delta \right) \right] \leq \left(\frac{\delta e}{(|\mathfrak{K}| - 1) \cdot |X|} \right)^{(|\mathfrak{K}| - 1) \cdot |X|} \exp(-\delta / (1 + \eta)).$$

Proof of Lemma 13. The proof of this lemma takes advantage of the decomposition result presented in Lemma 14. This lemma allows us to decompose $N_t(x)I(P_t(x), P(x))$ as follows

$$N_t(x)I(P_t(x), P(x)) = \sum_{r=1}^{k-1} N'_t(r, x) I_B(P'_t(r, x), P'(r, x)),$$

where $N'_t(r, x) := \sum_{r' \geq r} N_t(r', x)$, $r \in [1, \dots, k-1]$, $x \in X$, and

$$I_B(P'_t(r, x), P'(r, x)) := I((P'_t(r, x), 1 - P'_t(r, x)), (P'(r, x), 1 - P'(r, x)))$$

denotes the information distance between two Bernoulli distributions with mean $P'_t(r, x)$ and $P'(r, x)$. Recall that for any $r \in [1, \dots, k-1]$, $P'(r, x) = P(r, x) / \sum_{r' \geq r} P(r', x)$ and $P'_t(r, x) = P_t(r, x) / \sum_{r' \geq r} P_t(r', x)$; that is, $P'(r, x)$ is the probability of receiving a reward of r by pulling arm x , conditioned on the reward being greater than or equal to r . Having this decomposition in mind, we will show that for any $\xi(r, x) \geq 0$, $r \in [1, \dots, k-1]$ and $x \in X$, the following inequality holds

$$\begin{aligned} & \text{Prob} \left[\cap_{x \in X} \cap_{r=1}^{k-1} \{B \cap C \cap (N'_t(r, x) I_B(P'_t(r, x), P'(r, x)) \geq \xi(r, x))\} \right] \\ & \leq \exp(\sum_{x \in X} \sum_{r=1}^{k-1} \xi(r, x) / (1 + \eta)). \end{aligned} \tag{39}$$

Then, by applying the stochastic dominance bound, stated in Lemma 15, where we use

$Z(r, x) = \mathbb{1}\{B \cap C\} \cdot \mathbb{1}\{N'_t(r, x)I_B(P'_t(r, x), P'(r, x))\}$ and $a = 1/(1 + \eta)$, we get

$$\begin{aligned} \text{Prob} \left[B \cap C \cap \left(\sum_{x \in X} N_t(x) I(P_t(x), P(x)) \geq \delta \right) \right] &\leq \left(\frac{\delta e}{|X| (|\mathfrak{R}| - 1)(1 + \eta)} \right)^{|X|(|\mathfrak{R}| - 1)} \exp(-\delta/(1 + \eta)) \\ &\leq \left(\frac{\delta e}{|X| (|\mathfrak{R}| - 1)} \right)^{|X|(|\mathfrak{R}| - 1)} \exp(-\delta/(1 + \eta)) \end{aligned}$$

establishing the claim.

It remains to show Inequality (39). We have that

$$\begin{aligned} &\text{Prob} \left[\bigcap_{x \in X} \bigcap_{r=1}^{k-1} \{B \cap C \cap (N'_t(r, x)I_B(P'_t(r, x), P'(r, x)) \geq \xi(r, x))\} \right] \\ &\leq \text{Prob} \left[\bigcap_{x \in X} \bigcap_{r=1}^{k-1} \{B \cap C \cap (\bar{N}'(r, x)I_B(P'_t(r, x), P'(r, x)) \geq \xi(r, x)/(1 + \eta))\} \right]. \end{aligned}$$

We define now $Q'(r, x)$ for any $r \in [1, \dots, k-1]$ and $x \in X$ such that if there exists a $q' \in [0, 1]$ such that $q'y(r, x) \leq P'(r, x)y(r, x)$ and $\bar{N}'(r, x) \cdot I_B(q', P'(r, x)) = \xi(r, x)/(1 + \eta)$, then $Q'(r, x) = q'$. Evidently, if no such q' exists, then the event $(\bar{N}'(r, x)I_B(P'_t(r, x), P'(r, x)) \geq \xi(r, x)/(1 + \eta))$ can also not occur and Inequality (39) holds trivially. We now assume that we are not in this trivial case and we have

$$\bar{N}'(r, x)I_B(Q'(r, x), P'(r, x)) = \xi(r, x)/(1 + \eta) \quad \forall r \in [1, \dots, k-1], x \in X. \quad (40)$$

With this definition, we have

$$\begin{aligned} \bar{N}'(r, x)I_B(P'_t(r, x), P'(r, x)) \geq \xi(r, x)/(1 + \eta) \ \&\ \ P'_t(r, x)y(r, x) \leq P'(r, x)y(r, x) \\ \implies P'_t(r, x)y(r, x) \leq Q'(r, x)y(r, x) \end{aligned} \quad (41)$$

where this holds because of the convexity and continuity of the function I_B . Hence,

$$\begin{aligned} &\text{Prob} \left[\bigcap_{x \in X} \bigcap_{r=1}^{k-1} \{B \cap C \cap (N'_t(r, x)I_B(P'_t(r, x), P'(r, x)) \geq \xi(r, x))\} \right] \\ &\leq \text{Prob} \left[\bigcap_{x \in X} \bigcap_{r=1}^{k-1} \{B \cap (P'_t(r, x) \cdot y(r, x) \leq Q'(r, x) \cdot y(r, x))\} \right] \\ &\leq \prod_{x \in X} \prod_{r=1}^{k-1} \exp(-\bar{N}'(r, x)I_B(Q'(r, x), P'(r, x))) \\ &\leq \exp(-\sum_{x \in X, r \in [k-1]} \xi(r, x)/(1 + \eta)). \end{aligned}$$

The first inequality follows immediately from our implication in Equation (41). To obtain the penultimate inequality, we apply Lemma 16. The ultimate inequality follows from Equation

(40). □

Lemma 14. Let \widehat{P} and P be two distributions on the event space $\mathfrak{R} = [1, \dots, k]$. Let P' and \widehat{P}' , defined on the event space $[1, \dots, k-1]$, be their Bernoulli transformations:

$$P'(r) = \frac{P(r)}{1 - \sum_{r' < r} P(r')}, \quad \widehat{P}'(r) = \frac{\widehat{P}(r)}{1 - \sum_{r' < r} \widehat{P}(r')}, \quad \forall r \in [1, \dots, k-1]. \quad (42)$$

Let $N \geq 0$ and with a slight abuse of notation, define $N(r) := N \cdot \widehat{P}(r)$ for all $r \in [1, \dots, k]$. Further define $N'(r) := N - \sum_{r' < r} N(r')$ for all $r \in [1, \dots, k-1]$. Then,

$$N \cdot I(\widehat{P}, P) = \sum_{r=1}^{k-1} N'(r) \cdot I_B(\widehat{P}'(r), P'(r)),$$

where we use the shorthand notation

$$I_B(\widehat{P}'(r), P'(r)) := I((\widehat{P}'(r), 1 - \widehat{P}'(r)), (P'(r), 1 - P'(r)))$$

to denote the information distance between two Bernoulli distributions.

Proof of Lemma 14. We will use a proof by induction on the support size $|\mathfrak{R}| = k$. The case where the support size $k = 2$ is trivial. For the sake of induction, we assume that the statement is true for distributions supported on $k-1$ points. By the chain rule, we have

$$\begin{aligned} & N \cdot I(\widehat{P}, P) \\ &= N \cdot I_B(\widehat{P}(1), P(1)) + N(1 - \widehat{P}(1)) \cdot I\left(\frac{(\widehat{P}(2), \dots, \widehat{P}(k))}{(1 - \widehat{P}(1))}, \frac{(P(2), \dots, P(k))}{(1 - P(1))}\right). \end{aligned} \quad (43)$$

Define $\widehat{P}_s = (\widehat{P}(2), \dots, \widehat{P}(k))/(1 - \widehat{P}(1))$ and $P_s = (P(2), \dots, P(k))/(1 - P(1))$ and observe that both are distributions supported on $k-1$ points. That is, $\sum_{r=2}^k \widehat{P}_s(r) = \sum_{r=2}^k P_s(r) = 1$. Further, Let $N_s := N \cdot (1 - \widehat{P}(1))$ and define $N_s(r) := N_s \widehat{P}_s(r)$ and $N'_s(r) := N_s - \sum_{2 \leq r' < r} N_s(r')$ for $r \in \{2, \dots, k\}$. With these definitions, and by applying the induction assumption, Equation (43) can be written as

$$N \cdot I(\widehat{P}, P) = N \cdot I_B(\widehat{P}(1), P(1)) + \sum_{r=2}^{k-1} N'_s(r) \cdot I_B(\widehat{P}'_s(r), P'_s(r))$$

where

$$\begin{aligned}\widehat{P}'_s(r) &= \frac{\widehat{P}_s(r)}{1 - \sum_{2 \leq r' < r} \widehat{P}_s(r')} = \frac{\widehat{P}(r)}{(1 - P(1))(1 - \sum_{2 \leq r' < r} \widehat{P}_s(r)/(1 - P(1)))} \\ &= \frac{\widehat{P}(r)}{(1 - P(1) - \sum_{2 \leq r' < r} \widehat{P}_s(r))} = \frac{\widehat{P}(r)}{1 - \sum_{r' < r} \widehat{P}_s(r)} = \widehat{P}'(r)\end{aligned}$$

for all $r \in [2, \dots, k]$. Here, the last equality follows from Equation (42). The same argument can be made to argue that $P'_s(r) = P'(r)$ for all $r \in [2, \dots, k]$. So far, we established that

$$N \cdot I(\widehat{P}, P) = N \cdot I_B(\widehat{P}(1), P(1)) + \sum_{r=2}^{k-1} N'_s(r) \cdot I_B(\widehat{P}'(r), P'(r)),$$

To complete the proof, it suffices to show that for any $r \in [2, \dots, k]$, $N'_s(r) = N'(r)$:

$$\begin{aligned}N'_s(r) &= N_s - \sum_{2 \leq r' < r} N_s(r') \\ &= (1 - \widehat{P}(1)) \cdot (N - N \sum_{2 \leq r' < r} \widehat{P}_s(r')) \\ &= N(1 - \widehat{P}(1)) \cdot (1 - \sum_{2 \leq r' < r} \widehat{P}_s(r')) \\ &= N(1 - \widehat{P}(1)) \cdot (1 - \sum_{2 \leq r' < r} \widehat{P}(r')/(1 - \widehat{P}(1))) \\ &= N \cdot (1 - \widehat{P}(1) - \sum_{2 \leq r' < r} \widehat{P}(r')) \\ &= N \cdot (1 - \sum_{r' < r} \widehat{P}(r')) = N'(r).\end{aligned} \quad \square$$

Lemma 15 (Lemma 8 of Magureanu et al. (2014a)). *Let $a > 0$, $d \geq 2$, and $Z \in \mathbb{R}_+^d$ be a random variable such that for any $\xi \in \mathbb{R}_+^d$, we have that*

$$\text{Prob}[\cap_{i=1}^d (Z_i > \xi_i)] \leq \exp(-a \cdot \langle \mathbf{1}_d, \xi \rangle).$$

Then, for all $\delta \geq \frac{d}{a} > 0$ we have

$$\text{Prob}[\langle \mathbf{1}_d, Z \rangle \geq \delta] \leq \left(\frac{a\delta e}{d}\right)^d \exp(-a\delta),$$

where $\mathbf{1}_d$ is a d -dimensional all-ones vector.

Proof of Lemma 15. See the proof of Lemma 8 in Magureanu et al. (2014a). \square

Lemma 16. *For any $x \in X$ and $r \in [1, \dots, k-1]$, let $\bar{N}'(r, x) \in [0, \dots, t]$ and $y(r, x) \in \{-1, 1\}$. Consider the events*

$$B := \cap_{x \in X} \cap_{r=1}^{k-1} \{ \bar{N}'(r, x) \leq N'_t(r, x) \leq (1 + \eta) \bar{N}'(r, x) \}$$

$$C := \cap_{x \in X} \cap_{r=1}^{k-1} \{ P'_t(r, x) y(r, x) \leq Q'(r, x) y(r, x) \}$$

for $y(r, x) \in \{+1, -1\}$ for all $r \in [1, \dots, k-1]$ and $x \in X$. Assume that $Q'(r, x) y(r, x) \leq P'(r, x) y(r, x)$ for all $r \in [1, \dots, k-1]$ and $x \in X$. Then,

$$\text{Prob}[B \cap C] \leq \prod_{x \in X} \prod_{r=1}^{k-1} \exp(-\bar{N}'(r, x) I_B(Q'(r, x), P'(r, x))).$$

Proof of Lemma 16. We start by introducing a log moment-generating function related to the rewards for each arm. In particular, for all $x \in X$, $r \in [1, \dots, k-1]$, and λ , we define

$$\phi(r, x; \lambda) := \log(\mathbb{E}[\exp(\lambda \cdot \mathbb{1}\{R_t(x) = r\}) | R_t(x) \geq r]).$$

Evidently, we have that these functions $\phi(r, x)$ are not a function of time t as the reward distributions are identically distributed over time. Recall that

$$\begin{aligned} \text{Prob}[R_t(x) = r | R_t(x) \geq r] &= \frac{\text{Prob}[R_t(x) = r \ \& \ R_t(x) \geq r]}{\text{Prob}[R_t(x) \geq r]} = \frac{\text{Prob}[R_t(x) = r]}{\text{Prob}[R_t(x) \geq r]} = \frac{P(r, x)}{\sum_{r' \geq r} P(r', x)} \\ &= P'(r, x) \end{aligned}$$

for all $r \in [1, \dots, k-1]$ and $x \in X$. Therefore, we have that $\phi(r, x; \lambda)$ is the moment-generating function of a Bernoulli variable with success parameter $P'(r, x)$. That is,

$$\phi(r, x; \lambda) = \log(P'(r, x) \exp(\lambda) + (1 - P'(r, x))).$$

It is also easy to show that

$$I_B(Q'(r, x), P'(r, x)) = \max_{\lambda: \lambda y(r, x) \leq 0} \{Q'(r, x) - \phi(r, x; \lambda)\} \quad (44)$$

as indeed by assumption $Q'(r, x) y(r, x) \leq P'(r, x) y(r, x)$. That is, the relative entropy of a Bernoulli distribution is the Fenchel dual of its log moment-generating function. We then

define

$$\lambda^*(r, x) := \arg \max_{\lambda: \lambda y(r, x) \leq 0} \{Q'(r, x)\lambda - \phi(r, x; \lambda)\}, \quad (45)$$

and let

$$\begin{aligned} G_t &:= \exp \left(\sum_{x \in X} \sum_{r=1}^{k-1} (\lambda^*(r, x) \sum_{\tau=1}^{t-1} \mathbb{1}\{x_\tau = x, R_\tau(x) = r\} - \sum_{\tau=1}^{t-1} \mathbb{1}\{x_\tau = x, R_\tau(x) \geq r\}) \phi(r, x; \lambda^*(r, x)) \right) \\ &= \exp \left(\sum_{x \in X} \sum_{r=1}^{k-1} (\lambda^*(r, x) N_t(r, x) - N'_t(r, x) \phi(r, x; \lambda^*(r, x))) \right), \end{aligned} \quad (46)$$

where $N'_t(r, x) = \sum_{\tau=1}^{t-1} \mathbb{1}\{x_\tau = x, R_\tau(x) \geq r\}$, $r \in [1, \dots, k-1]$, $x \in X$, is the number of times before round t that arm x is played and we receive of reward of at least r . We will argue later that $\mathbb{E}[G_t] = 1$. Now assume that is indeed the case, then

$$\begin{aligned} \text{Prob}[B \cap C] &= \text{Prob} \left[\cap_{x \in X} \cap_{r=1}^{k-1} \{P'_t(r, x) y(r, x) \leq Q'(r, x) y(r, x), B\} \right] \\ &= \text{Prob} \left[\cap_{x \in X} \cap_{r=1}^{k-1} \left\{ \frac{N_t(r, x)}{N'_t(r, x)} y(r, x) \leq Q'(r, x) y(r, x), B \right\} \right] \\ &= \text{Prob} \left[\cap_{x \in X} \cap_{r=1}^{k-1} \{N_t(r, x) y(r, x) \leq Q'(r, x) y(r, x) N'_t(r, x), B\} \right]. \end{aligned}$$

Note in particular that the last characterization of $\text{Prob}[B \cap C]$ in fact allows for the event $N'_t(r, x) = 0$. Using the fact that, we have $\lambda^*(r, x) y(r, x) \leq 0$, we have

$$\begin{aligned} &\text{Prob}[B \cap C] \\ &\leq \text{Prob} \left[\sum_{x \in X} \sum_{r=1}^{k-1} \lambda^*(r, x) y^2(r, x) N_t(r, x) \geq \sum_{x \in X} \sum_{r=1}^{k-1} \lambda^*(r, x) y^2(r, x) Q'(r, x) N'_t(r, x), B \right] \\ &= \text{Prob} \left[\mathbb{1}\{B\} \exp \left(\sum_{x \in X} \sum_{r=1}^{k-1} \lambda^*(r, x) N_t(r, x) \right) \geq \exp \left(\sum_{x \in X} \sum_{r=1}^{k-1} \lambda^*(r, x) Q'(r, x) N'_t(r, x) \right) \right], \end{aligned}$$

where the last equations follows from the monotonously of the exponential function and the fact that $y^2(r, x) = 1$ as $y(r, x) = \{-1, 1\}$ for all $x \in X$ and $r \in [1, \dots, k-1]$. Then, the

definition of G_t gives us

$$\begin{aligned}
& \text{Prob}[B \cap C] \\
& \leq \text{Prob}\left[B \cap \left(G_t \geq \exp\left(\sum_{x \in X} \sum_{r=1}^{k-1} (\lambda^*(r, x) Q'(r, x) - \phi(r, x; \lambda^*(r, x))) N'_t(r, x)\right)\right)\right] \\
& = \text{Prob}\left[B \cap \left(G_t \geq \exp\left(\sum_{x \in X} \sum_{r=1}^{k-1} I_B(Q'(r, x), P'(r, x)) \cdot N'_t(r, x)\right)\right)\right] \\
& = \text{Prob}\left[G_t \cdot \mathbb{1}\{B\} \geq \exp\left(\sum_{x \in X} \sum_{r=1}^{k-1} I_B(Q'(r, x), P'(r, x)) \cdot N'_t(r, x)\right)\right] \\
& \leq \text{Prob}\left[G_t \cdot \mathbb{1}\{B\} \geq \exp\left(\sum_{x \in X} \sum_{r=1}^{k-1} I_B(Q'(r, x), P'(r, x)) \cdot \bar{N}'(r, x)\right)\right] \\
& \leq \mathbb{E}[G_t \cdot \mathbb{1}\{B\}] \exp\left(-\sum_{x \in X} \sum_{r=1}^{k-1} I_B(Q'(r, x), P'(r, x)) \cdot \bar{N}'(r, x)\right), \\
& \leq \exp\left(-\sum_{x \in X} \sum_{r=1}^{k-1} I_B(Q'(r, x), P'(r, x)) \cdot \bar{N}'(r, x)\right),
\end{aligned}$$

where the second equation follows from definition of λ^* , given in Equation (45), and the fact that $I_B(Q'(r, x), P'(r, x)) = \max_{\lambda: \lambda y(r, x) \leq 0} \{Q'(r, x) - \phi(r, x; \lambda)\}$ per Equation (44). The first inequality holds because under event B , the random variable $N'_t(r, x)$ is greater than or $\bar{N}'(r, x)$, and the second inequality follows from Markov's inequality. Finally, the last inequality follows from our earlier claim that $\mathbb{E}[G_t] = 1$ which implies $\mathbb{E}[G_t \cdot \mathbb{1}\{B\}] \leq 1$.

We still need to show that $\mathbb{E}[G_t] = 1$. Recall that G_t is a function of random variables $R_\tau(x_\tau)$ for $\tau \in [1, \dots, t]$; see Equation (46). In this following, for every $x \in X$, and $\tau \in [1, \dots, t]$ we define another random variable $R'_\tau(x)$ whose distribution is the same as $R_\tau(x)$. To construct $R'_\tau(x)$, we first define $k - 1$ independent Bernoulli random variables, $b_r(r, x) \sim \text{Bernoulli}(P'(r, x))$, $r \in [1, \dots, k-1]$, where $\text{Bernoulli}(p)$ is a Bernoulli random variable with a success rate of p . Then, if $b_r(r, x) = 0$ for all $r \in [1, \dots, k-1]$, we set $R'_\tau(x)$ to k . Otherwise, $R'_\tau(x) = \min\{r : b_r(r, x) = 1\}$. Note that for any τ_1 and τ_2 and $x \in X$, distribution of $R'_{\tau_1}(x)$ is the same as that of $R'_{\tau_2}(x)$, and the random variables $R'_{\tau_1}(x)$ and $R'_{\tau_2}(x)$ are independent of each other. Lemma 17, presented at the end of this section shows that for any $x \in X$ and $\tau \geq 0$, the distribution of $R'_\tau(x)$ is the same as $R_\tau(x)$. Consequently, the random variable

$$\begin{aligned}
& G'_t \\
& = \exp\left(\sum_{x \in X} \sum_{r=1}^{k-1} \left(\lambda^*(r, x) \sum_{\tau=1}^{t-1} \mathbb{1}\{x_\tau = x, R'_\tau(x) = r\} - \sum_{\tau=1}^{t-1} \mathbb{1}\{x_\tau = x, R'_\tau(x) \geq r\} \phi(r, x; \lambda^*(r, x))\right)\right)
\end{aligned}$$

has the same distribution as G_t . Hence, we will show that $\mathbb{E}[G_t] = \mathbb{E}[G'_t] = 1$ using a martingale argument.

Define the following filtrations

$$\begin{aligned}\mathcal{F}_\tau^{\pi'}(0) &:= \sigma((x_1, R'_1(x_1)), \dots, (x_{\tau-1}, R'_{\tau-1}(x_{\tau-1}))) \quad \forall \tau \in [1, \dots, t], \\ \mathcal{F}_\tau^{\pi'}(r) &:= \sigma((x_1, R'_1(x_1)), \dots, (x_{\tau-1}, R'_{\tau-1}(x_{\tau-1})), (x_\tau, b_\tau(1, x_\tau), \dots, b_\tau(r, x_\tau))) \\ &\quad \forall \tau \in [1, \dots, t], \quad r \in [1, \dots, k-1],\end{aligned}$$

where $b_\tau(r, x)$, $r \in [1, \dots, k-1]$, are the Bernoulli variables used to construct the rewards $R'_\tau(x)$ in round τ . Observe that $\mathcal{F}_\tau^{\pi'}(r)$ is independent from the Bernoulli variables $b_\tau(r', x_\tau)$ for $r' > r$. Evidently, we have the inclusions

$$\begin{aligned}\mathcal{F}_\tau^{\pi'}(r-1) &\subseteq \mathcal{F}_\tau^{\pi'}(r) \quad \forall \tau \in [1, \dots, t], \quad r \in [1, \dots, k-1], \\ \mathcal{F}_\tau^{\pi'}(k-1) &= \mathcal{F}_{\tau+1}^{\pi'}(0) \quad \forall \tau \in [1, \dots, t-1].\end{aligned}$$

The first inclusion follows immediately by construction. The second equality follows from the fact that x_τ is $\mathcal{F}_\tau^{\pi'}(0) \subseteq \mathcal{F}_\tau^{\pi'}(k-1)$ measurable and $R'_\tau(x_\tau)$ is revealed once the random variables $b_\tau(r, x_\tau)$, $r \in [1, \dots, k-1]$, are revealed and hence is $\mathcal{F}_\tau^{\pi'}(k-1)$ -measurable as well. Hence, the filtration sequence $\mathcal{F}_\tau^{\pi'}(r)$ for $\tau \in [1, \dots, t]$ and $r \in [0, \dots, k-1]$ is increasing; more precisely we have

$$\mathcal{F}_1^{\pi'}(0) \subseteq \mathcal{F}_1^{\pi'}(1) \subseteq \dots \subseteq \mathcal{F}_1^{\pi'}(k-1) = \mathcal{F}_2^{\pi'}(0) \subseteq \mathcal{F}_2^{\pi'}(1) \subseteq \dots \subseteq \mathcal{F}_t^{\pi'}(k-2) \subseteq \mathcal{F}_t^{\pi'}(k-1). \quad (47)$$

Define for any round $\tau \in [1, \dots, t]$ the random variables

$$\begin{aligned}G'_\tau(0) &= G'_\tau, \\ G'_\tau(r) &= G'_\tau \cdot \prod_{x \in X, r' \leq r} \exp\left(\mathbb{1}\{x_\tau = x\} \left[\lambda^*(r', x) \mathbb{1}\{R'_\tau(x) = r'\} - \mathbb{1}\{R'_\tau(x) \geq r'\} \phi(r', x; \lambda^*(r', x)) \right]\right) \\ &\quad \forall r \in [1, \dots, k-1].\end{aligned}$$

We will show that $G'_\tau(r)$ is a martingale with respect to the filtration sequence stated in Equation (47) and thus $\mathbb{E}[G'_t(r)] = \mathbb{E}[G'_1(0)] = 1$, which is the desired result.

First observe that for any round $\tau \in [1, \dots, t]$ we have

$$\mathbb{E}[G'_\tau(0) | \mathcal{F}_{\tau-1}^{\pi'}(k-1)] = \mathbb{E}[G'_\tau | \mathcal{F}_\tau^{\pi'}(0)] = \mathbb{E}[G'_{\tau-1}(k-1) | \mathcal{F}_\tau^{\pi'}(0)] = G'_{\tau-1}(k-1).$$

Furthermore, for any round $\tau \in [1, \dots, t]$ and $r \in [1, \dots, k-1]$, we have

$$\begin{aligned}
& \mathbb{E} [G'_\tau(r) | \mathcal{F}_\tau^{\pi'}(r-1)] \\
&= G'_\tau \cdot \mathbb{E} \left[\prod_{r' \leq r} \exp \left(\lambda^*(r', x_\tau) \mathbb{1}\{R'_\tau(x_\tau) = r'\} \right. \right. \\
&\quad \left. \left. - \mathbb{1}\{R'_\tau(x_\tau) \geq r'\} \phi(r', x_\tau; \lambda^*(r', x_\tau)) \right) \middle| \mathcal{F}_\tau^{\pi'}(r-1) \right] \\
&= G'_\tau \prod_{r' \leq r-1} \exp(\lambda^*(r', x_\tau) \mathbb{1}\{R'_\tau(x_\tau) = r'\} - \mathbb{1}\{R'_\tau(x_\tau) \geq r'\} \phi(r', x_\tau; \lambda^*(r', x_\tau))) \\
&\quad \cdot \mathbb{E} [\exp(\lambda^*(r, x_\tau) \mathbb{1}\{R'_\tau(x_\tau) = r\} - \mathbb{1}\{R'_\tau(x_\tau) \geq r\} \phi(r, x_\tau; \lambda^*(r, x_\tau))) \mid \mathcal{F}_\tau^{\pi'}(r-1)] \\
&= G'_\tau(r-1) \cdot \frac{\mathbb{E} [\exp(\lambda^*(r, x_\tau) \mathbb{1}\{R'_\tau(x_\tau) = r\}) \mid \mathcal{F}_\tau^{\pi'}(r-1)]}{\exp(\mathbb{1}\{R'_\tau(x_\tau) \geq r\} \phi(r, x_\tau; \lambda^*(r, x_\tau)))} \\
&= G'_\tau(r-1).
\end{aligned}$$

Here, the first equation follows from the fact that x_τ is $\mathcal{F}_\tau^{\pi'}(0) \subseteq \mathcal{F}_\tau^{\pi'}(r-1)$ measurable. The second equation follows from the fact that both $\mathbb{1}\{R'_\tau(x_\tau) = r'\}$ and $\mathbb{1}\{R'_\tau(x_\tau) \geq r'\}$ are $\mathcal{F}_\tau^{\pi'}(r-1)$ -measurable for $r' \leq r-1$. The third equation can be established as by remarking that $\mathbb{1}\{R'_\tau(x_\tau) \geq r\}$ is also $\mathcal{F}_\tau^{\pi'}(r-1)$ measurable. To establish the last equality we consider two cases. Conditioned on the event $\mathbb{1}\{R'_\tau(x_\tau) \geq r\} = 0$, we have $\mathbb{1}\{R'_\tau(x_\tau) = r\} = 0$ as well, and thus

$$\mathbb{1}\{R'_\tau(x_\tau) \geq r\} = 0 \implies \frac{\mathbb{E} [\exp(\lambda^*(r, x_\tau) \mathbb{1}\{R'_\tau(x_\tau) = r\}) \mid \mathcal{F}_\tau^{\pi'}(r-1)]}{\exp(\mathbb{1}\{R'_\tau(x_\tau) \geq r\} \phi(r, x_\tau; \lambda^*(r, x_\tau)))} = 1.$$

Conditioned on the event $\mathbb{1}\{R'_\tau(x_\tau) \geq r\} = 1$, we have that $\mathbb{1}\{R'_\tau(x_\tau) = r\}$ is by construction distributed as b_{r, x_τ} a Bernoulli random variable independent from $\mathcal{F}_\tau^{\pi'}(r-1)$. Hence,

$$\begin{aligned}
\mathbb{1}\{R'_\tau(x_\tau) \geq r\} = 1 &\implies \mathbb{E} [\exp(\lambda^*(r, x_\tau) \mathbb{1}\{R'_\tau(x_\tau) = r\}) \mid \mathcal{F}_\tau^{\pi'}(r-1)] \\
&= \mathbb{E} [\phi(r, x_\tau; \lambda^*(r, x_\tau)) \mid \mathcal{F}_\tau^{\pi'}(r-1)]
\end{aligned}$$

and we have

$$\mathbb{E} [\exp(\lambda^*(r, x_\tau) \mathbb{1}\{R'_\tau(x_\tau) = r\}) \mid \mathcal{F}_\tau^{\pi'}(r-1)] / \exp(\mathbb{1}\{R'_\tau(x_\tau) \geq r\} \phi(r, x_\tau; \lambda^*(r, x_\tau))) = 1.$$

Hence, $G'_\tau(r)$ is a martingale for the sequence $\mathcal{F}_\tau^{\pi'}(r)$ as stated in Equation (47). \square

Lemma 17. *Let $R \sim P$ be a random variable with finite support $\mathfrak{R} = \{1, 2, \dots, k\}$. Define distribution $P'(r) = \frac{P(r)}{\sum_{r' \geq r} P(r')}$, $r \in [1, \dots, k-1]$, and $k-1$ independent Bernoulli random*

variables, $b(r) \sim \text{Bernoulli}(P'(r))$, $r \in [1, \dots, k-1]$, where $\text{Bernoulli}(p)$ is a Bernoulli random variable with a success rate of p . Let, the random variable R' be

$$R' = \begin{cases} k & \text{if } \sum_{r=1}^{k-1} b(r) = 0; \\ \min\{r : b(r) = 1\} & \text{otherwise.} \end{cases}$$

Then, distribution of R' is the same as that of R .

Proof of Lemma 17. The proof makes use of the following claim.

Claim. For any $r \in [1, \dots, k]$, we have

$$\prod_{r' < r} (1 - P'(r')) = 1 - \sum_{r' < r} P(r') \quad \text{and} \quad P(r) = \prod_{r' < r} (1 - P'(r')) \cdot P'(r) .$$

This claim completes the proof because for any $r \in [k]$, we have

$$\text{Prob}[R' = r] = \prod_{r' < r} (1 - P'(r')) \cdot P'(r) = (1 - \sum_{r' < r} P(r')) P'(r) = P(r) = \text{Prob}[R = r] ,$$

where the first equation holds because $R' = r$ when $b(r') = 0$, $r' < r$, and $b(r) = 1$; the second equation follows from the claim; the third equation holds because by definition, $P'(r) = \frac{P(r)}{(1 - \sum_{r' < r} P(r'))}$.

Proof of claim. We will prove the claim using induction on r . Consider the base case of $r = 1$. We have $\prod_{r' < 1} (1 - P'(r')) = 1 - \sum_{r' < 1} P(r') = 1$ and $\prod_{r' < 1} (1 - P'(r')) \cdot P'(r) = P'(1) = P(1)$.

Assume for the sake of induction that $\prod_{r' < r-1} (1 - P'(r')) = 1 - \sum_{r' < r-1} P(r')$ and $P(r-1) = \prod_{r' < r-1} (1 - P'(r')) \cdot P'(r-1)$. First, we have that

$$\begin{aligned} \prod_{r' < r} (1 - P'(r')) &= (1 - \sum_{r' < r-1} P(r')) (1 - P'(r-1)) \\ &= 1 - \sum_{r' < r-1} P(r') - (1 - \sum_{r' < r-1} P(r')) P'(r-1) \\ &= 1 - \sum_{r' < r-1} P(r') - P(r-1) = 1 - \sum_{r' < r} P(r) , \end{aligned}$$

where the first equality follows from the induction hypothesis and the third equality follows from definition of $P'(\cdot)$. We have then also have that

$$\prod_{r' < r} (1 - P'(r')) P'(r) = (1 - \sum_{r' < r} P(r)) P'(r) = P(r) ,$$

where the first equality follows from the already established earlier relationship $\prod_{r' < r} (1 -$

$P'(r') = (1 - \sum_{r' < r} P(r))$ and the second equality follows again from definition of P' . \square

E Proof of Statements in Section 7.2

E.1 Proof of Lemma 5

By definition of the suboptimality gap of arms, when we are in the exploitation phase, i.e., \mathcal{E}_t holds, and the bad events $\bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)$ happen, the regret can be bounded as

$$\begin{aligned} \sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t, \bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x, P) &\leq \sum_{|X| < t \leq T} \sum_{x \in X} \mathbb{1}\{\mathcal{E}_t, x \in x^*(P_t), \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x, P) \\ &= \sum_{x \in X} \sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t, x \in x^*(P_t), \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x, P) \\ &= \sum_{x \in X} \sum_{1 \leq s \leq T} \mathbb{1}\{\mathcal{B}_{\tau(s,x)}^1(x, \kappa), \tau(s, x) \leq T\} \cdot \Delta(x, P), \end{aligned}$$

where the first inequality holds because we pull one of empirically optimal arms, i.e., $x \in x^*(P_t)$. Here, the stopping time $\tau(s, x) = \min\{\tau \in (|X|, \dots, T] : \sum_{|X| < t \leq \tau} \mathbb{1}\{\mathcal{E}_t, x \in x^*(P_t)\} \geq s\}$ indicates the first time in the exploitation phase that we deemed a certain arm x empirically optimal at least s times. When such an event never happens, we set $\tau(s, x) = T + 1$. The expected regret caused during the event of interest is hence bounded by

$$\begin{aligned} &\mathbb{E} \left[\sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t, \bigcup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x_t, P) \right] \\ &\leq \sum_{x \in X, r \in \mathfrak{R}} \sum_{1 \leq s \leq T} \text{Prob} \left[|P_{\tau(s,x)}(r, x) - P(r, x)| > \kappa \right] \cdot \Delta(x, P). \quad (48) \\ &\leq \sum_{1 \leq s \leq T} \sum_{x \in X, r \in \mathfrak{R}} \text{Prob} \left[|P_{\tau(s,x)}(r, x) - P(r, x)| > \kappa \right] \cdot \Delta(x, P). \end{aligned}$$

Now consider a stopping time $\tau(s, x)$ for some $1 \leq s \leq T$. We know that by construction of the stopping time $\tau(s, x)$, arm x was empirically optimal at least $s - 1$ times prior to round t . We will argue that $N_{\tau(s,x)}(x) > s/(4|X|)$ for $s \geq 8|X|$. For the sake of contradiction assume that $N_{\tau(s,x)}(x) \leq s/(4|X|)$. This implies that for at least $E \leq s - 1 - s/(4|X|)$ rounds $\{t_i\}$ prior to time $\tau(s, x)$ we had that there was an empirically optimal arm $x' \neq x$ such that $N_{t_i}(x') \leq N_{t_i}(x) \leq N_{\tau(s,x)}(x) \leq s/(4|X|)$. As in each of these rounds $N_{t_i}(x')$ is incremented by one there can only have been at most $E \leq (|X| - 1) \lceil s/(4|X|) \rceil \leq s/4 + |X|$ of such rounds. Hence, we must have

$$s/4 + |X| \geq E \geq s - 1 - s/(4|X|)$$

which is contradiction with $|X| \geq 1$ and $s \geq 8|X|$. Considering this and per Equation (48), we further upper bound our regret as follows:

$$\begin{aligned} & \mathbb{E} \left[\sum_{|X| < t \leq T} \mathbb{1}\{\mathcal{E}_t, \cup_{x \in x^*(P_t)} \mathcal{B}_t^1(x, \kappa)\} \cdot \Delta(x_t, P) \right] \\ & \leq 8|X| + |X| |\mathfrak{A}| \cdot \sum_{1 \leq s \leq T} \text{Prob} \left[|P_{\tau(s,x)}(r, x) - P(r, x)| > \kappa, N_{\tau(s,x)}(x) > s/(4|X|) \right] \end{aligned}$$

Applying the concentration bound in Lemma 3 gives

$$\begin{aligned} & |X| |\mathfrak{A}| \cdot \sum_{1 \leq s \leq T} \text{Prob} \left[|P_{\tau(s,x)}(r, x) - P(r, x)| > \kappa, N_{\tau(s,x)}(x) > s/(4|X|) \right] \\ & \leq |X| |\mathfrak{A}| \cdot \sum_{1 \leq s \leq T} 2 \exp(-s\kappa^2/(2|X|)) \end{aligned}$$

from which the claimed result follows. \square

E.2 Proof of Lemma 6

To show the result, we consider a reward distribution \bar{Q} such that $\bar{Q}(x^*(P_t)) = P_t(x^*(P_t))$ and $\bar{Q}(\tilde{X}(P_t)) = P(\tilde{X}(P_t))$ as visualized in Figure 5. Observe that the reward distribution \bar{Q} agrees with the actual reward distribution P at all the arms except for the empirically optimal arms. Now, considering our assumption that $P \in \text{int}(\mathcal{P})$ with a unique optimal arm, we can choose κ sufficiently small so that $\|\bar{Q}(x) - P(x)\|_\infty < \kappa$ for all $x \in x^*(P_t)$ implies (i) $\bar{Q} \in \text{int}(\mathcal{P})$ and (ii) $x^*(P) = x^*(\bar{Q})$ given that $\bar{Q}(x) = P(x)$ for any $x \in \tilde{X}(P_t)$; see Lemma 7. That is, if two distributions P and \bar{Q} are equal everywhere except at arms $x^*(P_t)$ and $\|\bar{Q}(x) - P(x)\|_\infty < \kappa$ for all $x \in x^*(P_t)$, then they are sufficiently close to have the same optimal arm and both be in the interior of \mathcal{P} . Given this, it is easy to see that distribution \bar{Q} is indeed deceitful, i.e., $\bar{Q} \in \text{DECEIT}(x^*(P), P_t)$, and that $x^*(P) \in \tilde{X}_d(P_t)$ as $\text{DECEIT}(x^*(P), P_t)$ is non-empty. Note that by Equation (3), $\bar{Q} \in \text{DECEIT}(x', P_t)$ if (i) $\bar{Q} \in \mathcal{P}$, (ii) $\bar{Q}(x^*(P_t)) = P_t(x^*(P_t))$, and (iii) $\sum_{r \in \mathfrak{R}} r \bar{Q}(r, x^*(P)) > \sum_{r \in \mathfrak{R}} r \bar{Q}(r, x)$ for some $x \in x^*(P_t)$. The two first conditions hold by construction and the final condition holds because the optimal arm under \bar{Q} is $x^*(P) \notin x^*(P_t)$; again see Figure 5. The constructed distribution $\bar{Q} \in \text{DECEIT}(x^*(P), P_t)$ guarantees that (26) occurred as

$$\begin{aligned} \sum_{x \in X} N_t(x) I(P_t(x), P(x)) & \geq \sum_{x \in \tilde{X}(P_t)} N_t(x) I(P_t(x), P(x)) \\ & \geq \sum_{x \in X} N_t(x) I(P_t(x), \bar{Q}(x)) \geq (1 + \epsilon) \log(t), \end{aligned}$$

where the first inequality follows from the construction of \bar{Q} (i.e., $\bar{Q}(x^*(P_t)) = P_t(x^*(P_t))$) and the second inequality follows from Equation (25) and the fact that $\bar{Q} \in \text{DECEIT}(x^*(P), P_t)$. \square

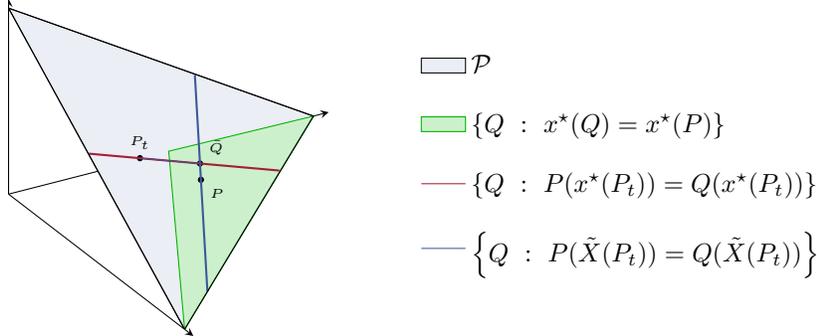


Figure 5: Construction of the reward distribution \bar{Q} . By assumption P is in the interior of both \mathcal{P} as well as $\{Q \in \mathcal{P} : x^*(Q) = x^*(P)\}$ as its optimal arm is unique.

E.3 Proof of Lemma 7

We first show that $\mathcal{P}_u(x_0^*)$ is an open set. By definition, for a given x_0^* , evidently we have that

$$\mathcal{P}_u(x_0^*) = \left\{ Q : 0 < \sum_{r \in \mathfrak{R}} Q(r, x_0^*) - \sum_{r \in \mathfrak{R}} Q(r, x) \quad \forall x \neq x_0^* \in X \right\}$$

is open as the functions $\sum_{r \in \mathfrak{R}} Q(r, x_0^*) - \sum_{r \in \mathfrak{R}} Q(r, x)$ are linear and hence continuous in Q .

Next, we show the second result of the lemma. Using the result we just showed, it follows that there exists a neighborhood \mathcal{N} such that

$$Q \in \mathcal{N} \implies x^*(Q) = x^*(P) =: x_0^*.$$

whenever we have $P \in \mathcal{P}_u(x_0^*)$. The maximum reward of arm x' for any $P \in \mathcal{N} \cap \mathcal{P}$ in this neighborhood hence simplifies to

$$\begin{aligned} \text{REW}_{\max}(x', Q) &:= \max_{Q'} \sum_{r \in \mathfrak{R}} r Q'(r, x') \\ \text{s.t. } &Q' \in \mathcal{P}, \\ &Q'(x_0^*) = Q(x_0^*). \end{aligned} \tag{49}$$

We wish to prove that $\text{REW}_{\max}(x', Q)$ is continuous at P by invoking Proposition 10 to its maximization characterization in stated in Equation (49). In order to do so we verify that all required conditions are met.

We first remark that our feasible set mapping is indeed lower semi-continuous.

Lemma 18. *The mapping $M(Q) := \{Q' \in \mathcal{P} : Q'(x_0^*) = Q(x_0^*)\}$ is lower semi-continuous at any $P \in \text{int}(\mathcal{P})$.*

Evidently, the set $M(Q)$ is always convex as it is an affine set for any $Q \in \mathcal{P}$. The mapping $M(Q)$ is closed at P as its graph $\{(Q, Q') : Q \in \mathcal{N}', Q' \in \mathcal{P}, Q'(x_0^*) = Q(x_0^*)\}$ is closed for all closed neighborhoods \mathcal{N}' containing P .

The objective function in the maximization problem stated in Equation (49) is continuous at (Q', Q) . The set of maximizers is always non-empty and bounded as its feasible set at $Q = P$, i.e., $M(P)$, is furthermore compact and non-empty. As the objective function of the maximization problem (49) is also convex in Q' , Proposition 10 implies the continuity of its maximum $\text{REW}_{\max}(x', Q)$ at P . \square

Proof of Lemma 18. The mapping is clearly nonempty as $Q \in M(Q)$. We can prove the lower-semicontinuity of this mapping at P directly from its definition. Consider indeed an arbitrary open set V such that $M(P) \cap V \neq \emptyset$. It remains to show that the set $\{Q : M(Q) \cap V \neq \emptyset\}$ contains an open neighborhood around P . From Section B, we know that without loss of generality we may assume that

$$V = \{Q' : \sum_{r \in \mathfrak{R}} Q'(r) = 1, \|Q' - \bar{Q}\|_\infty < \epsilon\}$$

for some $\epsilon > 0$ and $\bar{Q} \in \mathcal{P}_\Omega$. Hence, there exists U with

$$U \in \mathcal{P}, \quad \|U(x) - \bar{Q}(x)\|_\infty < \epsilon \quad \forall x \in X \quad \text{and} \quad U(x_0^*) = P(x_0^*).$$

In fact as $P \in M(P)$ and $P \in \text{int}(\mathcal{P})$ we may assume without loss of generality that $U \in \text{int}(\mathcal{P})$ as well. Indeed, assume this is not the case then we may simply consider the perturbation $(1 + \theta)U + \theta P$ instead for a sufficiently small $\theta > 0$ exploiting the convexity of $M(P)$ and the fact that V is an open set. Define now the parametric distribution $U_Q(x) = U(x)$ for $x \in \tilde{X}_0$ and $U_Q(x_0^*) = Q(x_0^*)$ for any Q . Clearly, $U_P = U$. By construction, we hence have $P \in \{Q : U_Q \in M(Q) \cap V\}$. Furthermore,

$$\begin{aligned} & \{Q : U_Q \in M(Q) \cap V\} \\ &= \{Q : U_Q \in \mathcal{P}\} \cap \{Q : \|Q(x_0^*) - \bar{Q}(x_0^*)\|_\infty < \epsilon\} \end{aligned}$$

The first set in the previous intersection contains an open neighborhood around P as $U_P = U \in \text{int}(\mathcal{P})$ and U_Q is a linear and hence continuous function in Q . The second set is an open set as

$\|Q(x_0^*) - \bar{Q}(x_0^*)\|_\infty$ is a continuous function in Q . As clearly we have $\{Q : U_Q \in M(Q) \cap V\} \subseteq \{Q : M(Q) \cap V \neq \emptyset\}$ the claim follows immediately. \square

F Proof of Statements in Section 7.3

F.1 Proof of Lemma 8

For the sake of contradiction, assume that this is false and hence $\min_{x \in X} N_t(x) < \epsilon \lceil s_t/4 \rceil$. When \mathcal{X}_t occurs then $\tau(s_t) = t$. We now argue that there must exist $\lceil 3s_t/4 \rceil$ rounds in the first t rounds in which $\min_{x \in X} N_t(x) \leq \epsilon s_t$. We have that $\min_{x \in X} N_{\tau(i)}(x) \leq s_{\tau(i)}\epsilon$ for any $i \in \{\lceil s_t/4 \rceil, \dots, s_t\}$. This is because for such rounds $i \in \{\lceil s_t/4 \rceil, \dots, s_t\}$ we have (i) $\min_{x \in X} N_{\tau(i)}(x) \leq \min_{x \in X} N_t(x)$, (ii) $s_{\tau(i)}\epsilon = i\epsilon \geq \lceil s_t/4 \rceil\epsilon$, and (iii) by our assumption, $\min_{x \in X} N_t(x) < \epsilon \lceil s_t/4 \rceil$. So far we established that there exist at least $\lceil 3s_t/4 \rceil \leq s_t - \lceil s_t/4 \rceil$ rounds before time t in which $\min_{x \in X} N_t(x) \leq \epsilon s_t$. After every $|X|$ of such rounds $\min_{x \in X} N_t(x)$ is incremented at least by one. Hence,

$$\min_{x \in X} N_t(x) \geq \lceil \lceil 3s_t/4 \rceil / |X| \rceil.$$

As we choose $\epsilon < 1/|X|$, we have that $\min_{x \in X} N_t(x) \geq \lceil \epsilon \lceil 3s_t/4 \rceil \rceil$. This is in contradiction with $\min_{x \in X} N_t(x) < \epsilon \lceil s_t/4 \rceil$ whenever $s_t > 2/\epsilon$. \square

F.2 Proof of Lemma 9

We first show that when we are in the exploration phase, there exists an empirically sub-optimal arm $x' \in \tilde{X}(P_t)$ such that $N_t(x') \leq (1 + \epsilon)\eta_t(x') \log(t) \leq (1 + \epsilon)\|\eta_t\|_\infty \log(T)$ where the target exploration rate $\eta_t = \text{SU}(P_t, \eta'_t; \mu_t, \epsilon)$ is determined with the help of the shallow update Algorithm 2. Here, η'_t is the reference logarithmic rate in round t . Assume that this was not the case, then $N_t(x) \geq (1 + \epsilon)\eta_t(x) \log(t)$ for all $x \in X$. All empirically deceitful arms $x \in \tilde{X}_d(P_t)$ would

$$\begin{aligned} \overline{\text{DUAL}}(N_t/\log(t), x, P_t; \mu_t(x)) &\geq \overline{\text{DUAL}}((1 + \epsilon)\eta_t, x, P_t; \mu_t(x)) \\ &= (1 + \epsilon) \cdot \overline{\text{DUAL}}(\eta_t, x, P_t; \mu_t(x)) \geq 1 + \epsilon \end{aligned}$$

pass the sufficient information test; a contradiction. Note that the first inequality follows from the fact that the dual-test function $\overline{\text{DUAL}}(\eta, x, P_t; \mu_t(x'))$ is an increasing function of the target rates η ; see Lemma 19, stated below. The subsequent equality follows from

the fact that the dual-test function $\overline{\text{DUAL}}(\eta, x, P_t; \mu_t(x'))$ is positively homogeneous in the target rates η ; see again Lemma 19. The final inequality holds as $\eta_t = \text{SU}(P_t, \eta'_t; \mu_t, \epsilon)$ and consequently $\overline{\text{DUAL}}(\eta_t, x, P_t; \mu_t(x)) \geq 1$ for all empirically deceitful arms $x \in \tilde{X}_d(P_t)$; see the shallow update Algorithm 2.

So far we established that when we are in the exploration phase, there exists an empirically suboptimal arm $x' \in \tilde{X}(P_t)$ such that $N_t(x') \leq (1 + \epsilon)\eta_t(x') \log(t) \leq (1 + \epsilon)\|\eta_t\|_\infty \log(T)$. Next, we use this to show the desired result. Recall that whenever we are in the exploration round, arm x_t we play is either \underline{x}_t or \bar{x}_t . If $x_t = \underline{x}_t = \arg \min_{x \in X} N_t(x)$, then our earlier observation implies that $N_t(x_t) \leq N_t(x') \leq (1 + \epsilon)\|\eta_t\|_\infty \log(T)$. We now assume that $x_t = \bar{x}_t$, where $\bar{x}_t = \arg \min_{x \in \tilde{X}(P_t)} N_t(x)/\eta_t(x)$. Again, by our earlier observation, there must exist at least one empirically suboptimal arm x' such that $N_t(x')/\eta_t(x') \leq (1 + \epsilon) \log(t)$. This implies that $N_t(x_t) \leq (1 + \epsilon)\eta_t(x_t) \log(t) \leq (1 + \epsilon)\|\eta_t\|_\infty \log(T)$ as well. Thus, in both cases, we have the inequality

$$N_t(x_t) \leq (1 + \epsilon)\|\eta_t\|_\infty \log(T),$$

as desired. □

Lemma 19 (Dual-test Function). *The dual-test function $\overline{\text{DUAL}}(\eta, x, P; \mu)$ is non-decreasing and positively homogeneous in the exploration rate η . That is, we have*

$$\begin{aligned} \overline{\text{DUAL}}(\eta_1, x', P; \mu) &\leq \overline{\text{DUAL}}(\eta_2, x', P; \mu) && \forall 0 \leq \eta_1 \leq \eta_2, \\ \overline{\text{DUAL}}(a \cdot \eta, x', P; \mu) &= a \cdot \overline{\text{DUAL}}(\eta, x', P; \mu) && \forall 0 \leq a, \end{aligned}$$

for any dual feasible $\mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*$.

Proof of Lemma 19. The dual function $\text{DUAL}(\eta, x', P; \mu)$ in η as defined in Equation (13) is a non-decreasing function of η for all arms $x' \in X$, reward distributions $P \in \mathcal{P}_\Omega$, and dual variables $\mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*$; see Lemma 20. Evidently, it follows that $\max_{\rho \geq 0} \text{DUAL}(\eta, x', P; \rho \cdot \mu)$ is a non-decreasing in η for all $x' \in X$ and $P \in \mathcal{P}_\Omega$ too. Positive homogeneity follows because

$$\begin{aligned} \overline{\text{DUAL}}(a \cdot \eta, x', P; \mu) &= \max_{\rho \geq 0} \text{DUAL}(a \cdot \eta, x', P; \rho \cdot \mu) = \max_{\rho \geq 0} \text{DUAL}(a \cdot \eta, x', P; \rho \cdot a \cdot \mu) \\ &= \max_{\rho \geq 0} a \cdot \text{DUAL}(\eta, x', P; \rho \cdot \mu) = a \cdot \overline{\text{DUAL}}(\eta, x', P; \mu), \end{aligned}$$

where the third equality again follows from Lemma 20. □

Lemma 20 (Dual Function). *For any $x' \in \tilde{X}$, the dual function $\text{DUAL}(\eta, x', P; \mu)$ is nondecreasing in the exploration rate η and positively homogeneous in (η, μ) . We have indeed*

for any dual feasible $\mu \in \mathbb{R}_+ \times \mathbb{R} \times \mathcal{K}^*$ that

$$\begin{aligned} \text{DUAL}(\eta_1, x', P ; \mu) &\leq \text{DUAL}(\eta_2, x', P ; \mu) && \forall 0 \leq \eta_1 \leq \eta_2, \\ \text{DUAL}(a \cdot \eta, x', P ; a \cdot \mu) &= a \cdot \text{DUAL}(a \cdot \eta, x', P ; a \cdot \mu) && \forall 0 \leq a. \end{aligned}$$

Proof of Lemma 20. Both properties can be proved simply by verification on the explicit representation

$$\begin{aligned} \text{DUAL}(\eta, x', Q ; \mu) &= \sum_{x \in \tilde{X}(Q), r \in \mathfrak{R}} \eta(x) \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbf{1}(x = x')}{\eta(x)} \right) Q(r, x) \\ &\quad - \sum_{x \in x^*(Q), r \in \mathfrak{R}} \lambda(r, x) Q(r, x) \\ &\quad + \alpha \text{REW}^*(Q) + \beta |\tilde{X}(Q)| + \chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbf{1}(x = x')). \end{aligned}$$

The term $\chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbf{1}(x = x'))$ is clearly nondecreasing in η . The nonlinear terms $\eta(x) \log((\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbf{1}(x = x'))/\eta(x)) Q(r, x)$ can easily be observed to be nondecreasing in η by noting that the functions of the form $u \log((u - v)/u)$ are increasing in $u \in \mathbb{R}$ for any given $v \in \mathbb{R}$. This proves the first part of the lemma. The second part is trivial and follows from

$$\begin{aligned} &\text{DUAL}(a\eta, x', Q ; a\mu) \\ &= \sum_{x \in \tilde{X}(Q), r \in \mathfrak{R}} a\eta(x) \log \left(\frac{a\eta(x) - a\lambda(r, x) - a\beta - a\alpha \cdot r \mathbf{1}(x = x')}{a\eta(x)} \right) Q(r, x) - \sum_{x \in x^*(Q), r \in \mathfrak{R}} a\lambda(r, x) Q(r, x) \\ &\quad + a\alpha \text{REW}^*(Q) + a\beta |\tilde{X}(Q)| + \chi_{-\infty}(a\eta(x) \geq a\lambda(r, x) + a\beta + a\alpha \cdot r \mathbf{1}(x = x')) \\ &= \sum_{x \in \tilde{X}(Q), r \in \mathfrak{R}} a\eta(x) \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbf{1}(x = x')}{\eta(x)} \right) Q(r, x) - \sum_{x \in x^*(Q), r \in \mathfrak{R}} a\lambda(r, x) Q(r, x) \\ &\quad + a\alpha \text{REW}^*(Q) + a\beta |\tilde{X}(Q)| + a\chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbf{1}(x = x')) \\ &= a \cdot \text{DUAL}(\eta, x', Q ; \mu). \end{aligned}$$

□

G Continuity Results

We suggest the readers who are not familiar with Slater's Constraint Qualification Condition to review Section I.2 before reading the proofs presented in this section.

G.1 Proof Proposition 2

The proof has two parts. In the first part, we show that the distance function $\text{DIST}(\eta, x', P)$, defined in Equation (6), is finite if and only if x' is deceitful, i.e., $\text{REW}^*(P) < \text{REW}_{\max}(x', P)$. In the second part, we show that problem (6) and problem (9) are identical when arm x' is deceitful.

First part. By definition of the distance function DIST in Equation (6), a necessary requirement for the distance function $\text{DIST}(\eta, x', P)$ to be finite is that $\text{REW}^*(P) < \text{REW}_{\max}(x', P)$. To see why, note that the set of deceitful distributions $\text{DECEIT}(x', P)$ is empty and hence $\text{DIST}(\eta, x', P)$ infinite if $\text{REW}^*(P) \geq \text{REW}_{\max}(x', P)$. We will show that this condition, i.e., $\text{REW}^*(P) < \text{REW}_{\max}(x', P)$, is in fact sufficient too. Therefore, we have that the effective domain of the information distance function is given by $\{P \in \mathcal{P} : \text{DIST}(\eta, x', P) < \infty\} = \{P : \text{REW}^*(P) < \text{REW}_{\max}(x', P)\}$, as desired.

When $\text{REW}^*(P) < \text{REW}_{\max}(x', P)$, we will argue that the distance function is finite, i.e.,

$$\text{DIST}(\eta, x', P) \leq \sum_{x \in \bar{X}(P)} \eta(x) I(P(x), \bar{Q}(x)) < \infty$$

by constructing a distribution $\bar{Q} \in \text{DECEIT}(x', P)$ which satisfies $\bar{Q} \ll P$; see Lemma 11. By definition of $\text{REW}_{\max}(x', P)$ stated in Equation (7), there exists a maximizing distribution $\hat{Q} \in \mathcal{P}$ such that $\hat{Q}(x^*(P)) = P(x^*(P))$ and

$$\text{REW}^*(P) < \sum_{r \in \mathfrak{R}} r \hat{Q}(r, x')$$

which does ensure that $\hat{Q} \in \text{DECEIT}(x', P)$. Unfortunately, we are not ensured that $\hat{Q} \ll P$. However, we may simply perturb \hat{Q} as $\bar{Q} = (1 - \theta)\hat{Q} + \theta P$ for some $\theta \in (0, 1)$ and observe that from convexity we have $\bar{Q} \in \mathcal{P}$ and $\bar{Q}(x^*(P)) = P(x^*(P))$. Furthermore, for sufficiently small perturbations θ we have from continuity of the expected reward of arm x' in the reward distribution that $\text{REW}^*(P) < \sum_{r \in \mathfrak{R}} r \bar{Q}(r, x')$ as well. Thus, $\bar{Q} \in \text{DECEIT}(x', P)$ and $\bar{Q} \ll P$.

Second part. Here, we show that problem (6) and problem (9) are identical when arm x' is deceitful. It is easy to see that the optimal value of problem (6), i.e., $\text{DIST}(\eta, x', P)$, is greater than or equal to that of problem (9) as the latter is characterized as a minimization problem of the same objective function over a larger feasible set. Next, we show that the optimal value of problem (6) is less than or equal to that of problem (9). Let Q^* be the optimal solution to problem (9). Note that an optimal solution of this problem exists because its objective function is lower semicontinuous per Lemma 11 and its feasible set is a compact set; see for instance Bertsekas (2009, Prop. 3.2.1). Now, consider the perturbed reward distributions

$Q'_\theta = \theta \cdot \bar{Q} + (1 - \theta) \cdot Q^* \in \mathcal{P}$ parameterized in $\theta \in [0, 1]$. Again we have that $Q'_\theta \in \mathcal{P}$ as the set \mathcal{P} is convex and contains both reward distribution Q^* and \bar{Q} . Furthermore, we have that $\sum_{r \in \mathfrak{R}} r Q'_\theta(r, x') > \text{REW}^*(P)$ for all $\theta \in (0, 1]$ as we have both $\sum_{r \in \mathfrak{R}} r Q^*(r, x') \geq \text{REW}^*(P)$ and $\sum_{r \in \mathfrak{R}} r \bar{Q}(r, x') > \text{REW}^*(P)$. Hence, the reward distributions Q'_θ are feasible in the minimization problem (6) and consequently for any $\theta \in (0, 1]$ we have

$$\begin{aligned} (6) &\leq \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q'_\theta(x)) \\ &\leq \theta \cdot \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), \bar{Q}(x)) + (1 - \theta) \cdot \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q^*(x)). \end{aligned}$$

Considering the fact that the term $\sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), \bar{Q}(x))$ is finite, and by letting θ go to zero, we have

$$\begin{aligned} &\lim_{\theta \downarrow 0} \theta \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), \bar{Q}(x)) + (1 - \theta) \sum_{x \in \tilde{X}(P)} \eta(x) I(P(x), Q^*(x)) \\ &= \sum_{x \in \tilde{X}(P)} \eta(x) \cdot I(P(x), Q^*(x)). \end{aligned}$$

The proof is then completed by observing that the left hand side is greater than or equal to the optimal value of problem (6), and the right hand side is equal to that of problem (9). \square

G.2 Proof of Proposition 3

Here, we first show that the lower regret bound function $C(Q)$ is continuous at any reward distribution $P \in \mathcal{P}'$. We show that the lower regret bound function $C(Q)$ is continuous with the help of Lemmas 7 and Lemma 22, which are found at the end of this section and two classical continuity results (Propositions 9 and 10) regarding the minimizers of parametric optimization problems restated in Section I.1 for completeness.

From Lemma 7, we can find a neighborhood $\mathcal{N} \subseteq \mathcal{P}'$ around any $P \in \mathcal{P}'$ such that

$$Q \in \mathcal{N} \implies x^*(Q) = x^*(P) := x_0^*, \tilde{X}(Q) = \tilde{X}(P) := \tilde{X}_0 \text{ and } \tilde{X}_d(Q) = \tilde{X}_d(P) := \tilde{X}_d^0.$$

That is, all reward distributions Q in the neighborhood \mathcal{N} share their optimal, nondeceitful and deceitful arms with the reward distribution P . From Proposition 2, for any $Q \in \mathcal{N}$, the regret lower bound function is given as

$$\begin{aligned} C(Q) &:= \inf_{\eta \geq 0} \sum_{x \in \tilde{X}_0} \eta(x) \Delta(x, Q) \\ \text{s.t. } &1 \leq \text{DIST}(\eta, x, Q) \quad \forall x \in \tilde{X}_d^0 \\ &\eta(x) = 0 \quad \forall x \in x_0^* \end{aligned} \tag{50}$$

where the value of $\eta(x_0^*)$ for the optimal arm is in fact arbitrary and can be set to zero. We would like to establish that the lower regret bound function $C(Q)$ is continuous at P by applying Proposition 10 to the minimization problem (50). To apply this proposition, we need to show that its feasible set mapping is a lower semi-continuous mapping. One can hope to establish lower semi-continuity by verifying the conditions of Proposition 9. One of these conditions demands the inequality constraint functions $\text{DIST}(\eta, x, Q)$ to be upper semi-continuous at points (η_0, P) , which may not be the case when η_0 is not strictly positive. To circumvent this challenge, we introduce the following family of perturbed regret lower bound problem

$$\begin{aligned}
C_\delta(Q) &:= \inf_{\eta} \sum_{x \in \tilde{X}_0} \eta(x) \Delta(x, Q) \\
\text{s.t.} \quad & 1 \leq \text{DIST}(\eta, x, Q) \quad \forall x \in \tilde{X}_d^0, \\
& \eta(x) \geq \delta \quad \forall x \in X, \\
& \eta(x) = \delta \quad \forall x \in x_0^*
\end{aligned} \tag{51}$$

and show that continuity of $C_\delta(Q)$ at $Q = P$ implies continuity of the regret lower bound at P . We then proceed to show $C_\delta(Q)$ is indeed continuous at P .

Continuity of the perturbed regret lower bound implies that of regret lower bound. By definition of the perturbed regret lower bound, we have the following inequalities

$$C_\delta(Q) - \delta |\tilde{X}_0| \leq C_\delta(Q) - \delta \sum_{x \in \tilde{X}_0} \Delta(x, Q) \leq C(Q) \leq C_\delta(Q)$$

for any $\delta > 0$ uniformly in $Q \in \mathcal{N}$. The first inequality holds because $\Delta(x, Q) \in [0, 1]$ for any arm x and reward distribution Q . The second inequality follows from the fact that we have

$$\begin{aligned}
C_\delta(Q) &= \inf_{\eta} \sum_{x \in \tilde{X}_0} (\eta(x) - \delta) \Delta(x, Q) + \delta \sum_{x \in \tilde{X}_0} \Delta(x, Q) \\
\text{s.t.} \quad & 1 \leq \text{DIST}(\eta, x, Q) \quad \forall x \in \tilde{X}_d^0 \\
& \eta(x) - \delta \geq 0 \quad \forall x \in X, \\
& \eta(x) - \delta = 0 \quad \forall x \in x_0^* \\
& \leq \inf_{\eta} \sum_{x \in \tilde{X}_0} (\eta(x) - \delta) \Delta(x, Q) + \delta \sum_{x \in \tilde{X}_0} \Delta(x, Q) \\
\text{s.t.} \quad & 1 \leq \text{DIST}(\eta - \delta \cdot \mathbf{1}, x, Q) \quad \forall x \in \tilde{X}_d^0 \\
& \eta(x) - \delta \geq 0 \quad \forall x \in X, \\
& \eta(x) - \delta = 0 \quad \forall x \in x_0^* \\
& = C(Q) + \delta \sum_{x \in \tilde{X}_0} \Delta(x, Q).
\end{aligned} \tag{52}$$

Hence, we have that $C(Q) = \lim_{\delta \downarrow 0} C_\delta(Q)$ uniformly in $Q \in \mathcal{N}$. It remains to be shown that the perturbed regret lower bound functions $C_\delta(Q)$ are continuous. The claimed result then indeed follows as uniform limits of continuous functions are themselves continuous.

The perturbed regret lower bound is continuous. We show the perturbed regret lower functions are continuous by applying Proposition 10 to the minimization problem (51). Let $M_\delta(Q)$ be the feasible set of the perturbed minimization problem (51) where we use as suggested in Section I the decomposition

$$\begin{aligned} M_\delta(Q) &:= \Gamma(Q) \cap \{\eta : g_i(\eta, Q) \leq 0\} \\ &= \{\eta : \eta(x) \geq \delta \ \forall x \in X, \eta(x) = \delta \ \forall x \in x_0^*\} \cap \{\eta : 1 \leq \text{DIST}(\eta, x, Q) \ \forall x \in \tilde{X}_d^0\}. \end{aligned}$$

We will invoke Proposition 10 with parameter $\theta = Q$, decision variable $y = \eta$, $\Gamma(Q) = \{\eta : \eta(x) \geq \delta \ \forall x \in X, \eta(x) = \delta \ \forall x \in x_0^*\}$, and $\{\eta : g_i(\eta, Q) \leq 0\} = \{\eta : 1 - \text{DIST}(\eta, x, Q) \leq 0 \ \forall x \in \tilde{X}_d^0\}$. To invoke Proposition 10, and establish our claim, it simply remains to verify that all conditions of this proposition are met:

1. **Feasible region $M_\delta(Q)$ is a lower semi-continuous mapping at P .** We establish this result by invoking Proposition 9.

(a) *The Slater's constraint qualification condition is satisfied ($M_{0,\delta}(P) \neq \emptyset$).* This condition, which is presented in Definition 3, holds if we can show that the set of Slater points

$$M_{0,\delta}(P) := \{\eta : \eta(x) \geq \delta \ \forall x \in X, \eta(x_0^*) = \delta\} \cap \{\eta : 1 < \text{DIST}(\eta, x, P) \ \forall x \in \tilde{X}_d^0\}$$

is non-empty. To show this, we will first point out that $\text{DIST}(\hat{\eta}, x, P) > 0$ for any $\hat{\eta} \geq \delta > 0$ and any deceitful arm $x \in \tilde{X}_d^0$. We can then use this result to construct a Slater point in $M_{0,\delta}(P)$.

By Proposition 2, for any deceitful arm $x' \in \tilde{X}_d^0$, we can characterize the distance function as the minimum

$$\begin{aligned} \text{DIST}(\hat{\eta}, x', P) &= \min_{Q \in \mathcal{P}} \sum_{x \in \tilde{X}_0} \hat{\eta}(x) I(P(x), Q(x)) \\ &\text{s.t. } Q(x_0^*) = P(x_0^*), \\ &\quad \sum_{r \in \mathfrak{A}} r Q(r, x') \geq \text{REW}^*(P). \end{aligned}$$

By the classical extreme value theorem, the minimum of the above optimization problem must be achieved at some Q^* as its feasible set is compact and its objective function is lower semi-continuous due to Lemma 11. That is, $\text{DIST}(\hat{\eta}, x', P) = \sum_{x \in \tilde{X}_0} \hat{\eta}(x) I(P(x), Q^*(x))$ where $\eta(x) \geq 0$ for all $x \in \tilde{X}_0$. Because there is only one unique optimal arm x_0^* , which is

trivially not in \tilde{X}_d^0 , it follows that $Q^* \neq P$ and hence $\text{DIST}(\hat{\eta}, x', P) > 0$ due to the first result of Lemma 11. This allows us to define the following Slater (feasible) points for the perturbed problem (51) at $Q = P$:

$$\bar{\eta}(x) = \delta + 2\hat{\eta}(x) / \min_{x' \in \tilde{X}_d^0} \text{DIST}(\hat{\eta}, x', P) \quad \forall x \in \tilde{X}^0 \quad \text{and} \quad \bar{\eta}(x_0^*) = \delta.$$

To see why these are Slater points, first of all note that $\bar{\eta} \geq \delta$ and second of all, for any $x \in \tilde{X}_d^0$, we have

$$\begin{aligned} \text{DIST}(\bar{\eta}, x, P) &= \text{DIST}(\delta \mathbf{1} + 2\hat{\eta} / \min_{x' \in \tilde{X}_d^0} \text{DIST}(\hat{\eta}, x', P), x, Q) \\ &\geq \text{DIST}(2\hat{\eta} / \min_{x' \in \tilde{X}_d^0} \text{DIST}(\hat{\eta}, x', P), x, P) \\ &= \frac{2\text{DIST}(\hat{\eta}, x, P)}{\min_{x' \in \tilde{X}_d^0} \text{DIST}(\hat{\eta}, x', P)} \geq 2, \end{aligned}$$

where the first equality follows from the fact that the value of $\eta(x_0^*)$ is immaterial and the first inequality follows from Lemma 21 where we show $0 < \text{DIST}(\alpha\eta, x, Q) = \alpha\text{DIST}(\eta, x, Q)$ for $\alpha > 0$. Thus, $\bar{\eta} \in M_{0,\delta}(P)$.

(b) *The distance function $\text{DIST}(\eta, x, Q)$ is upper semi-continuous on $\{M_{0,\delta}(P)\} \times \{P\}$.* First remark that $M_{0,\delta}(P) \subseteq \{\eta : \eta \geq \delta \cdot \mathbf{1}\}$. In Lemma 22, presented at the end of this section, we show that $\text{DIST}(\eta, x, Q)$ is in fact continuous on $\{\eta : \eta \geq \delta \cdot \mathbf{1}\} \times \{P\}$ for any deceitful arm $x \in \tilde{X}_d^0$ and $P \in \text{int}(\mathcal{P})$.

(c) *The distance function $\eta \mapsto \text{DIST}(\eta, x, P)$ is concave on $\Gamma(P)$.* First remark that $\Gamma(P) \subseteq \{\eta : \eta \geq 0\}$. From its definition, it is also immediately clear (c.f., Boyd and Vandenberghe (2004, Section 3.2.5)) that the distance function $\text{DIST}(\eta, x, P)$ is concave on $\{\eta : \eta \geq 0\}$ for any deceitful arm $x \in \tilde{X}_d^0$.

2. Feasible region $M_\delta(Q)$ is convex for all $Q \in \mathcal{P}$ and the mapping $M_\delta(Q)$ is closed at P . The feasible region $M_\delta(P)$ is evidently convex as the distance functions $\text{DIST}(\eta, x, Q)$ are concave in η . The mapping $M_\delta(Q)$ is closed at P if its graph $\{(\eta, Q) : Q \in \mathcal{N}', \eta \in M_\delta(Q)\}$ is a closed set for some closed neighborhood \mathcal{N}' containing P and itself sufficiently small to be a subset of \mathcal{P}' . Closedness of this graph follows immediately from Lemma 22 in which we indicate that the distance function $\text{DIST}(\eta, x, Q)$ is continuous at any point $\eta_0 > 0$ and $Q \in \mathcal{P}'$ for any deceitful arm $x \in \tilde{X}_d^0$.

3. The set of minimizers of problem (51) at P is nonempty and compact. The set of minimizers in the optimization problem (51) for P is nonempty and compact as its associated objective function is continuous and its sublevel sets restricted to the feasible

domain are bounded, c.f., Bertsekas (2009, Prop. 3.2.1).

4. The objective function of problem (51) is continuous and convex. The objective function of the perturbed optimization problem (51), i.e., $\sum_{x \in \tilde{X}_0} \eta(x) \Delta(x, Q)$, is continuous at any point (η, Q) . This is so because the term $\Delta(x, Q) := \sum_{r \in \mathfrak{R}} rQ(r, x) - \text{REW}^*(Q)$ is continuous at any Q . Observe that the first term in $\Delta(x, Q)$, i.e., $\sum_{r \in \mathfrak{R}} rQ(r, x)$ is linear in Q , and hence continuous. The second term, i.e., $\text{REW}^*(Q) = \max_x \sum_{r \in \mathfrak{R}} Q(r, x)$, is also continuous as it is the maximum of continuous (linear) functions. The product $\sum_{x \in \tilde{X}_0} \eta(x) \Delta(x, Q)$ will consequently be continuous at any point (η, Q) . Evidently, the objective function is linear and hence convex in η .

The proof of the existence of an $0 < \epsilon$ -suboptimal rate selection $\eta_\epsilon^c(Q)$ is referred to Proposition 4. \square

Lemma 21 (Distance Function). *For any $x' \in \tilde{X}$, the distance function $\text{DIST}(\eta, x', P)$ is non-decreasing and positively homogeneous in the exploration rate η . That is, we have*

$$\begin{aligned} \text{DIST}(\eta_1, x', P) &\leq \text{DIST}(\eta_2, x', P) && \forall 0 \leq \eta_1 \leq \eta_2, \\ \text{DIST}(a \cdot \eta, x', P) &= a \cdot \text{DIST}(\eta, x', P) && \forall 0 \leq a. \end{aligned}$$

Proof of Lemma 21. Both statements are an immediate consequence of the fact that the objective function in the information minimization problem (6), defining our distance function $\text{DIST}(\eta, x', P)$, is nondecreasing and positively homogeneous in the logarithmic rate η for any reward distribution Q . \square

Lemma 22 (Continuity of the Distance Function). *Assume that the reward distribution $P \in \text{int}(\mathcal{P})$ has a unique optimal arm $x^*(P)$ and consider any of its deceitful arms $x' \in \tilde{X}_d(P)$. The distance function $\text{DIST}(\eta, x', Q)$ is continuous at any point (η_0, P) with $\eta_0 > 0$.*

G.2.1 Proof of Lemma 22

Recall that an arm $x' \in \tilde{X}_d(Q)$ is denoted as deceitful for some reward distribution Q if we have that $\text{REW}_{\max}(x', Q) > \text{REW}^*(Q)$. From Lemma 7, we know there exists a neighborhood $\mathcal{N} \subseteq \mathcal{P}$ around P so that

$$Q \in \mathcal{N} \implies x^*(Q) = x^*(P) := x_0^*, \tilde{X}(Q) = \tilde{X}(P) := \tilde{X}_0 \text{ and } \tilde{X}_d(P) = \tilde{X}_d^0 \subseteq \tilde{X}_d(Q).$$

Note that in general we may observe $\tilde{X}_d(P) \subset \tilde{X}_d(Q)$ as a nondeceitful arm $x \in \tilde{X}_n(P)$ for which $\text{REW}^*(x, P) = \text{REW}^*(P)$ may become deceitful $x \in \tilde{X}_d(Q)$ with $Q \in \mathcal{N}$. For each

reward distribution $Q \in \mathcal{N}$ in a neighborhood around P , arm $x^*(P)$ remains optimal, arms in set $\tilde{X}(P)$ remain suboptimal, while any deceitful arm in $\tilde{X}_d(P)$ remains deceitful. Let $\eta_0 > 0$ be any arbitrary positive logarithmic exploration rate. From Proposition 2, it follows that for any $Q \in \mathcal{N}$, we have that the distance function for any deceitful arm $x' \in \tilde{X}_d^0$ simplifies to

$$\begin{aligned} \text{DIST}(\eta, x', Q) &= \min_{Q' \in \mathcal{P}} \sum_{x \in \tilde{X}_0} \eta(x) I(Q(x), Q'(x)) \\ \text{s.t. } &\sum_{r \in \mathfrak{R}} r Q'(r, x') \geq \sum_{r \in \mathfrak{R}} r Q(r, x_0^*), \\ &Q'(x_0^*) = Q(x_0^*). \end{aligned} \quad (53)$$

We would like to show that function $\text{DIST}(\eta, x', Q)$ is continuous at (η_0, P) by applying Proposition 10 to its minimization characterization stated in Equation (53). To do so, we will verify all the conditions in this proposition. Before that, let us decompose as suggested in Section I the feasible set mapping of problem (53) as follows:

$$\begin{aligned} M(\eta, Q) &=: \Gamma(\eta, Q) \cap \{Q' : g(Q', (\eta, Q)) \leq 0\} \\ &= \{Q' \in \mathcal{P} : Q'(x_0^*) = Q(x_0^*)\} \cap \{Q' : \sum_{r \in \mathfrak{R}} r Q'(r, x') \geq \sum_{r \in \mathfrak{R}} r Q(r, x_0^*)\}. \end{aligned}$$

We now proceed to verifying all the conditions of Proposition 10.

1. **Feasible region $M(\eta, Q)$ is a lower semi-continuous mapping at (η_0, P) .** We show this result by invoking Proposition 9. We need to verify the following conditions to invoke Proposition 9.

(a) *The mapping $\Gamma(\eta, Q)$ is lower semi-continuous at (η_0, P) .* Clearly, the mapping $\Gamma(\eta, Q) := \{Q' \in \mathcal{P} : Q'(x_0^*) = Q(x_0^*)\}$ is independent of η_0 . Hence, this claim is already established in Lemma 18.

(b) *Set $\Gamma(\eta_0, P)$ is convex and non-empty.* Solutions sets of linear systems are well known to be affine and hence convex. Evidently, we have $P \in \Gamma(\eta_0, P)$ and hence $\Gamma(\eta_0, P)$ must be non-empty.

(c) *The Slater's constraint qualification condition is satisfied ($M_0(\eta_0, P) \neq \emptyset$).* Consider a deceitful arm $x' \in \tilde{X}_d(P)$ and let

$$\begin{aligned} Q^* &\in \arg \max \sum_{r \in \mathfrak{R}} r Q(r, x') \\ \text{s.t. } &Q \in \mathcal{P}, \\ &Q(x_0^*) = P(x_0^*). \end{aligned}$$

Note that the maximum of the above optimization problem is attained at some distribution Q^* because its objective function is continuous and its feasible region is a compact set.

Because arm x' is deceitful, i.e., $\sum_{r \in \mathfrak{R}} r Q^*(r, x') > \sum_{r \in \mathfrak{R}} r P(r, x_0^*)$, it follows immediately that $Q^* \in M_0(\eta_0, P) = \{Q' \in \mathcal{P} : Q'(x_0^*) = Q(x_0^*)\} \cap \{Q' : \sum_{r \in \mathfrak{R}} r Q'(r, x') > \sum_{r \in \mathfrak{R}} r Q(r, x_0^*)\}$, which is the desired result.

2. Feasible region $M(\eta, Q)$ is convex for all $Q \in \mathcal{P}$ and $M(\eta, Q)$ is a closed mapping at (η_0, P) . The set \mathcal{P} is assumed to be convex and closed. The $\{Q' : Q'(x_0^*) = Q(x_0^*), \sum_{r \in \mathfrak{R}} r Q'(r, x') \geq \sum_{r \in \mathfrak{R}} r Q(r, x_0^*)\}$ is a closed and convex polyhedral set. Consequently, the intersection $M(\eta, P) = \mathcal{P} \cap \{Q' : Q'(x_0^*) = P(x_0^*), \sum_{r \in \mathfrak{R}} r Q'(r, x') \geq \sum_{r \in \mathfrak{R}} r P(r, x_0^*)\}$ is a closed convex set for any (η, Q) . The mapping $M(\eta, Q)$ is closed at (η_0, P) as its graph

$$\{(Q', \eta, Q) : Q' \in \mathcal{P}, Q'(x_0^*) = Q(x_0^*), \sum_{r \in \mathfrak{R}} r Q'(r, x') \geq \sum_{r \in \mathfrak{R}} r Q(r, x_0^*), (\eta, Q) \in \mathcal{N}'\}$$

is closed for any closed neighborhood \mathcal{N}' containing (η_0, P) .

3. The set of minimizers of problem (53) at (η_0, P) is nonempty and compact. The set of minimizers of problem (53) is nonempty and compact as its feasible set $P \in M(\eta_0, P)$ is compact (recall that \mathcal{P} is a compact set) and its objective function is lower semi-continuous because of the last result of Lemma 11, stated in Section B; see also Bertsekas (2009, Proposition 3.2.1).

4. The objective function of problem (53) is continuous at (η_0, P) and convex in Q' . With a slight abuse of notation, let Q^* be an arbitrary minimizer to problem (53). As the target rate satisfies $\eta_0 > 0$, we will now argue that we must have that $P \ll Q^* \in \text{int}(\mathcal{P})$. Indeed, otherwise $\text{DIST}(\eta_0, x', P) = \sum_{x \in \tilde{X}_0} \eta_0(x) I(P(x), Q^*(x)) = +\infty$ which would contradict Proposition 2 in which we show that for any deceitful arm $x' \in \tilde{X}_d(P)$ the distance function $\text{DIST}(\eta_0, x', P)$ is always finite. Hence, the objective function is continuous at the point $(\eta_0, P, Q^*) \in \mathbb{R}_{++} \times \text{int}(\mathcal{P}) \times \text{int}(\mathcal{P})$. Finally, the objective function of problem (53) is also convex in Q' , per Lemma 11. \square

G.3 Proof of Proposition 4

We now proceed to show that the selection functions $(\eta'_\epsilon(Q), \mu_\epsilon(Q))$, defined in Equation (19), are continuous at any reward distribution $P \in \mathcal{P}'$ satisfying Assumption 1. We will prove that the non-empty, closed and convex mapping representing the feasible set $M(Q)$ of problem (19) is lower semi-continuous at P . By Michael's theorem, c.f., Aubin and Frankowska (2009, Corollary 9.1.2), there exists a selection $(\eta_\epsilon^c(Q), \mu_\epsilon^c(Q))$ which is continuous on $\mathcal{N}' \subseteq \mathcal{P}'$, a sufficiently small closed neighborhood of P . The existence of such

neighborhood is guaranteed by Lemma 7. Remark that by our minimal selection definition, we must have for all $Q \in \mathcal{N}'$ that

$$\sum_{x \in X} \eta'_\epsilon(x, Q)^2 + \sum_{x \in X} \|\mu_\epsilon(x, Q)\|^2 \leq \max_{Q \in \mathcal{N}'} \sum_{x \in X} \eta_\epsilon^c(x, Q)^2 + \sum_{x \in X} \|\mu_\epsilon^c(x, Q)\|^2 < \infty$$

where the last inequality follows from the fact that continuous functions on compact sets attain a finite maximum. The minimum norm selection $(\eta'_\epsilon(Q), \mu_\epsilon(Q))$ is unique and continuous at such P following Aubin and Frankowska (2009, Corollary 9.3.2).

From Lemma 7, there exists an open neighborhood $\mathcal{N} \subseteq \mathcal{P}'$ around P satisfying Assumption 1 so that

$$Q \in \mathcal{N} \implies x^*(Q) = x^*(P) := x_0^*, \tilde{X}(Q) = \tilde{X}(P) := \tilde{X}_0 \text{ and } \tilde{X}_d(Q) = \tilde{X}_d(P) := \tilde{X}_d^0.$$

That is, for each reward distribution $Q \in \mathcal{N}$ in a neighborhood around P , the arm $x^*(P)$ remains optimal, the arms $\tilde{X}(P)$ remain suboptimal, and any (non)deceitful arm remains (non)deceitful. Hence, the feasible set for $Q \in \mathcal{N}$ of the minimization problem (19) is captured by the following mapping:

$$M(Q) = \left\{ (\eta, \mu) : \begin{array}{l} \eta(x) \in \mathbb{R}_+, \mu(x) \in \mathcal{K}^* \times \mathbb{R}_+ \times \mathbb{R} \quad \forall x \in X, \\ 1 \leq \text{DUAL}(\eta, x, Q; \mu(x)) \quad \forall x \in \tilde{X}_d^0, \\ \sum_{x \in \tilde{X}_0} \eta(x) \Delta(x, Q) \leq C(Q) + \epsilon, \\ \eta(x) \geq \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, Q)) \quad \forall x \in \tilde{X}_0, \\ \eta(x) \geq \lambda(r, x, x') + \beta(x') + \alpha(x') \cdot r \mathbb{1}(x = x') + \frac{\epsilon}{2 \sum_{x \in \tilde{X}_0} \Delta(x, Q)} \\ \forall x \in \tilde{X}_d^0, x \in \tilde{X}_0 \end{array} \right\}.$$

Mapping M is lower-semicontinuous at P . We will establish that the mapping M is lower semi-continuity by verifying the following conditions of Proposition 9 where we use the decomposition

$$\begin{aligned} M(Q) &=: \Gamma(Q) \cap \{(\eta, \mu) : g_i((\eta, \mu), Q) \leq 0 \quad \forall i \in [1, \dots, m_2]\} \\ &= \{\mu(x) \in \mathcal{K}^* \times \mathbb{R}_+ \times \mathbb{R} \quad \forall x \in X\} \cap \\ &\quad \left\{ (\eta, \mu) : \begin{array}{l} 1 \leq \text{DUAL}(\eta, x, Q; \mu(x)) \quad \forall x \in \tilde{X}_d^0, \\ \sum_{x \in \tilde{X}_0} \eta(x) \Delta(x, Q) \leq C(Q) + \epsilon, \\ \eta(x) \geq \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, Q)) \quad \forall x \in \tilde{X}_0, \eta(x) \geq 0 \quad \forall x \in X, \\ \eta(x) \geq \lambda(r, x, x') + \beta(x') + \alpha(x') \cdot r \mathbb{1}(x = x') + \frac{\epsilon}{2 \sum_{x \in \tilde{X}_0} \Delta(x, Q)} \\ \forall x \in \tilde{X}_d^0, x \in \tilde{X}_0 \end{array} \right\}. \end{aligned}$$

- **The Slater's constraint qualification condition.** To verify this condition, we need to show that the set of Slater points

$$M_0(P) = \left\{ (\eta, \mu) : \begin{array}{l} \eta(x) > 0, \mu(x) \in \mathcal{K}^* \times \mathbb{R}_+ \times \mathbb{R} \quad \forall x \in X, \\ \text{DUAL}(\eta, x, P; \mu(x)) > 1 \quad \forall x \in \tilde{X}_d^0, \\ \sum_{x \in \tilde{X}_0} \eta(x) \Delta(x, P) < C(P) + \epsilon, \\ \eta(x) > \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, P)) \quad \forall x \in \tilde{X}_0, \\ \eta(x) > \lambda(r, x, x') + \beta(x') + \alpha(x') \cdot r \mathbf{1}(x = x') + \frac{\epsilon}{2 \sum_{x \in \tilde{X}_0} \Delta(x, P)} \\ \forall x \in \tilde{X}_d^0, x \in \tilde{X}_0 \end{array} \right\}$$

is nonempty. Consider to that end first the following $\epsilon/4$ -suboptimal logarithmic exploration rates

$$\left\{ (\eta, \mu) : \begin{array}{l} \eta \geq 0, \mu(x) \in \mathcal{K}^* \times \mathbb{R}_+ \times \mathbb{R} \quad \forall x \in \tilde{X}_d^0, \\ \text{DUAL}(\eta, x, P; \mu(x)) \geq 1 \quad \forall x \in \tilde{X}_d^0, \\ \sum_{x \in \tilde{X}_0} \eta(x) \Delta(x, P) \leq C(P) + \epsilon/4, \\ \eta(x) \geq \lambda(r, x, x') + \beta(x') + \alpha(x') \cdot r \mathbf{1}(x = x') \quad \forall x \in \tilde{X}_d^0, x \in \tilde{X}_0 \end{array} \right\}$$

in the optimization problem (15). From the definition of the infimum in problem (15), the previous set is nonempty for all $\epsilon > 0$. Consider an arbitrary element $(\eta_{\epsilon/4}, \mu_{\epsilon/4})$ in the previous set. We can easily verify that the point $(\eta_{\epsilon/4} + \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, P)) \cdot \mathbf{1}, \mu_{\epsilon/4})$ is in $M(P)$, where $\mathbf{1}$ is the all one vector. Indeed, from Lemma 20, we have that $\text{DUAL}(\eta_{\epsilon/4} + \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, P)) \cdot \mathbf{1}, x, P; \mu_{\epsilon/4}) \geq \text{DUAL}(\eta_{\epsilon/4}, x, P; \mu_{\epsilon/4}) \geq 1$ and $\sum_{x \in \tilde{X}_0} (\eta_{\epsilon/4}(x) + \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, P))) \Delta(x, P) = \sum_{x \in \tilde{X}_0} \eta_{\epsilon/4}(x) \Delta(x, P) + \epsilon/2 \leq \epsilon/4 + \epsilon/2 = 3\epsilon/4 < \epsilon$.

We show that a similar construction can be used to construct a Slater point in $M_0(P)$ as well. Consider a point $(\eta'', \mu_0 = (\alpha_0, \beta_0, \lambda_0)) = (1 + \delta) \cdot (\eta_{\epsilon/4}(P) + \epsilon / (2 \sum_{x \in \tilde{X}(P)} \Delta(x, P)) \cdot \mathbf{1}), \mu_{\epsilon/4}) \in M(P)$ for some $\delta > 0$ small enough so that $\delta / (1 - 3\delta) < \epsilon / (4C(P))$. Using again Lemma 20, it is trivial to verify that we have

$$\begin{aligned} \eta'' &\geq (1 + \delta) \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, P)) \cdot \mathbf{1} > \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, P)) \cdot \mathbf{1} > 0, \\ \text{DUAL}(\eta'', x, P; \mu_0(x)) &\geq 1 + \delta > 1 \quad \forall x \in \tilde{X}_d^0, \\ \sum_{x \in \tilde{X}_0} \eta''(x) \Delta(x, P) &\leq (1 + \delta) \cdot (C(P) + 3\epsilon/4) < C(P) + \epsilon, \\ \eta''(x) &> \lambda_0(r, x, x') + \beta_0(x') + \alpha_0(x') \cdot r \mathbf{1}(x = x') + \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, P)) \\ &\forall x' \in \tilde{X}_d^0, x \in \tilde{X}_0. \end{aligned}$$

- **The rest of conditions.** We first verify the required convexity conditions. It is clear that the constraint functions $g_i((\eta, \mu), P)$ which appear in the feasible set mapping M are linear and hence convex jointly in η and μ with the exception of those related to the constraints $\text{DUAL}(\eta, x, P ; \mu(x)) \geq 1$ for all $\forall x \in \tilde{X}_d^0$. As pointed out previously, the dual function $\text{DUAL}(\eta, x, P ; \mu(x))$ is concave jointly in η and μ and consequently all constraint functions are convex. Finally, we verify all required continuity conditions. As $\eta(x) \geq \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, Q)) > 0$ and

$$\eta(x) - \lambda(r, x, x') - \beta(x') - \alpha(x') \cdot r \mathbb{1}(x = x') \geq \epsilon / (2 \sum_{x \in \tilde{X}_0} \Delta(x, Q)) > 0$$

for all $x' \in \tilde{X}_d^0$ and $x \in \tilde{X}_0$, the constraint functions $-\text{DUAL}(\eta, x, Q ; \mu(x))$ for all $\forall x \in \tilde{X}_d^0$ are continuous on $M_0(P) \times \{P\}$; see Proposition 6 below. The constraint functions related to all other constraints characterizing the feasible set mapping M are linear and hence continuous. \square

Proposition 6 (Continuity of the Dual Function). *Consider any distribution $P \in \mathcal{P}'$ which satisfies Assumption 1. Then, the dual function $\text{DUAL}(\eta, x', P ; \mu)$ for a fixed arm x' and reward distribution P is upper semi-continuous at any point (η, μ) in its domain. Furthermore, the dual function is also continuous for any fixed arm x' at any point (η_0, P, μ_0) with $\eta_0(x) > 0$ and $\eta_0(x) - \lambda_0(r, x) - \beta_0 - \alpha_0 \cdot r \mathbb{1}(x = x') > 0$ for all $x \in \tilde{X}(P)$ and $r \in \mathfrak{R}$.*

G.3.1 Proof of Proposition 6

We first argue that the dual function is upper semi-continuous⁷ at any point in its domain. We decompose the dual function into

$$\begin{aligned} \text{DUAL}(\eta, x', Q ; \mu) = & \sum_{x \in \tilde{X}(Q), r \in \mathfrak{R}} \eta(x) \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')}{\eta(x)} \right) Q(r, x) \\ & - \sum_{x \in x^*(Q), r \in \mathfrak{R}} \lambda(r, x) Q(r, x) + \alpha \text{REW}^*(Q) + \beta |\tilde{X}(Q)| \\ & + \chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \end{aligned}$$

and will show that each of the terms is an upper semi-continuous function at any point of their respective domains. The first part of proposition than follows from the fact that the sum of upper semi-continuous functions is itself upper semi-continuous on the intersection of the domains.

⁷A function $y \mapsto f(y)$ is upper semi-continuous at any point in its domain if and only if its superlevel sets $\{y : f(y) \geq \gamma\}$ are closed for any $\gamma \in \mathbb{R}$.

It is clear that the characteristic function $\chi_{-\infty}(S)$ of a closed set S is upper semi-continuous on its entire domain as its γ -superlevel sets are either the closed set S for $\gamma \leq 0$ or the empty set when $\gamma > 0$. Hence, the term

$$\chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x'))$$

as an affine composition of $\chi_{-\infty}(\mathbb{R}_+)$ and $\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')$ is an upper semi-continuous function, c.f., Bertsekas (2009, Proposition 1.1.4).

The log perspective function $(u, v) \mapsto u \log(v/u)$ is upper semi-continuous on its entire domain; see Combettes (2018, Proposition 2.3). Hence, the affine composition function

$$\eta(x) \log((\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x'))/\eta(x))$$

of the log perspective function and $(\eta(x), \eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x'))$ is upper semi-continuous too; c.f., Bertsekas (2009, Proposition 1.1.4). As we have that $P(r, x) \geq 0$, the second term

$$\eta(x) \log\left(\frac{(\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x'))}{\eta(x)}\right) P(r, x) \quad (54)$$

is upper semi-continuous as well, c.f., Bertsekas (2009, Proposition 1.1.5). The remaining terms $\beta|\tilde{X}(P)|$, $\alpha \text{REW}^*(P)$ and $-\sum_{x \in x^*(P), r \in \mathfrak{R}} \lambda(r, x) P(r, x)$ are linear and thus upper semi-continuous functions as well.

Next, we argue that the dual function is continuous as well. This follows immediately from the fact that the problematic terms, presented in Equation (54), are finite and continuous when $\eta(x) > 0$ and $\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x') > 0$ for all $x \in X$ and $r \in \mathfrak{R}$. Indeed, the terms $\alpha \text{REW}^*(Q)$ and $\beta|\tilde{X}(Q)|$ are continuous via Lemma 7. \square

G.4 Proof of Proposition 5

Let $(\eta'_\epsilon(Q), \mu_\epsilon(Q))$ be the ϵ -suboptimal selection from the feasible set of problem presented in Equation (19). The selection $(\eta'_\epsilon(Q), \mu_\epsilon(Q))$ is continuous on any sufficiently small neighborhood $\{Q : \|Q - P\|_\infty \leq \kappa\}$ around P due to Proposition 4 and the fact that \mathcal{P}' is an open set. We can hence assume $\kappa > 0$ is sufficiently small so that $\max\{\|\eta'_\epsilon(Q) - \eta'_\epsilon(P)\| : \|Q - P\|_\infty \leq \kappa\} \leq \delta$, where $\delta > 0$ is an arbitrary positive number. Then, for any Q_2 so that

$\|Q_2 - P\|_\infty \leq \kappa$, we have

$$\begin{aligned}
& \|\text{SU}(Q_1, \eta'_\epsilon(Q_2) ; \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(P)\|_2 \\
& \leq \|\text{SU}(Q_1, \eta'_\epsilon(Q_2) ; \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(Q_2)\|_2 + \|\eta'_\epsilon(Q_2) - \eta'_\epsilon(P)\|_2 \\
& \leq \|\text{SU}(Q_1, \eta'_\epsilon(Q_2) ; \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(Q_2)\|_2 + \delta.
\end{aligned} \tag{55}$$

It remains to be shown that the term $\|\text{SU}(Q_1, \eta'_\epsilon(Q_2) ; \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(Q_2)\|_2$ diminishes to zero for all Q_1 and Q_2 such that $\|Q_1 - P\| \leq \kappa$ and $\|Q_2 - P\| \leq \kappa$. From Lemma 7, we may assume that $\kappa > 0$ is sufficiently small to guarantee

$$\tilde{X}_d(Q_2) = \tilde{X}_d(Q_1) = \tilde{X}_d(P).$$

It should be remarked that the term $\|\text{SU}(Q_1, \eta'_\epsilon(Q_2) ; \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(Q_2)\|_2$ coincides with the minimum in problem (18). Hence, the objective value of any feasible selection from problem (18) immediately serves as an upper bound on $\|\text{SU}(Q_1, \eta'_\epsilon(Q_2) ; \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(Q_2)\|_2$ as we discuss next. We will require the following supporting Lemma. \square

Lemma 23. *Let $\overline{\text{DUAL}}(\eta'_\epsilon(Q_2), x', Q_2 ; \mu_\epsilon(x', Q_2)) \geq 1$ for some arm $x' \in X$ and all Q_2 with $\|Q_2 - P\| \leq \kappa$. There exists $\delta(\kappa, x, P)$ (independent of Q_2) such that*

$$\overline{\text{DUAL}}(\eta'_\epsilon(Q_2), x', Q_1 ; \mu_\epsilon(x', Q_2)) \geq 1 - \delta(\kappa, x', P)$$

for all Q_1 with $\|Q_1 - P\|_\infty \leq \kappa$ and $\lim_{\kappa \rightarrow 0} \delta(\kappa, x', P) = 0$.

Since $(\eta'_\epsilon(Q_2), \mu_\epsilon(Q_2))$ is an ϵ -suboptimal solution to problem (15) with $Q = Q_2$, we have the information condition $\overline{\text{DUAL}}(\eta'_\epsilon(Q_2), x, Q_2 ; \mu_\epsilon(x, Q_2)) \geq \text{DUAL}(\eta'_\epsilon(Q_2), x, Q_2 ; \mu_\epsilon(x, Q_2)) \geq 1 \quad \forall x \in \tilde{X}_d(P)$ for all Q_2 such that $\|Q_2 - P\| \leq \kappa$ where we select κ to be sufficiently small to ensure that $\tilde{X}_d(Q_2) = \tilde{X}_d(P)$ which is possible as $P \in \mathcal{P}'$. From positive homogeneity of the test dual function and Lemma 23 it follows that we can consider for any P a small enough $\kappa > 0$ such that

$$\overline{\text{DUAL}}(\eta'_\epsilon(Q_2), x, Q_1 ; \mu_\epsilon(x, Q_2)) \geq 1 - \delta(\kappa, P) \quad \forall x \in \tilde{X}_d(Q_1)$$

for all Q_1 and Q_2 with $\|Q_1 - P\| \leq \kappa$ and $\|Q_2 - Q_1\| \leq \kappa$ where $\delta(\kappa, P) = \max_{x \in \tilde{X}_d(P)} \delta(\kappa, x, P)$ and $\lim_{\kappa \rightarrow 0} \delta(\kappa, P) = 0$. Suppose that we consider κ small enough so that $\delta(\kappa, P) \leq \delta$, where $\delta > 0$ is taken to be such that $(C(P) + \epsilon)/(1 - \delta) \leq C(P) + 2\epsilon$. Then, we have that $\eta'_\epsilon(Q_2)/(1 - \delta)$ is feasible in minimization problem (18). Hence, the term of interest is bounded by $\|\text{SU}(Q_1, \eta'_\epsilon(Q_2) ; \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(Q_2)\|_2 \leq \|\eta'_\epsilon(Q_2)/(1 - \delta) - \eta'_\epsilon(Q_2)\|_2 \leq \|\eta'_\epsilon(Q_2)\|_2(1/(1 - \delta) - 1)$.

We thus have from Equation (55) finally that

$$\begin{aligned} \|\text{SU}(Q_1, \eta'_\epsilon(Q_2) ; \mu_\epsilon(Q_2), \epsilon) - \eta'_\epsilon(Q_2)\|_2 &\leq \|\eta'_\epsilon(Q_2)\|_2(1/(1-\delta) - 1) + \delta \\ &\leq (\|\eta'_\epsilon(P)\|_2 + \delta)(1/(1-\delta) - 1) + \delta. \end{aligned}$$

The result follows as $\delta > 0$ can be made arbitrarily small and $\|\eta'_\epsilon(P)\|_2$ is bounded. \square

G.4.1 Proof of Lemma 23

Fix some arm $x' \in X$. For Q such $\|Q - P\| \leq \kappa$, we can simplify the dual function

$$\begin{aligned} \text{DUAL}(\eta, x', Q ; \mu) &= \sum_{x \in \tilde{X}(Q), r \in \mathfrak{R}} \eta(x) \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')}{\eta(x)} \right) Q(r, x) \\ &\quad - \sum_{x \in x^*(Q), r \in \mathfrak{R}} \lambda(r, x) Q(r, x) + \alpha \text{REW}^*(Q) + \beta |\tilde{X}(Q)| \\ &\quad + \chi_{-\infty}(\eta(x) \geq \lambda(r, x) + \beta + \alpha \cdot r \mathbb{1}(x = x')) \\ &= \sum_{x \in \tilde{X}_0, r \in \mathfrak{R}} \eta(x) \log \left(\frac{\eta(x) - \lambda(r, x) - \beta - \alpha \cdot r \mathbb{1}(x = x')}{\eta(x)} \right) Q(r, x) \\ &\quad + \sum_{r \in \mathfrak{R}} (\alpha r - \lambda(r, x_0^*)) Q(r, x_0^*) + \beta |\tilde{X}(P)| \end{aligned}$$

using the fact that $|\tilde{X}(Q)| = |\tilde{X}(P)| = |X| - 1$ and $\text{REW}^*(Q) = \sum_{r \in \mathfrak{R}} r Q(r, x_0^*)$. As a function of the reward distribution the dual is hence recognized to be affine in Q for $\|Q - P\| \leq \kappa$. We may use the standard inequality $u^\top v \geq -\|u\|_1 \|v\|_\infty$ to bound the dual function as

$$\begin{aligned} &\overline{\text{DUAL}}(\eta'_\epsilon(Q_2), x, Q_1 ; \mu_\epsilon(x, Q_2)) \\ &\geq \text{DUAL}(\eta'_\epsilon(Q_2), x, Q_1 ; \mu_\epsilon(x, Q_2)) \\ &\geq \min_{\|Q - Q_2\|_\infty \leq 2\kappa} \text{DUAL}(\eta'_\epsilon(Q_2), x, Q ; \mu_\epsilon(x, Q_2)) \\ &\geq \text{DUAL}(\eta'_\epsilon(Q_2), x, Q_2 ; \mu_\epsilon(x, Q_2)) - \\ &\quad \sum_{x \in \tilde{X}_0, r \in \mathfrak{R}} 2\kappa \left| \eta'_\epsilon(x, Q_2) \log \left(\frac{\eta'_\epsilon(x, Q_2) - \lambda_\epsilon(r, x, Q_2) - \beta_\epsilon(Q_2) - \alpha_\epsilon(Q_2) \cdot r \mathbb{1}(x = x')}{\eta'_\epsilon(x, Q_2)} \right) \right| \\ &\quad - \sum_{r \in \mathfrak{R}} 2\kappa |\alpha_\epsilon(Q_2) - \lambda_\epsilon(r, x_0^*, Q_2)| \\ &\geq 1 - \sum_{x \in \tilde{X}_0, r \in \mathfrak{R}} 2\kappa \left| \eta'_\epsilon(x, Q_2) \log \left(\frac{\eta'_\epsilon(x, Q_2) - \lambda_\epsilon(r, x, Q_2) - \beta_\epsilon(Q_2) - \alpha_\epsilon(Q_2) \cdot r \mathbb{1}(x = x')}{\eta'_\epsilon(x, Q_2)} \right) \right| \\ &\quad - \sum_{r \in \mathfrak{R}} 2\kappa |\alpha_\epsilon(Q_2) - \lambda_\epsilon(r, x_0^*, Q_2)|, \end{aligned}$$

where we exploit the triangle inequality $\|Q_2 - Q_1\|_\infty \leq \|Q_2 - P\|_\infty + \|P - Q_1\|_\infty \leq 2\kappa$. We will now indicate that absolute values appearing in the previous inequality can be bounded uniformly for all Q_2 with $\|Q_2 - P\|_\infty \leq \kappa$ from which the claim follows. That is, we take

$$\begin{aligned} & \delta(\kappa, x, P) \\ & := \max_{\|Q_2 - P\|_\infty \leq \kappa} \sum_{x \in \tilde{X}_0, r \in \mathfrak{R}} 2\kappa \left| \eta'_\epsilon(x, Q_2) \log \left(\frac{\eta'_\epsilon(x, Q_2) - \lambda_\epsilon(r, x, Q_2) - \beta_\epsilon(Q_2) - \alpha_\epsilon(Q_2) r \mathbb{1}(x = x')}{\eta'_\epsilon(x, Q_2)} \right) \right| \\ & \quad + \sum_{r \in \mathfrak{R}} 2\kappa |\alpha_\epsilon(Q_2) - \lambda_\epsilon(r, x_0^*, Q_2)|. \end{aligned}$$

We have indeed that $\lim_{\kappa \rightarrow 0} \delta(\kappa, x, P) = 0$ if we can show

$$\begin{aligned} & \min_{\|Q_2 - P\|_\infty \leq \kappa} \eta'_\epsilon(x, Q_2) > 0, \\ & \min_{\|Q_2 - P\|_\infty \leq \kappa} \eta'_\epsilon(x, Q_2) - \lambda_\epsilon(r, x, Q_2) - \beta_\epsilon(Q_2) - \alpha_\epsilon(Q_2) \cdot r \mathbb{1}(x = x') > 0, \\ & \max_{\|Q_2 - P\|_\infty \leq \kappa} \eta'_\epsilon(x, Q_2) - \lambda_\epsilon(r, x, Q_2) - \beta_\epsilon(Q_2) - \alpha_\epsilon(Q_2) \cdot r \mathbb{1}(x = x') < \infty, \end{aligned}$$

and $\max_{\|Q_2 - P\|_\infty \leq \kappa} |\alpha_\epsilon(Q_2) - \lambda_\epsilon(r, x_0^*, Q_2)| < \infty$ for all $r \in \mathfrak{R}$ and $x \in \tilde{X}_0$.

Define

$$R(\kappa) = \max_{\|Q - P\|_\infty \leq \kappa} \|\mu_\epsilon(Q)\|_2^2 + \|\eta'_\epsilon(Q)\|_2^2.$$

The radius $R(\kappa)$ is finite due to the classical boundedness theorem which insures that continuous functions on compact sets are bounded. We have that both $\|\mu_\epsilon(Q_2)\|_2^2 \leq R(\kappa)$ and $\|\eta'_\epsilon(Q_2)\|_2^2 \leq R(\kappa)$ are uniformly bounded for all Q_2 such that $\|Q_2 - P\| \leq \kappa$. Hence, we have $\max_{\|Q_2 - P\|_\infty \leq \kappa} \eta'_\epsilon(x, Q_2) - \lambda_\epsilon(r, x, Q_2) - \beta_\epsilon(Q_2) - \alpha_\epsilon(Q_2) \cdot r \mathbb{1}(x = x') < \infty$ as well as $\max_{\|Q_2 - P\|_\infty \leq \kappa} |\alpha_\epsilon(Q_2) - \lambda_\epsilon(r, x_0^*, Q_2)| < \infty$ for all $r \in \mathfrak{R}$ and $x \in \tilde{X}_0$. Furthermore, by construction, the logarithmic rate $\eta'_\epsilon(Q_2)$ is uniformly bounded away from zero for all Q_2 such that $\|Q_2 - P\| \leq \kappa$. Likewise, we have that for any $x \in \tilde{X}$, $x' \in \tilde{X}_d$, and $r \in \mathfrak{R}$, the terms $\eta'_\epsilon(x, Q_2) - \lambda_\epsilon(r, x, Q_2) - \beta_\epsilon(Q_2) - \alpha_\epsilon(Q_2) \cdot r \mathbb{1}(x = x') > 0$ are also uniformly bounded away from zero for all Q_2 such that $\|Q_2 - P\| \leq \kappa$. \square

H Complexity of DUSA

The main computational bottleneck of DUSA is solving the deep update Algorithm 3 in which we need to find an ϵ -suboptimal exploration rate $\eta'_\epsilon(P_t)$ to the minimization problem (15) characterizing the lower regret bound $C(P_t)$. As we have mentioned in Section 6, we

Its last two sets of constraints replace the last constraint of the dual problem (15), i.e., $1 \leq \text{DUAL}(\eta, x', P ; \mu(x'))$, using exponential cones. To make it clear, we start with examining the second set of constraints that involves the exponential cone K_{exp} . Note that by Chandrasekaran and Shah (2017), we have $(u, v, w) \in K_{exp} \iff \{(u, v, w) : v \log(w/v) \geq u\}$. Considering this, the second set of constraints in problem (56) can be written as

$$\eta(x) \log \left(\frac{\eta(x) - \lambda(r, x)(x') - \beta(x') - \alpha(x') \cdot r \mathbb{1}(x = x')}{\eta(x)} \right) \geq -\ell(r, x)(x').$$

Observe that the left hand side is $\eta(x) \log(\omega)$, where variable $\omega = (\eta(x) - \lambda(r, x)(x') - \beta(x') - \alpha(x') r \mathbb{1}(x = x')) / \eta(x)$. Thus, $-\ell(r, x)(x')$ is a lower bound on $\eta(x) \log(\omega)$, where this term, i.e., $\eta(x) \log(\omega)$, appears in the definition of the dual function DUAL , given in Equation (13). Having this interoperation of variable ℓ 's in mind, it is easy to see that the last set of constraints of problem (56) ensures that $1 \leq \text{DUAL}(\eta, x', P ; \mu(x'))$. \square

Proposition 7 presents a conic representation of the convex problem (15). However, not all convex optimization problems are easy to solve. Convex problems do admit polynomial time algorithms when their conic representation involves exclusively cones which admit an “efficient” representation.⁹ Efficiently representable cones, such as the positive orthant, second-order cone, or positive semi-definite cone, are precisely those for which an efficient barrier function is known and consequently the associated conic optimization problem can be tractably solved using a modern interior point method. That is, the conic minimization problem can be solved in time polynomial in the number of arms $|X|$, the number of rewards $|\mathfrak{R}|$ as well as in the number of required accurate digits of the optimal solution. As each of the cones \mathbb{R} , \mathbb{R}_+ and K_{exp} admits an efficient barrier, the dual cone \mathcal{K}^* will ultimately determine whether the dual deep update can be carried out tractably or not.

Recall that the dual cone \mathcal{K}^* is defined as the dual of the conic hull \mathcal{K} of the set of all reward distributions \mathcal{P} and captures all structural information in the bandit problem. It is well known, c.f., Nesterov (2004), that $\mathcal{K}^* = \text{cone}(\mathcal{P})^*$ is efficiently representable if and only if its dual $\mathcal{K} = \text{cone}(\mathcal{P})$ is too. Hence, whether or not the convex problem (15) can be solved tractably depends only on whether we can represent the structural information regarding the bandit reward distributions efficiently.

⁹One particular class of NP-hard convex optimization problems are the copositive optimization problems discussed by Burer (2009).

I Supporting Results on Duality and Continuity

In many of the proofs in this paper, we need some basic results concerning the existence of the strong duality and continuity of the minimum of parametric convex optimization problems. The fact that strong duality and continuity are intimately related has been long established. See for instance Bank et al. (1982, Section 4.4) for a more in depth discussion of this connection. In this section, for readers convenience, we present the most relevant results on these subjects to our paper.

I.1 Duality

Consider the following standard minimization problem

$$\begin{aligned}
 v^*(\theta) &:= \inf_y f(y, \theta) \\
 \text{s.t. } & g(y, \theta) \leq 0, && [\text{dual variable } \mu_1 \in \mathbb{R}_+^{m_1}] \\
 & h(y, \theta) = 0, && [\text{dual variable } \mu_2 \in \mathbb{R}^{m_2}] \\
 & y \in K, && [\text{dual variable } \mu_3 \in K^*] \\
 & y \in \Gamma(\theta) \subseteq Y && [\text{not dualized}]
 \end{aligned} \tag{57}$$

parameterized in $\theta \in \Theta$. In what follows we assume that Y and Θ are finite dimensional linear spaces. The inequality and equality constraint functions g and h , respectively, characterize the first two functional constraints. Note that there are m_1 inequality constraints, where the i^{th} inequality constraint is associated with the inequality constraint function g_i . Similarly, there are m_2 equality constraints, where the i^{th} equality constraint is associated with equality constraint function h_i . The third constraint characterizes a conic set constraint. As before, we denote with K^* the dual cone of the closed convex cone K . The mapping $\Gamma(\theta)$, also referred to as set-valued mapping, characterizing the ultimate parametric set constraint is assumed to be non-empty for all $\theta \in \Theta$. The objective function $f(y, \theta)$ itself implicitly constrains the optimization problem to its domain $\text{dom}(f(y, \theta)) := \{y : f(y, \theta) < +\infty\}$. We will denote with $y^*(\theta)$ the set of all minimizers in problem (57). We remark that $y^*(\theta)$ may be an empty set as the infimum is not necessarily attained. We will denote $M(\theta)$ the feasible set mapping in problem (57). Consequently, we have that $y^*(\theta) \subseteq M(\theta)$. The following standard constraint qualification condition plays a pivotal role in establishing both strong duality for our minimization problem (57) as well as the continuity of its minimum $v^*(\theta)$ in the parameter θ .

Definition 3 (Slater's Constraint Qualification Condition). *We say that the feasible region*

of the primal optimization problem (57) satisfies Slater's constraint qualification condition at some given $\theta_0 \in \Theta$ if there exists a Slater point $y_0 \in \Gamma(\theta_0)$ such that $g(y_0, \theta_0) < 0$, $h(y_0, \theta_0) = 0$ and $y_0 \in \text{int}(K)$. We will denote the set of all such Slater points as

$$M_0(\theta) := \Gamma(\theta) \cap \{y \in \text{int}(K) : g(y, \theta) < 0, h(y, \theta) = 0\}.$$

Consider the Lagrangian function $L(y, \theta ; \mu = (\mu_1, \mu_2, \mu_3)) = f(y, \theta) + \mu_1^\top g(y, \theta) + \mu_2^\top h(y, \theta) - \mu_3^\top y$ and the associated dual problem

$$\begin{aligned} u^*(\theta) = \sup_{\mu} \quad & q(\theta ; \mu = (\mu_1, \mu_2, \mu_3)) \\ \text{s.t.} \quad & \mu_1 \in \mathbb{R}_+^{m_1}, \\ & \mu_2 \in \mathbb{R}^{m_2}, \\ & \mu_3 \in K^* \end{aligned} \tag{58}$$

for the dual function $q(\theta ; \mu) := \inf\{L(y, \theta ; \mu) : y \in \Gamma(\theta)\}$. As indicated in the minimization problem (57), the dual variables $\mu = (\mu_1, \mu_2, \mu_3)$ are associated with the primal functional constraints $g(y, \theta) \leq 0$, $h(y, \theta) = 0$ and the primal set constraint $y \in K$. The other primal set constraint $y \in \Gamma(\theta)$ is not dualized and consequently has no associated dual variables. Note that we trivially have the weak duality inequality $u^*(\theta) \leq v^*(\theta)$. In order to have strong duality, i.e., $u^*(\theta) = v^*(\theta)$, one typically needs to require convexity and some type of constraint qualification condition. We recall the following classical result which can be derived immediately from Boyd and Vandenberghe (2004, Section 5.9).

Proposition 8 (Strong Duality). *Fix a parameter value $\theta_0 \in \Theta$. Let the objective $f(y, \theta_0)$ be a convex function in y . Likewise, the inequality constraint functions $g_i(y, \theta_0)$, $i \in \{1, \dots, m_1\}$ are convex in y while the equality constraint functions $h_i(y, \theta_0)$, $i \in \{1, \dots, m_2\}$ are affine in y . Assume that the primal minimum in problem (57) is bounded from above and that*

$$M_0(\theta_0) \cap \text{relint}(\text{dom}(f(y, \theta_0))) \cap \Gamma(\theta_0) \neq \emptyset.$$

That is, there is a Slater point which is in the intersection of the relative interior of the domain of $f(y, \theta_0)$ and the parametric set constraint $\Gamma(\theta_0)$. Then, we have the strong duality equality $u^(\theta_0) = v^*(\theta_0)$ and the set of dual optimal solutions is non-empty and compact.*

Proof. It is evident that the primal can be characterized as the convex optimization problem

$$\begin{aligned} v^*(\theta_0) &= \inf_y f(y, \theta_0) + \chi_{+\infty}(\Gamma(\theta_0))(y) \\ \text{s.t. } & g(y, \theta_0) \leq 0, \\ & h(y, \theta_0) = 0, \\ & y \in K, \end{aligned}$$

where the convex characteristic function $\chi_{+\infty}(\Gamma(\theta_0))(y)$ is zero if $y \in \Gamma(\theta_0)$ and positive infinity otherwise. From Boyd and Vandenberghe (2004, Section 5.9), we know that strong duality, i.e., $u^*(\theta_0) = v^*(\theta_0)$, holds if we can find a point

$$y_0 \in \text{relint}(\text{dom}(f(y, \theta_0)) \cap \Gamma(\theta_0) \cap \text{dom}(g(y, \theta_0)) \cap K)$$

such that $y_0 \in \text{int}(K)$ and $g(y_0, \theta_0) < 0$ and $h(y_0, \theta_0) = 0$. From Bertsekas (2009, Proposition 1.3.8) and the fact that $\text{relint}(C) \subseteq \text{int}(C)$ for any set C it follows that

$$M_0(\theta_0) \cap \text{relint}(\text{dom}(f(y, \theta_0)) \cap \Gamma(\theta_0)) \neq \emptyset$$

is sufficient to guarantee that such a point y_0 exists. □

I.2 Continuity

The same constraint qualification condition stated in Definition 3 also plays an important role for the minimum $v^*(\theta)$ of problem (57) to be a continuous function. Establishing continuity of a minimum typically requires its objective function to lower semi-continuous and its feasible set mapping to be a lower semi-continuous mapping in addition to some type of constraint qualification condition. As the definition of what constitutes semi continuous mappings is not universally consistent¹⁰ throughout the literature we state for completeness the definition taken in this paper explicitly.

Definition 4 (Semi-continuous Mappings). *A mapping $G : A \rightarrow B$ is said to be upper semi-continuous at point $a \in A$ if for any open neighborhood V of $G(a)$, there exists a neighbourhood U such that all $x \in U$, $G(x)$ is a subset of V . A mapping $G : A \rightarrow B$ is said to be lower semi-continuous at the point $a \in A$ if for any open set V intersecting $G(a)$, there exists a neighbourhood U such that for any $x \in U$, the intersection of $G(x)$ and V is non-empty. A mapping is lower semi-continuous if it is lower semi-continuous at any point $a \in A$ and upper*

¹⁰For instance, Berge (1997) requires upper semi-continuous mappings to be compact valued while we following Aubin and Frankowska (2009) do not.

semi-continuous if it is upper semi-continuous at any point $a \in A$.

In this paper, we only have a need to show continuity of the minimum $v^*(\theta)$ in problem (57) when only the first and final constraint are nontrivial. That is, it suffices here to consider problems with no explicit cone and equality constraints. We have indeed the following direct corollary of a more general result by Bank et al. (1982, Theorem 3.1.6) in arbitrary normed linear spaces.

Proposition 9 (Lower Semi-continuity of Feasible Sets). *Let the mapping $\Gamma(\theta)$ be a lower semi-continuous mapping at $\theta_0 \in \Theta$, the set $\Gamma(\theta_0)$ be convex and the feasible set in problem (57) satisfy the Slater's constraint qualification, i.e., $M_0(\theta_0) \neq \emptyset$, with no cone or equality constraints ($K = \mathbb{R}^n$ and $m_2 = 0$). Further,*

(i) *let the inequality constraint functions $g_i(y, \theta)$, $i \in \{1, \dots, m_1\}$, be upper semi-continuous on $M_0(\theta_0) \times \{\theta_0\}$ and*

(ii) *let the inequality constraint functions $y \mapsto g_i(y, \theta_0)$, $i \in \{1, \dots, m_1\}$, be convex on $\Gamma(\theta_0)$.*

Then, the mapping $M(\theta)$ is lower semi-continuous at θ_0 .

The following standard result by the same authors Bank et al. (1982, Theorem 4.3.3) can then be invoked to prove continuity of the primal minimum.

Proposition 10 (Continuity of Optimal Value). *Let the feasible set mapping $M(\theta)$ be lower semi-continuous at $\theta_0 \in \Theta$ in the primal minimization problem (57) with no cone or equality constraints ($K = \mathbb{R}^n$ and $m_2 = 0$). Furthermore, assume that its feasible region $M(\theta)$ is convex for all $\theta \in \Theta$ and the mapping $M(\theta)$ is closed at θ_0 . Finally, assume that $y^*(\theta_0)$ is a non-empty bounded set. Let the objective function $(y, \theta) \mapsto f(y, \theta)$ satisfy the following technical conditions:*

1. *the function $f(y, \theta)$ is lower semi-continuous on $Y \times \{\theta_0\}$.*
2. *the function $f(y, \theta)$ is upper semi-continuous at some (y_0^*, θ_0) with $y_0^* \in y^*(\theta_0)$.*
3. *the function $f(y, \theta)$ is convex in $y \in Y$ for each parameter value $\theta \in \Theta$.*

Then, the minimum $v^(\theta)$ is continuous at θ_0 .*