
When Combinatorial Thompson Sampling meets Approximation Regret

Pierre Perrault

Idemia

pierre.perrault@idemia.com

Abstract

We study the Combinatorial Thompson Sampling policy (CTS) for combinatorial multi-armed bandit problems (CMAB), within an *approximation regret* setting. Although CTS has attracted a lot of interest, it has a drawback that other usual CMAB policies do not have when considering non-exact oracles: for some oracles, CTS has a poor approximation regret (scaling linearly with the time horizon T) [Wang and Chen, 2018]. A study is then necessary to discriminate the oracles on which CTS could learn. This study was started by Kong et al. [2021]: they gave the first approximation regret analysis of CTS for the *greedy oracle*, obtaining an upper bound of order $\mathcal{O}(\log(T)/\Delta^2)$, where Δ is some minimal reward gap. In this paper, our objective is to push this study further than the simple case of the greedy oracle. We provide the first $\mathcal{O}(\log(T)/\Delta)$ approximation regret upper bound for CTS, obtained under a specific condition on the approximation oracle, allowing a reduction to the exact oracle analysis. We thus term this condition REDUCE2EXACT, and observe that it is satisfied in many *concrete* examples. Moreover, it can be extended to the *probabilistically triggered arms* setting, thus capturing even more problems, such as *online influence maximization*.

1 Introduction

Stochastic multi-armed bandits (MAB) [Robbins, 1952, Berry and Fristedt, 1985, Lai and Robbins, 1985] are decision-making problems in which an *agent* acts sequentially in an uncertain environment. At each round $t \in \mathbb{N}^*$, the agent must select one arm from a fixed set of n arms, denoted by $[n] \triangleq \{1, \dots, n\}$, using a *policy*, based on the feedback from the previous rounds. Then it gets as feedback an *outcome* $X_{i,t} \in \mathbb{R}$ — a random variable sampled from \mathbb{P}_{X_i} , independently from previous rounds — where i is the selected arm and \mathbb{P}_{X_i} is a probability distribution — unknown to the agent — of mean μ_i^* . The goal for the agent is to maximize the *cumulative reward* over a total of T rounds (T is the *time horizon* and may be unknown). The performance metric of a policy is its *regret* R_T , which is the expectation of the difference over T rounds between the cumulative reward of the policy that always picked the arm with the highest expected reward and the cumulative reward of the learning policy. MAB models the so called dilemma between exploration and exploitation, i.e., whether to continue exploring arms to obtain more information (and thus strengthen the confidence in the estimates of the distributions \mathbb{P}_{X_i}), or to use the information gathered by playing the best arm according to the observations so far.

In this paper, we study stochastic combinatorial multi-armed bandit (CMAB), with *semi-bandit feedback*, a.k.a. stochastic semi-bandit (still abbreviated as CMAB in this paper), an extension of MAB where the agent plays an *action* (also called *super-arm*) $A_t \in \mathcal{A} \subset \mathcal{P}([n])$ at each round t , where \mathcal{A} is fixed and called *action space*. The feedback includes the outcomes of all base arms in

the played super-arm.¹ The expected reward, given A_t , is assumed to be in the form² $r(A_t, \boldsymbol{\mu}^*)$, where $\boldsymbol{\mu}^* \in \mathbb{R}^n$ is the unknown vector of expectations (traditionally, the reward is linear and equal to $\mathbf{e}_{A_t}^\top \boldsymbol{\mu}^*$). In recent years, CMAB has attracted a lot of interest (see e.g. [Cesa-Bianchi and Lugosi \[2012\]](#), [Gai et al. \[2012\]](#), [Chen et al. \[2013, 2016\]](#), [Kveton et al. \[2015\]](#), [Wang and Chen \[2017\]](#), [Perrault \[2020\]](#)), particularly due to its wide applications in network routing, online advertising, recommender system, influence marketing, etc.

Many CMAB policies are based on the *Upper Confidence Bound* (UCB) approach, extending the classical UCB policy [[Auer et al., 2002](#)] from MAB to CMAB. This type of approach uses an optimistic estimate $\boldsymbol{\mu}_t$ of $\boldsymbol{\mu}^*$ (i.e., for which the reward function is overestimated), lying in a well-chosen confidence region. Then, the action is chosen by plugging $\boldsymbol{\mu}_t$ inside an *oracle* (typically, $\text{Oracle}(\boldsymbol{\mu}^*)$ is a maximizer of the reward function $A \mapsto r(A, \boldsymbol{\mu}^*)$). An example of such policy is *Combinatorial Upper Confidence Bound* (CUCB) [[Chen et al., 2013](#), [Kveton et al., 2015](#)], that uses a Cartesian product of the individual confidence intervals of each arm as a confidence region. For mutually independent arms, [Combes et al. \[2015\]](#) provided the UCB-style policy *Efficient Sampling for Combinatorial Bandit* (ESCB), building a tighter axis-aligned ellipsoidal confidence region around the empirical mean, which helps to better restrict the exploration. [Degenne and Perchet \[2016\]](#) provided a policy called OLS-UCB, leveraging a *sub-Gaussianity* assumption on the arms to generalize the ESCB approach. These policies have been further extended to more general settings afterwards [[Perrault et al., 2020c,d, 2019a](#)]. Although improving CUCB, all these generalizations are inefficient in terms of computation time.

Another paradigm that has recently gained interest (and which will be our focus in this paper) is to rely on *Thompson Sampling* (TS) instead of UCB, still targeting *frequentist regret*. Although introduced much earlier by [Thompson \[1933\]](#), the theoretical analysis of TS for MAB is quite recent: [Kaufmann et al. \[2012\]](#), [Agrawal and Goyal \[2012\]](#) gave a regret bound matching the UCB policy theoretically. Moreover, TS often performs better than UCB in practice, making TS an attractive policy for further investigations. For CMAB, TS extends to *Combinatorial Thompson Sampling* (CTS). In CTS, the unknown mean $\boldsymbol{\mu}^*$ is associated with a belief (a prior distribution, that could be e.g. a product of Beta or Gaussian distributions) updated to a posterior with the Bayes' rule, each time a feedback is received. In order to choose an action at round t , CTS draws a sample $\boldsymbol{\theta}_t$ from the current belief, and plays the action given by $\text{Oracle}(\boldsymbol{\theta}_t)$. CTS is an attractive policy because it has similar advantages to the previously mentioned policies working with ellipsoidal confidence regions while being, like CUCB, computationally efficient. Indeed, recently, for mutually independent arms and sub-Gaussian arms respectively, [Wang and Chen \[2018\]](#), [Perrault et al. \[2020a\]](#) proposed tight analyses of CTS.

Unlike UCB-based policies, the analysis of CTS is valid only when Oracle is exact, i.e., when

$$\text{Oracle}(\boldsymbol{\mu}) \in \arg \max_{A \in \mathcal{A}} r(A, \boldsymbol{\mu}).$$

Although this holds true for many combinatorial problem described by the pair (r, \mathcal{A}) (we recall that r is usually linear), there exist some problems where the requirement on Oracle has to be relaxed in order to make it tractable. This is usually done considering an α -approximation oracle [[Chen et al., 2013, 2016](#), [Wen et al., 2016](#)], for $\alpha \in (0, 1)$:

$$r(\text{Oracle}(\boldsymbol{\mu}), \boldsymbol{\mu}) \geq \alpha \max_{A \in \mathcal{A}} r(A, \boldsymbol{\mu}). \quad (1)$$

Under an α -approximation oracle, the benchmark cumulative reward is the α -fraction of the optimal reward, leading to the notion of *approximation regret* [[Kakade et al., 2009](#), [Streeter and Golovin, 2009](#), [Chen et al., 2016](#)].

Aware of the limitation of their CTS analysis (that works only with exact oracles), [Wang and Chen \[2018\]](#) also proved, in their Theorem 2, that this limitation is not a technical artifact. More precisely, they provided a specific CMAB instance with an associated approximation oracle, such that CTS on this instance and with this oracle must have a regret scaling linearly in T . Although this negative result is of great interest to the research community, some concerns limit its consideration. Indeed,

¹Note that we will consider the *probabilistically triggered arms* extension in this paper, i.e., where the feedback is on triggered arms. For brevity, we do not present this generalization in the introduction.

²Henceforth, we typeset vectors and matrices in bold and indicate components with indices, e.g., $\mathbf{a} = (a_i)_{i \in [n]} \in \mathbb{R}^n$. We also let \mathbf{e}_i be the i^{th} canonical unit vector of \mathbb{R}^n , and define the *incidence vector* of any subset $A \subset [n]$ as $\mathbf{e}_A \triangleq \sum_{i \in A} \mathbf{e}_i$. We denote by $\mathbf{a} \odot \mathbf{b} \triangleq (a_i b_i)$ the Hadamard product of two vectors \mathbf{a} and \mathbf{b} .

not only the CMAB instance provided by Wang and Chen [2018] is actually a MAB one (meaning that there is an efficient oracle that simply enumerates the arms), so the use of an approximation regret is not justified, but above all, both the oracle and the instance are uncommon and designed for the proof.

Interested in the question of whether the example provided by Wang and Chen [2018] is pathological or generalizable, Kong et al. [2021] recently initiated a study, which revealed that linear approximation regret for CTS seems to be pathological. More precisely, they derived a $\mathcal{O}(\log(T)/\Delta^2)$ bound for the specific case of a *greedy* oracle³, where Δ is some reward gap. This result is obtained by bounding the approximation regret by a *greedy regret*, that simply replaces $\alpha \max_{A \in \mathcal{A}} r(A, \mu^*)$ with $r(\text{Oracle}(\mu^*), \mu^*)$, using equation (1). They also gave a tight lower bound on the greedy regret.

In this paper, we want to explore another class of oracles covering more problems in practice. Our goal is to demonstrate that although there are instrumental examples of problems where CTS has a linear regret, the majority of concrete problems do not follow this regime, and are in fact similar to the exact oracle case.

Contributions With a specific condition on Oracle, we describe a general set of CMAB problems where the approximation regret of CTS has a $\mathcal{O}(\log(T)/\Delta)$ bound, improving by a factor $1/\Delta$ over the bound of Kong et al. [2021]. This does not contradict their lower bound, which is focused on the greedy regret. We call this set of CMAB problems REDUCE2EXACT, because, as we will see, a REDUCE2EXACT problem can be approximated using a reduction to sub-problems that can be solved exactly. Our main result is the approximation regret guaranty for REDUCE2EXACT problems, provided in Theorem 2. The REDUCE2EXACT condition on Oracle is structural and applies notably to greedy algorithms for submodular maximization. In particular, it allows to deal with problems such as *probabilistic maximum coverage*. We note that REDUCE2EXACT is compatible with the *probabilistically triggered arms* setting, which allows us to capture even more problems, such as *online influence maximization* [Wen et al., 2017, Wang and Chen, 2017]. As we want to focus on concrete problems, we provide several other examples that belongs to REDUCE2EXACT: *Metric k-center*, *Vertex cover*, *Max-Cut* and *Travelling salesman problem*.

Further related work We refer the reader to Wang and Chen [2018] for more related work on TS for combinatorial bandits. Briefly, one can mention Gopalan et al. [2014], that gave a frequentist high-probability regret bounds for TS with a general action space and feedback model — Komiyama et al. [2015], that studied TS for the m -sets action space — Wen et al. [2015], that studied TS for contextual CMAB problems, using the Bayesian regret metric (see also Russo and Van Roy [2016]).

Other known limitations of CTS Apart from the limitation related to the approximation regret that interests us in this paper, there are some other existing limitations of CTS highlighted in the literature, which we review here: The CTS policy has an exponential constant term in its regret upper bound [Wang and Chen, 2018, Perrault et al., 2020a], and Wang and Chen [2018] proved in their Theorem 3 that this is unavoidable. A similar behavior have been demonstrated in Zhang and Combes [2021], where it is shown that CTS does not scale polynomially in the ambient dimension n in general. In addition, Zhang and Combes [2021] also proved that CTS is not minimax optimal. Actually, they even proved that in high dimensions, the minimax regret of CTS is almost linear in T .

The strengths of CTS Despite the weaknesses mentioned above, CTS remains a widely used policy, mainly because of its empirical performance. Indeed, CTS generally outperforms other policies such as CUCB and ESCB [Wang and Chen, 2018, Perrault et al., 2020a]. Moreover, it is relatively simple to implement, and is computationally efficient (just like CUCB). On the theory side, another advantage is that for an exact oracle, CTS is asymptotically quasi-optimal⁴ for many settings where CUCB is not, and where ESCB is computationally inefficient [Perrault et al., 2020a]. It would be desirable that these advantages also apply to the case of approximation regret, thus motivating our investigations.

³It is worth mentioning that this oracle is one of the most common, so it is logical to focus on it first.

⁴This means that it has a distribution-dependent regret upper bound whose leading term in T has an optimal rate, up to a poly-logarithmic factor in n .

2 Model and definitions

For more generality, we consider the *probabilistically triggered arms* extension of CMAB [Chen et al., 2016, Wang and Chen, 2017], abbreviated to CMAB-T. In this context, the action $A \in \mathcal{A}$ selected is not necessarily equal to the triggered super-arm S . More precisely, the action space \mathcal{A} is no longer necessarily a subset of $\mathcal{P}([n])$ and can be infinite. At round t , the agent selects $A_t \in \mathcal{A}$, based on the history of observations $\mathcal{H}_t \triangleq \sigma(\mathbf{X}_1 \odot \mathbf{e}_{S_1}, \dots, \mathbf{X}_{t-1} \odot \mathbf{e}_{S_{t-1}})$ and a possible extra source of randomness (we denote by \mathcal{F}_t the filtration containing \mathcal{H}_t and the extra randomness of round t — in particular, A_t is \mathcal{F}_t -measurable). Then, an independent sample $\mathbf{X}_t \sim \mathbb{P}_{\mathbf{X}}$, $\mathbf{X}_t \in \mathbb{R}^n$ is drawn and a random subset $S_t \in \mathcal{S} \subset \mathcal{P}([n])$ of arms are triggered (\mathcal{S} is called *super-arm space* or *subset space*). We assume that S_t is drawn independently from a distribution $D_{\text{trig}}(A_t, \mathbf{X}_t)$ and that the outcome of an arm does not depend on whether it is triggered. In addition, if we don't have $\mathbb{P}_{\mathbf{X}} = \otimes_{i \in [n]} \mathbb{P}_{X_i}$, we assume that D_{trig} doesn't depend on \mathbf{X}_t . For the feedback, the outcome of each triggered arm is observed, i.e., $\mathbf{e}_{S_t} \odot \mathbf{X}_t$ is observed. The expected reward is of the form $r(A_t, \boldsymbol{\mu}^*)$, where r is a function defined on a domain $\mathcal{A} \times \mathcal{M}$, with $\mathcal{M} \subset \mathbb{R}^n$. The objects $\mathcal{A}, D_{\text{trig}}, r$ are known to the agent. We assume that $r(\cdot, \boldsymbol{\mu}^*)$ admits a maximum $r(A^*, \boldsymbol{\mu}^*)$ on \mathcal{A} . In the following, we give the definition of the probability that an arm $i \in [n]$ is triggered (and thus that a feedback from i is obtained) by having played a certain action $A \in \mathcal{A}$.

Definition 1 (Triggering probabilities). *The triggering probabilities are defined for all $i \in [n]$ and $A \in \mathcal{A}$ as $p_i(A) \triangleq \mathbb{P}[i \in S]$, where $S \sim D_{\text{trig}}(A, \mathbf{X})$, $\mathbf{X} \sim \mathbb{P}_{\mathbf{X}}$.*

Under an α -approximation Oracle, we use the approximation regret to evaluate the performance of a policy π , defined as follows.

Definition 2 (Approximation regret). *The T -round α -approximation regret of a learning policy π that selects action $A_t \in \mathcal{A}$ at round t is defined as follows, where the approximation gap is defined as $\Delta_t = \Delta(A_t) \triangleq 0 \vee (\alpha r(A^*, \boldsymbol{\mu}^*) - r(A_t, \boldsymbol{\mu}^*))$, with $A^* \in \arg \max_{A \in \mathcal{A}} r(A, \boldsymbol{\mu}^*)$.*

$$R_{T,\alpha}(\pi) \triangleq \mathbb{E} \left[\sum_{t \in [T]} \Delta_t \right].$$

To approach the problem of minimizing $R_{T,\alpha}$, we consider the following standard assumptions [Wang and Chen, 2017].

Assumption 1 (Approximation oracle). *The agent has access to an Oracle such that for any mean vector $\boldsymbol{\mu} \in \mathcal{M}$,*

$$r(\text{Oracle}(\boldsymbol{\mu}), \boldsymbol{\mu}) \geq \alpha r(A^*, \boldsymbol{\mu}).$$

Assumption 2 (1-norm triggering probability modulated bounded smoothness). *There exists $\mathbf{B} \in \mathbb{R}_+^n$ such that for all $A \in \mathcal{A}$, for all $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathcal{M}$,*

$$|r(A, \boldsymbol{\mu}) - r(A, \boldsymbol{\mu}')| \leq \sum_{i \in [n]} p_i(A) B_i |\mu_i - \mu'_i|.$$

Assumption 3 (Sub-Gaussianity of the outcome distribution). *$\mathbb{P}_{\mathbf{X}}$ is such that $\forall \boldsymbol{\lambda} \in \mathbb{R}^n$,*

$$\mathbb{E} \left[e^{\boldsymbol{\lambda}^\top (\mathbf{X} - \boldsymbol{\mu}^*)} \right] \leq e^{\|\boldsymbol{\lambda}\|_2^2 / 8}.$$

For example, $\mathbb{P}_{\mathbf{X}} = \otimes_{i \in [n]} \mathbb{P}_{X_i}$, and $X_i \stackrel{\text{a.s.}}{\in} [0, 1]$ (from Hoeffding's Lemma [Hoeffding, 1963]).

Definition 3 (Other definitions). *We define, for $i \in [n]$, the minimal gap of an action containing i as*

$$\Delta_{i,\min} \triangleq \inf_{A \in \mathcal{A}: p_i(A) > 0, \Delta(A) > 0} \Delta(A).$$

The minimal and maximal gaps are defined as

$$\Delta_{\min} \triangleq \min_{i \in [n]} \Delta_{i,\min} \quad \text{and} \quad \Delta_{\max} \triangleq \sup_{A \in \mathcal{A}} \Delta(A).$$

For $A \in \mathcal{A}$, we let $\mathsf{T}(A) \triangleq \{i \in [n] : p_i(A) > 0\}$ be the set of arms that are triggerable by selecting action A . We finally define

$$m \triangleq \sup_{A \in \mathcal{A}} |\mathsf{T}(A)|, \quad m^* \triangleq |\mathsf{T}(A^*)|, \quad \text{and} \quad p^* \triangleq \inf_{i \in [n], A \in \mathcal{A}: p_i(A) > 0} p_i(A).$$

Algorithm 1 CTS-BETA

Initialization: For each arm i , let $\gamma_i = \delta_i = 1$.
For all $t \geq 1$:
 Draw $\boldsymbol{\theta}_t \sim \otimes_{i \in [n]} \text{Beta}(\gamma_i, \delta_i)$.
 Play $A_t = \text{Oracle}(\boldsymbol{\theta}_t)$.
 Get the observation $\mathbf{X}_t \odot \mathbf{e}_{S_t}$, and draw $\mathbf{Y}_t \sim \otimes_{i \in S_t} \text{Bernoulli}(X_{i,t})$.
 For all $i \in S_t$ update $\gamma_i \leftarrow \gamma_i + Y_{i,t}$ and $\delta_i \leftarrow \delta_i + 1 - Y_{i,t}$.

Algorithm 2 CTS-GAUSSIAN

Input: $\beta > 1$.
Initialization: Play each arm once (if the agent knows that $\boldsymbol{\mu}^* \in [a, b]^n$, this might be skipped)
For every subsequent round t :
 Draw $\boldsymbol{\theta}_t \sim \otimes_{i \in [n]} \mathcal{N}(\bar{\mu}_{i,t-1}, N_{i,t-1}^{-1} \beta/4)$ ($\theta_{i,t} \sim \mathcal{U}[a, b]$ if $N_{i,t-1} = 0$).
 Play $A_t = \text{Oracle}(\boldsymbol{\theta}_t)$.
 Get the observations $\mathbf{X}_t \odot \mathbf{e}_{S_t}$ and let $\mathbf{Y}_t = \mathbf{X}_t$.
 Update $\bar{\mu}_{i,t-1}$ and counters accordingly.

3 Combinatorial Thompson Sampling and exact oracle analysis

In this section, we present the CTS policy, focusing on two versions, one working with a Beta prior, and the other with a Gaussian prior. Then, we present an associated analysis for the exact oracle case (i.e., with $\alpha = 1$).

3.1 Algorithms

Based on the above assumptions, we focus on two versions of CTS. The first version is CTS-BETA [Wang and Chen, 2018] (Algorithm 1), working when we assume furthermore that $\mathbb{P}_{\mathbf{X}} = \otimes_{i \in [n]} \mathbb{P}_{X_i}$ and $\mathbf{X} \stackrel{a.s.}{\in} [0, 1]^n$ (note that this actually covers the case of bounded \mathbf{X} , by adjusting the parameter \mathbf{B}). For each arm $i \in [n]$, CTS-BETA maintains a Beta prior distribution with parameters γ_i and δ_i (initialized to 1). At each round t , for each arm i , the algorithm sample $\theta_{i,t}$ from the corresponding prior, representing the current estimate of μ_i^* . Then the oracle outputs the action A_t to play according to the input vector $\boldsymbol{\theta}_t$. Based on the observation feedback, the algorithm then updates the corresponding Beta distributions. The second version is CTS-GAUSSIAN [Perrault et al., 2020a] (Algorithm 2), that works under the more general Assumption 3. It is essentially the same as CTS-BETA, except that the prior distributions are Gaussian. For both Algorithm 1 and Algorithm 2, and an arm $i \in [n]$, we define the number of time i has been triggered at the beginning of round t , called counter of arm i , as

$$N_{i,t-1} \triangleq \sum_{t' \in [t-1]} \mathbb{I}\{i \in S_{t'}\}.$$

We also define the *empirical mean* at the beginning of round t as

$$\bar{\mu}_{i,t-1} \triangleq \sum_{t' \in [t-1]} \frac{\mathbb{I}\{i \in S_{t'}\} Y_{i,t'}}{N_{i,t-1}}.$$

3.2 Analysis of CTS: the $\alpha = 1$ case

Although this is close to some known results in the current literature [Huyuk and Tekin, 2019, Perrault et al., 2020a], there is no proof for the classical $\mathcal{O}\left(\log(m) \log(T) \sum_{i \in [n]} B_i^2 / \Delta_{i,\min}\right)$ regret bound under the above assumptions in the CMAB-T setting, either for CTS-BETA (Algorithm 1) or CTS-GAUSSIAN (Algorithm 2). We thus provide such a result in Theorem 1 (the proof is postponed to Appendix A). We can notice a difference with the work of Huyuk and Tekin [2019] concerning the Assumption 2, where the triggering probabilities do not appear (and are present in the main term of their final regret bound). The main difference with Perrault et al. [2020a] is that they do not consider probabilistically triggered arms.

Theorem 1. If $\Delta_{\min} > 0$ and $p^* > 0$, the policy π described in Algorithm 1 (under Assumptions 1, 2 and $\mathbb{P}_{\mathbf{X}} = \otimes_{i \in [n]} \mathbb{P}_{X_i}$, $\mathbf{X} \stackrel{a.s.}{\in} [0, 1]^n$) or Algorithm 2 (under Assumptions 1, 2 and 3) has a regret of order

$$R_{T,1}(\pi) = \mathcal{O} \left(\sum_{i \in [n]} \frac{B_i^2 \log(m) \log(T)}{\Delta_{i,\min}} \right).$$

In addition to being a new result in itself, Theorem 1 will be useful for the $\alpha < 1$ case. Concerning the bound, after changing the CMAB setting to CMAB-T, it should be noticed that the T -independent additive constant depends on $1/p^*$, the cause being the use of the Lipschitz condition (Assumption 2) without weighting by probabilities. We think that this dependence should be avoidable, but it seems that another technique has to be considered. Finally, we remark that this kind of bound can be usually transformed into a gap-independent \sqrt{T} bound [Chen et al., 2013], however, for CTS, achieving this transformation is impossible since CTS is not minimax optimal (as we mentioned in the paragraph "Other known limitations of CTS").

4 The $\alpha < 1$ case for REDUCE2EXACT problems

We will now look at the $\alpha < 1$ case. Our strategy is based on the following observation: many approximation algorithms involve a relaxation, or a reduction to one or more problems that can be solved exactly (in this paper, we use the terminology sub-problems to refer to them). The approximation guarantee for such an approximation algorithm is thus obtained by linking the original problem to those sub-problems. We give a simple abstract example to illustrate this idea. Let's say we want to maximize a function f on some set F , using an α -approximation algorithm. Assume there exist two other functions g and h defined on sets G and H respectively, such that $G \times H \subset F$ and such that g and h can be maximized exactly on G and H respectively. Finally, assume that for all $(x, y) \in G \times H$,

$$\alpha \max_F f - f(x, y) \leq \max_G g - g(x) + \max_H h - h(y). \quad (2)$$

This means that an α -approximation algorithm to maximize f can simply output the feasible solution $(\arg \max_G g, \arg \max_H h) \in F$. This example may seem very basic and artificial at first glance, but it turns out that many approximation algorithms rely on the same principle, as we will see in subsection 4.2. To see how this can be exploited for approximation regret minimization, we can notice that the LHS of (2) is an approximation gap (as defined in Definition 2, taking the max with 0 and considering that the choice $(x, y) \in F$ is that of a policy at a given round, with F playing the role of the action space) and that the RHS is the sum of two gaps, with G (respectively H) playing the role of the action space, but this time without approximation factor. Summing over the rounds, we find that the corresponding approximation regret is bounded by the sum of two "classical" regrets (in this paper, we use the terminology sub-regrets). Thus, the bounds we obtain on these two sub-regrets using CTS with the corresponding exact oracles translate into an bound on the approximation regret.

Avoiding the "mismatch" phenomenon Kong et al. [2021] identified the reason why their regret bound has a Δ^2 in the denominator, while the usual CMAB algorithms only have a Δ . They term it a "mismatch" between the estimated gaps that need to be eliminated by exploration and the actual regret the algorithm needs to pay. We argue that, in fact, this mismatch phenomenon should exist in principle, even for non-approximation regret. It is usually avoided using a smoothness assumption like Assumption 2, linking our estimations (here the arms that generate the feedback) and what is paid. In the above example, although we are in an approximation context, the situation is fundamentally no different. We see that (2) links the paid approximation regret with two exact sub-regrets which are themselves assumed to be related to our estimates through a smoothness-like property. There is thus an indirect link between the approximation regret and the outcomes, which enables one of the Δ present in the denominator of the exploration term $\log(T)/\Delta^2$ to be cancelled out by the actual regret paid, thus avoiding the mismatch phenomenon.

To summarize, just as the approximation relation is obtained by linking the original problem to sub-problems that can be solved exactly, the idea behind REDUCE2EXACT problems is to link the approximation regret to several sub-regrets, each satisfying the appropriate properties for CTS, namely a smoothness relationship and the availability of an exact oracle. We formalize this in the following assumption.

Assumption 4 (REDUCE2EXACT). *There exist $\ell \in \mathbb{N}^*$, $\mathbf{c} \in \mathbb{R}_+^\ell$ and $\mathbf{B}_j \in \mathbb{R}_+^n$ for all $j \in [\ell]$ such that the following is true. Oracle is of the form $\text{Oracle} = \text{Oracle}_2 \circ \text{Oracle}_1$, where Oracle_1 and Oracle_2 are described as follows.*

- Oracle_1 : For $\boldsymbol{\mu} \in \mathcal{M}$, $\text{Oracle}_1(\boldsymbol{\mu})$ must output a sequence (E_1, \dots, E_ℓ) described as follows. For each $j \in [\ell]$, let $\mathcal{E}_j = \mathcal{E}_j(E_1, \dots, E_{j-1})$ be a sub-action space which may depend on E_1, \dots, E_{j-1} and let $r_j(\cdot, \boldsymbol{\mu}) : \mathcal{E}_j \rightarrow \mathbb{R}$ be a reward sub-function. Then, we require that $E_j \in \arg \max_{E \in \mathcal{E}_j} r_j(E, \boldsymbol{\mu})$.
- Oracle_2 : For an input $E_1 \in \mathcal{E}_1, \dots, E_\ell \in \mathcal{E}_\ell$, $\text{Oracle}_2(E_1, \dots, E_\ell)$ must output an action in \mathcal{A} such that:

$$\Delta(\text{Oracle}_2(E_1, \dots, E_\ell)) \leq \sum_{j \in [\ell]} (r_j(E_j^*, \boldsymbol{\mu}^*) - r_j(E_j, \boldsymbol{\mu}^*)) \cdot c_j, \quad (3)$$

where for all $j \in [\ell]$, $E_j^* \in \arg \max_{E \in \mathcal{E}_j} r_j(E, \boldsymbol{\mu}^*)$.

- Finally, in addition to the above constraints on Oracle_1 and Oracle_2 , for each $j \in [\ell]$, we require that the reward sub-function r_j satisfies Assumption 2 with the constants \mathbf{B}_j and with the triggering probabilities $p_i(\text{Oracle}_2(E_1, \dots, E_\ell))$, $i \in [n]$.

Informally, in the above Assumption 4, for $\boldsymbol{\mu} \in \mathcal{M}$, $\text{Oracle}_1(\boldsymbol{\mu})$ exactly solves a finite sequence of recursively defined optimization sub-problems and Oracle_2 builds an action in \mathcal{A} for the original approximation problem using the intermediate solutions provided by Oracle_1 . At first sight, Assumption 4 seems very specific and rather difficult to fulfill. However, we will see in subsection 4.2 that many concrete problems satisfy it.

4.1 Analysis

In this subsection, we give in Theorem 2 the main result of this paper. It basically states that under Assumptions 4, the CTS policy have a regret bound comparable to the exact oracle case. The proof is postponed to Appendix B.

Theorem 2. *If $\Delta_{\min} > 0$ and $p^* > 0$, the policy π described in Algorithm 1 (under Assumption 4) and $\mathbb{P}_{\mathbf{X}} = \otimes_{i \in [n]} \mathbb{P}_{X_i}$, $\mathbf{X} \stackrel{\text{a.s.}}{\in} [0, 1]^n$ or Algorithm 2 (under Assumptions 3 and 4) has regret of order*

$$R_{T,\alpha}(\pi) = \mathcal{O} \left(\sum_{i \in [n]} \frac{\left(\sum_{j \in [\ell]} B_{ij} c_j \right)^2 \log(m) \log(T)}{\Delta_{i,\min}} \right).$$

The idea of the proof is quite simple once Assumption 4 has been made. We can see that the approximation regret can be decomposed into ℓ sub-regrets, according to equation (3). Then, one must focus on the fact that the sub-regrets may dependent on each other and that the right gap (defined with the original reward function) must be obtained in the denominator of the final bound.

4.2 Examples of REDUCE2EXACT problems

Here, we present several problems belonging to REDUCE2EXACT. Each time, after a quick introduction of the problem, we translate it into our CMAB-T context, and finally show how it satisfies the REDUCE2EXACT criteria. The Travelling salesman problem (TSP) is treated in Appendix C.

Submodular maximization (e.g., probabilistic maximum coverage (PMC)) Here, we only expose the PMC example, noting that the same derivation can be applied to a monotone submodular⁵ function. PMC is one of the main examples proposed by Kong et al. [2021]. Given a weighted bipartite graph $G = (L, R, E)$, with weights $\boldsymbol{\mu}^* \triangleq (\mu_{(u,v)}^*)_{(u,v) \in E}$ (notice there are thus $n = |E|$ arms, recalling that for us an arm is only something that produces an outcome, not something we can choose as an action, explaining why arms are not indexed by vertices here), the goal is to find an

⁵ f is monotone if for every $A \subset B$, we have $f(A) \leq f(B)$. It is submodular if for every A, B we have that $f(A \cup B) + f(A \cap B) \leq f(A) + f(B)$.

action $A \in \mathcal{A} \triangleq \{A \subset L : |A| = k\}$, $k \in \mathbb{N}^*$, maximizing the expected number of influenced nodes in R , where each node $v \in R$ can be independently influenced by $u \in A$ with probability $\mu_{(u,v)}^*$, i.e., maximizing $f(A, \boldsymbol{\mu}^*) \triangleq \sum_{v \in R} \left(1 - \prod_{u \in A: (u,v) \in E} (1 - \mu_{(u,v)}^*)\right)$. This problem can be applied to the semi-bandit framework called *the ad placement problem*, where L are the web pages, R are the users and $\mu_{(u,v)}^*$ is the probability that user v clicks on the ad on web page u . In this application, the user's click probabilities are unknown and must be learned as the rounds progress. The Greedy oracle can provide an approximate solution with approximation ratio $\alpha = 1 - 1/e$ [Nemhauser et al., 1978]. This setting fits REDUCE2EXACT as follows:

- Oracle₁($\boldsymbol{\mu}$) : For $i \in [k]$, let $\mathcal{E}_i \triangleq \{(a_1, \dots, a_i) : (a_1, \dots, a_{i-1}) = E_{i-1}, a_i \in L \setminus E_{i-1}\}$ and $r_i((a_1, \dots, a_i), \boldsymbol{\mu}) \triangleq f(\{a_1, \dots, a_i\}, \boldsymbol{\mu})$.
- Oracle₂(E_1, \dots, E_k) : Let $(a_1, \dots, a_k) \triangleq E_k$. Oracle₂ returns $A = \{a_1, \dots, a_k\}$.

Let $A^i = \{a_1, \dots, a_i\}$ for some $i \in [k]$ and by abuse of notation, let $f = f(\cdot, \boldsymbol{\mu}^*)$. Informally, we see that in the above decomposition of the oracle, at each step, Oracle₁ maximizes $f(A^i)$ with A^{i-1} fixed, i.e., Oracle₁ optimizes only on a_i (we thus recover the greedy algorithm). It is precisely these sub-problems of finding a_i that can be solved exactly. Then, we see that Oracle₂ simply returns the last A^i constructed. We will now prove the relation (3). The following is true using that f is monotone submodular (this is actually the way Nemhauser et al. [1978] proved the approximation guarantee, and is true for any monotone submodular function):

$$f(A^*) - f(A^i) \leq \sum_{a \in A^* \setminus A^i} (f(\{a\} \cup A^i) - f(A^{i+1})) + k(f(A^{i+1}) - f(A^i)).$$

Once the above relation is obtained, we can actually continue the original proof from Nemhauser et al. [1978], skipping each step where we would need to use the property of Oracle₁, thus leaving a term in the right-hand side for each time we skipped.

$$\begin{aligned} f(A^*) - f(A^k) &= f(A^*) - f(A^{k-1}) - (f(A^k) - f(A^{k-1})) \\ &\leq (f(A^*) - f(A^{k-1})) \left(1 - \frac{1}{k}\right) + \frac{1}{k} \sum_{a \in A^* \setminus A^{k-1}} (f(\{a\} \cup A^{k-1}) - f(A^k)) \\ \dots &\leq f(A^*) \left(1 - \frac{1}{k}\right)^k + \sum_{i=1}^k \frac{\left(1 - \frac{1}{k}\right)^{i-1}}{k} \sum_{a \in A^* \setminus A^{k-i}} (f(\{a\} \cup A^{k-i}) - f(A^{k-i+1})). \end{aligned}$$

Finally, since $\left(1 - \frac{1}{k}\right)^k \leq e^{-1}$, we get that $\Delta(\text{Oracle}_2(E_1, \dots, E_k)) = (1 - e^{-1})f(A^*) - f(A)$ is bounded by

$$\sum_{i=1}^k \frac{\left(1 - \frac{1}{k}\right)^{i-1}}{k} |A^* \setminus A^{k-i}| (r_{k-i+1}(E_{k-i+1}^*, \boldsymbol{\mu}^*) - r_{k-i+1}(E_{k-i+1}, \boldsymbol{\mu}^*)).$$

Note in passing that we recover the classical approximation if the right-hand-side was equal to 0. It is easy to see that the reward function r_j satisfies Assumption 2 with $\mathbf{B}_j = \mathbf{e}_{[n]}$. We thus finally get our Assumption 4.

Online influence maximization (OIM) As the analysis mentioned above only uses submodularity, it can be extended to the problem of online influence maximization in a social network. A social network is modeled as a directed graph $G = (V, E)$, with nodes V representing users and edges E representing connections. For a node $i \in V$, a subset $A \subset V$, and a vector $\mathbf{x} \in \{0, 1\}^E$, let the predicate $A \overset{\mathbf{x}}{\rightsquigarrow} i$ hold if, in the graph defined by $G_{\mathbf{x}} \triangleq (V, \{ij \in E, x_{ij} = 1\})$, there is a forward path from a node in A to the node i . If it holds, we say that i is influenced by A under \mathbf{x} . The goal is to find an action $A \in \mathcal{A} \triangleq \{A \subset V : |A| = k\}$ maximizing the *influence spread* $\sigma(A, \boldsymbol{\mu}^*) \triangleq \mathbb{E} \left[\left| \left\{ i \in V, A \overset{\mathbf{X}}{\rightsquigarrow} i \right\} \right| \right]$, where $\mathbf{X} \sim \otimes_{(u,v) \in E} \text{Bernoulli}(\mu_{(u,v)}^*)$. This model is called the *independent cascade model* [Kempe et al., 2003, 2015]. A notable property to use the greedy oracle is that σ is monotone submodular. As the exact calculation of σ is prohibitive, it is estimated by simulating the diffusion process, resulting in an approximation factor $\alpha = 1 - e^{-1} - \varepsilon$ in the

above greedy oracle analysis [Kempe et al., 2015, Feige, 1998, Chen et al., 2010], with $\varepsilon > 0$. In OIM, arms may be probabilistically triggered, and Assumption 2 holds with the constants being all equal to $\max_{u \in V} \left| \left\{ v \in V, \{u\} \stackrel{(\mathbb{I}\{\mu_e^* > 0\})_{e \in E}}{\rightsquigarrow} v \right\} \right|$, which is the largest number of nodes any node can reach [Wang and Chen, 2017]. As previously, we thus get Assumption 4 with \mathbf{B}_j being $e_{[n]}$ times this constant.

Metric k -center This example and the following ones are less common for CMAB, but allow to well illustrate REDUCE2EXACT. Given a set of cities, one wants to build k warehouses in different cities and minimize the maximum distance of a city to a warehouse. Formally, given a complete undirected weighted graph $G = (V, E)$ whose distances $d(v_i, v_j)$ satisfy the triangle inequality, the goal is to find an action $A \in \mathcal{A} \triangleq \{A \subset V : |A| = k\}$ that minimizes $\max_{v \in V} d(v, A)$. We can consider the semi-bandit setting where the set of base arms is E , $\boldsymbol{\mu}^* \triangleq (d(v_i, v_j))_{(v_i, v_j) \in E}$ and the feedback set S includes the edges of the graph induced by the chosen action. We can target an approximation regret with $\alpha = 1/2$ using the following oracle (which is simply the standard greedy algorithm for this problem).

- Oracle₁($\boldsymbol{\mu}$) : For $i \in [k]$, let $\mathcal{E}_i \triangleq \{(a_1, \dots, a_i) : (a_1, \dots, a_{i-1}) = E_{i-1}, a_i \in V \setminus E_{i-1}\}$ and $r_i((a_1, \dots, a_i), \boldsymbol{\mu}) \triangleq \min_{j \in [i-1]} \mu_{a_i, a_j}$ (with $r_1 = 0$). One can notice we thus have $E_i^* \in \arg \max_{(a_1, \dots, a_i) \in \mathcal{E}_i} d(a_i, \{a_1, \dots, a_{i-1}\}) = \arg \max_{(a_1, \dots, a_i) \in \mathcal{E}_i} d(a_i, E_{i-1})$.
- Oracle₂(E_1, \dots, E_k) : Let $(a_1, \dots, a_k) \triangleq E_k$. Oracle₂ returns $A = \{a_1, \dots, a_k\}$. Let $w \in \arg \max_{v \in V} d(v, A)$.

We thus have:

$$\begin{aligned} \frac{1}{2} \max_{v \in V} d(v, A) - \max_{v \in V} d(v, A^*) &\leq \frac{1}{2} \left(\max_{v \in V} d(v, A) - \min_{a \in A \cup \{w\}} \min_{a' \in A \setminus \{a\}} d(a, a') \right) \\ &= \frac{1}{2} \max_{j \in [k-1]} \left(\max_{v \in V} d(v, A) - d(a_{j+1}, E_j) \right) \\ &\leq \frac{1}{2} \max_{j \in [k-1]} \left(\max_{v \in V} d(v, E_j) - d(a_{j+1}, E_j) \right) \\ &= \frac{1}{2} \max_{j \in [k-1]} (r_{j+1}(E_{j+1}^*, \boldsymbol{\mu}^*) - r_{j+1}(E_{j+1}, \boldsymbol{\mu}^*)). \end{aligned}$$

Where the first inequality is deduced as follows: the map $f : v \mapsto \arg \min_{a^* \in A^*} d(a^*, v)$ defines a partition of V into $k = |A^*|$ clusters. By the pigeonhole principle, one cluster contains 2 different points $a, a' \in A \cup \{w\}$ (simply because its size is $k + 1$). We can assume $a \in A$ without loss of generality. Thus, since $f(a) = f(a')$, we get $d(a, a') \leq d(a, f(a)) + d(a', f(a')) = d(a, A^*) + d(a', A^*) \leq 2 \max_{v \in V} d(v, A^*)$.

Vertex cover The problem consists, given an undirected graph $G = (V, E)$, in finding a set of vertices with minimal cost to cover all the edges of E . Formally, with $\boldsymbol{\mu}^* \in \mathbb{R}_+^V$, the goal is to find an action $A \in \mathcal{A} \triangleq \{A \subset V : \forall (u, v) \in E, u \in A \text{ or } v \in A\}$ that minimizes $\mathbf{e}_A^\top \boldsymbol{\mu}^*$. The semi-bandit feedback is defined directly as $S = A$. We can target an approximation regret with $\alpha = 1/2$ using the following linear programming (LP) relaxation oracle.

- Oracle₁($\boldsymbol{\mu}$) : Let⁶ $\mathcal{E}_1 \triangleq \{\mathbf{x} \in \{0, 1/2, 1\}^V : \forall (u, v) \in E, x_u + x_v \geq 1\}$ and $r_1(\mathbf{x}, \boldsymbol{\mu}) \triangleq -\mathbf{x}^\top \boldsymbol{\mu}$.
- Oracle₂(E_1) : Let $\mathbf{x} \triangleq E_1$. Oracle₂ returns $A = \{v \in V : x_v \geq 1/2\} \in \mathcal{A}$ (since $\forall (u, v) \in E, x_u + x_v \geq 1$, so $x_u \geq 1/2$ or $x_v \geq 1/2$, so $u \in A$ or $v \in A$). Let $\mathbf{x}^* \triangleq E_1^*$.

We have:

$$\frac{1}{2} \mathbf{e}_A^\top \boldsymbol{\mu}^* - \mathbf{e}_A^\top \boldsymbol{\mu}^* \leq \mathbf{x}^\top \boldsymbol{\mu}^* - \mathbf{e}_A^\top \boldsymbol{\mu}^* \leq \mathbf{x}^\top \boldsymbol{\mu}^* - \mathbf{x}^{*\top} \boldsymbol{\mu}^* = r_1(E_1^*, \boldsymbol{\mu}^*) - r_1(E_1, \boldsymbol{\mu}^*).$$

⁶The LP relaxation of vertex cover is half-integral, so that we can allow each variable to be in $\{0, 1/2, 1\}$ rather than the interval from 0 to 1.

Notice, r_1 satisfies Assumption 2 with the constants being 1, using $\mathbf{x} \in \{0, 1/2, 1\}^V$, so that $\mathbf{x} \leq \mathbf{e}_A$.

Max-Cut Given an undirected weighted graph $G = (V, E)$, with weights $\boldsymbol{\mu}^* \triangleq (\mu_{(u,v)}^*)_{(u,v) \in E}$, the goal is to find an action $A \in \mathcal{A} \triangleq \{A \subset V\}$ maximizing the total weight of the edges between A and its complement, i.e., $\frac{1}{2} \sum_{(u,v) \in E} \mu_{(u,v)}^* (1 - y_u y_v)$, where $\mathbf{y} \triangleq \mathbf{e}_A - \mathbf{e}_{V \setminus A}$. We consider the semi-bandit context where E is the set of arms, and the feedback set includes the edges between A and its complement, i.e., $S = \{(u, v) \in E, y_u y_v = -1\}$. To include randomization within the oracle, we can extend the action space \mathcal{A} to the set of probability measures on $\{A \subset V\}$, replacing $\Delta(A)$ by its expectation on A , as a function of the distribution of A . The polynomial-time approximation algorithm for Max-Cut with the best known approximation ratio [Goemans and Williamson, 1995] uses semidefinite programming and randomized rounding, and achieves an approximation ratio of $\alpha = \frac{2}{\pi} \min_{0 \leq \theta \leq \pi} \frac{\theta}{1 - \cos \theta} \approx 0.878$. In our context, it can be defined as follows.

- Oracle₁($\boldsymbol{\mu}$) : We define $\mathcal{E}_1 \triangleq \{(\mathbf{v}_u)_{u \in V} \in \{\mathbf{v} \in \mathbb{R}^V : \|\mathbf{v}\|_2 = 1\}^V\}$ and $r_1((\mathbf{v}_u)_{u \in V}, \boldsymbol{\mu}) \triangleq \frac{1}{2} \sum_{(u,v) \in E} \mu_{(u,v)} (1 - \mathbf{v}_u^\top \mathbf{v}_v)$.
- Oracle₂(E_1) : Let $(\mathbf{v}_u)_{u \in V} \triangleq E_1$. Oracle₂ returns the distribution of $A = \{u \in V, \mathbf{v}_u^\top \mathbf{Z} \geq 0\}$, where $\mathbf{Z} \sim \mathcal{U}(\{\mathbf{v} \in \mathbb{R}^V : \|\mathbf{v}\|_2 = 1\})$.

The following is proved by Goemans and Williamson [1995]:

$$\begin{aligned} \sum_{(u,v) \in E} \mu_{(u,v)} \mathbb{P}[(u, v) \in S] &= \mathbb{E} \left[\sum_{(u,v) \in E} \mu_{(u,v)} \frac{1 - \text{sign}(\mathbf{v}_u^\top \mathbf{Z}) \text{sign}(\mathbf{v}_v^\top \mathbf{Z})}{2} \right] \\ &= \sum_{(u,v) \in E} \mu_{(u,v)} \frac{\arccos(\mathbf{v}_u^\top \mathbf{v}_v)}{\pi} \\ &\geq \alpha \sum_{(u,v) \in E} \mu_{(u,v)} \frac{1 - \mathbf{v}_u^\top \mathbf{v}_v}{2} = \alpha r_1((\mathbf{v}_u)_{u \in V}, \boldsymbol{\mu}). \end{aligned}$$

Thus, r_1 satisfies Assumption 2 with the constants being $1/\alpha$. We also get, with $\mathbf{y}^* \triangleq \mathbf{e}_{A^*} - \mathbf{e}_{V \setminus A^*}$:

$$\begin{aligned} \Delta(\text{Oracle}_2(E_1)) &= \alpha \sum_{(u,v) \in E} \mu_{(u,v)}^* \frac{1 - y_u^* y_v^*}{2} - \mathbb{E} \left[\sum_{(u,v) \in E} \mu_{(u,v)}^* \frac{1 - \text{sign}(\mathbf{v}_u^\top \mathbf{Z}) \text{sign}(\mathbf{v}_v^\top \mathbf{Z})}{2} \right] \\ &\leq \alpha (r_1(E_1^*, \boldsymbol{\mu}^*) - r_1(E_1, \boldsymbol{\mu}^*)). \end{aligned}$$

5 Conclusion

In this article, our main objective is to further expand the "approximation regret scope" of the CTS policy. We not only expand it to probabilistically triggered arms (which was one of the open questions by Kong et al. [2021]), but we also consider a broader class of oracles compatible with CTS. More precisely, we propose a condition, REDUCE2EXACT, which may seem unnatural at first, but which in fact simply expresses that sub-problems that can be solved exactly must be hidden in the original approximation problem, and that the approximation oracle exploit them to output the final solution. Knowing that the majority of approximation algorithms use one or more relaxations to an exact problem (e.g., solving a convex programming relaxation to obtain a fractional solution and then rounding this fractional solution to get a feasible solution), our assumption falls within the range of many CMAB-T settings. From this reduction, we naturally obtain the standard tight regret bound $\mathcal{O}(\log(T)/\Delta_{\min})$. This is the first tight bound for the approximation regret on non-exact oracles.

As future work, it may be interesting to explore other CMAB-T problems where the REDUCE2EXACT condition does (or doesn't) hold. We also think our setting should be generalizable to the *budgeted regret* setting without much difficulty (see Perrault et al. [2019b, 2020c,b] for examples with an approximation oracle). Finally, we have that the approximation regret is in some way conservative compared to the greedy regret. For a given oracle, we can easily consider the equivalent of the greedy regret for that oracle. An interesting investigation would then be to extend the work of Kong et al. [2021] in this direction, considering other types of oracle.

References

- S. Agrawal and N. Goyal. Thompson Sampling for Contextual Bandits with Linear Payoffs. *CoRR*, abs/1209.3352, <http://arxiv.org/abs/1209.3352>, sep 2012. URL <http://arxiv.org/abs/1209.3352>.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- D. A. Berry and B. Fristedt. *Bandit Problems: Sequential Allocation of Experiments*, volume 38 of *Monographs on statistics and applied probability*. Chapman and Hall, 1985.
- N. Cesa-Bianchi and G. Lugosi. Combinatorial bandits. In *Journal of Computer and System Sciences*, volume 78, pages 1404–1422, 2012.
- S.-H. Chang, P. C. Cosman, and L. B. Milstein. Chernoff-type bounds for the gaussian error function. *IEEE Transactions on Communications*, 59(11):2939–2944, 2011.
- W. Chen, C. Wang, and Y. Wang. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Knowledge Discovery and Data Mining*, 2010.
- W. Chen, Y. Wang, and Y. Yuan. Combinatorial multi-armed bandit: General framework and applications. In S. Dasgupta and D. McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 151–159, Atlanta, Georgia, USA, 2013. PMLR. URL <http://proceedings.mlr.press/v28/chen13a.html>.
- W. Chen, Y. Wang, and Y. Yuan. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *Journal of Machine Learning Research*, 17, 2016.
- N. Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.
- R. Combes, M. S. Talebi Mazraeh Shahi, A. Proutiere, and m. Ielarge. Combinatorial bandits revisited. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2116–2124. Curran Associates, Inc., 2015. URL <http://papers.nips.cc/paper/5831-combinatorial-bandits-revisited.pdf>.
- R. Degenne and V. Perchet. Combinatorial semi-bandit with known covariance. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2972–2980. Curran Associates, Inc., 2016. URL <http://papers.nips.cc/paper/6137-combinatorial-semi-bandit-with-known-covariance.pdf>.
- R. Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.
- U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)*, 45(4): 634–652, 1998.
- Y. Gai, B. Krishnamachari, and R. Jain. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *Transactions on Networking*, 20(5):1466–1478, 2012.
- M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6): 1115–1145, 1995.
- A. Gopalan, S. Mannor, and Y. Mansour. Thompson sampling for complex bandit problems. In *International Conference on Machine Learning*, 2014.
- W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.
- A. Huyuk and C. Tekin. Analysis of Thompson Sampling for combinatorial multi-armed bandit with probabilistically triggered arms. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1322–1330. PMLR, 2019.
- I. M. Jacobs and J. Wozencraft. Principles of communication engineering. 1965.
- S. M. Kakade, A. T. Kalai, and K. Ligett. Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106, 2009.

- E. Kaufmann, N. Korda, and R. Munos. Thompson Sampling: An Asymptotically Optimal Finite Time Analysis. *Algorithmic Learning Theory*, 2012.
- D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.
- D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. *Theory of Computing*, 11(4):105–147, 2015.
- J. Komiyama, J. Honda, and H. Nakagawa. Optimal Regret Analysis of Thompson Sampling in Stochastic Multi-armed Bandit Problem with Multiple Plays. jun 2015. URL <http://arxiv.org/abs/1506.00779>.
- F. Kong, Y. Yang, W. Chen, and S. Li. The Hardness Analysis of Thompson Sampling for Combinatorial Semi-bandits with Greedy Oracle, 2021.
- B. Kveton, Z. Wen, A. Ashkan, and C. Szepesvari. Tight regret bounds for stochastic combinatorial semi-bandits. In *International Conference on Artificial Intelligence and Statistics*, 2015.
- T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- O. Marchal, J. Arbel, et al. On the sub-gaussianity of the beta and dirichlet distributions. *Electronic Communications in Probability*, 22, 2017.
- G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14(1):265–294, 1978.
- P. Perrault. *Efficient Learning in Stochastic Combinatorial Semi-Bandits*. Theses, Univeristé Paris-Saclay, Nov. 2020. URL <https://tel.archives-ouvertes.fr/tel-03093268>.
- P. Perrault, V. Perchet, and M. Valko. Exploiting structure of uncertainty for efficient matroid semi-bandits. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5123–5132, Long Beach, California, USA, 2019a. PMLR. URL <http://proceedings.mlr.press/v97/perrault19a.html>.
- P. Perrault, V. Perchet, and M. Valko. Finding the bandit in a graph: Sequential search-and-stop. In K. Chaudhuri and M. Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 1668–1677. PMLR, 2019b. URL <http://proceedings.mlr.press/v89/perrault19a.html>.
- P. Perrault, E. Boursier, V. Perchet, and M. Valko. Statistical Efficiency of Thompson Sampling for Combinatorial Semi-Bandits. *arXiv preprint arXiv:2006.06613*, 2020a.
- P. Perrault, J. Healey, Z. Wen, and M. Valko. On the approximation relationship between optimizing ratio of submodular (rs) and difference of submodular (ds) functions, 2020b.
- P. Perrault, J. Healey, Z. Wen, and M. Valko. Budgeted online influence maximization. In *Proceedings of the 37th International Conference on Machine Learning*, pages 6588–6599. 2020c.
- P. Perrault, V. Perchet, and M. Valko. Covariance-adapting algorithm for semi-bandits with application to sparse rewards. In *Conference on Learning Theory*, 2020d.
- H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematics Society*, 58:527–535, 1952.
- D. Russo and B. Van Roy. An information-theoretic analysis of thompson sampling. *The Journal of Machine Learning Research*, 17(1):2442–2471, 2016.
- M. Streeter and D. Golovin. An online algorithm for maximizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 1577–1584, 2009.
- W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294, 1933.
- Q. Wang and W. Chen. Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. In *Neural Information Processing Systems*, 2017. URL

<http://arxiv.org/abs/1703.01610>.

- S. Wang and W. Chen. Thompson Sampling for Combinatorial Semi-Bandits. mar 2018. URL <http://arxiv.org/abs/1803.04623>.
- Z. Wen, B. Kveton, and A. Ashkan. Efficient learning in large-scale combinatorial semi-bandits. In *International Conference on Machine Learning*, pages 1113–1122, 2015.
- Z. Wen, B. Kveton, and M. Valko. Influence maximization with semi-bandit feedback. Technical report, 2016.
- Z. Wen, B. Kveton, M. Valko, and S. Vaswani. Online influence maximization under independent cascade model with semi-bandit feedback. In *Neural Information Processing Systems*, 2017.
- R. Zhang and R. Combes. On the Suboptimality of Thompson Sampling in High Dimensions. *arXiv preprint arXiv:2102.05502*, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] See Abstract and section 1.
 - (b) Did you describe the limitations of your work? [Yes] See the paragraph "Other known limitations of CTS" and the second paragraph in section 5.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] This work does not have any potential negative societal impacts.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See section 2.
 - (b) Did you include complete proofs of all theoretical results? [Yes] See all sections in Appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Proof of Theorem 1

In the following proof, we will treat both algorithms at the same time, detailing the steps where there is a difference. Thus, for Algorithm 1, we use the convention $\beta = 1$ in the following. We let $B \triangleq \|\mathbf{B}\|_\infty$ and $0 < \varepsilon < \Delta_{\min}/(2B(m^{*2} + 1))$. For the two algorithms, we consider the following events for any time step $t \in \mathbb{N}^*$:

- $\mathfrak{Z}_t \triangleq \{\Delta_t > 0\}$,
- $\mathfrak{B}_t \triangleq \{\|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\bar{\boldsymbol{\mu}}_{t-1} - \boldsymbol{\mu}^*)\|_1 > \Delta_{\min}/2 - B(m^{*2} + 1)\varepsilon\}$,
- $\mathfrak{C}_t \triangleq \{\|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \boldsymbol{\mu}^*)\|_1 > \Delta_t - B(m^{*2} + 1)\varepsilon\}$,
- $\mathfrak{D}_t \triangleq \left\{ \|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1 \geq \sqrt{\sum_{i \in [n]} \frac{2 \log(2^n (1 + \lceil \log_2(1/p^*) \rceil)^n T) p_i(A_t)^2 B_i^2 \beta}{N_{i,t-1}}} \right\}$.

We decompose the regret analysis into several steps, each step corresponding to a filtration of the regret against a combination of these events.

Step 1: bound under $\mathfrak{Z}_t \wedge \neg \mathfrak{C}_t$ The filtered regret bound

$$\mathbb{E} \left[\sum_{t=1}^T \Delta_t \mathbb{I}\{\mathfrak{Z}_t \wedge \neg \mathfrak{C}_t\} \right] \leq \Delta_{\max} \frac{cm^*}{p^* \varepsilon^2} \left(\frac{c'}{\varepsilon^4} \right)^{m^*}$$

is deduced from the following two lemmas, considering the following events for a subset $Z \subset [n]$:

$$\mathfrak{R}(\boldsymbol{\theta}', Z) \triangleq$$

$$\left\{ Z \subset \mathsf{T}(A) \text{ s.t. } A = \text{Oracle}(\boldsymbol{\theta}'), \|\mathbf{p}(A) \odot \mathbf{B} \odot (\boldsymbol{\theta}' - \boldsymbol{\mu}^*)\|_1 > \Delta(A) - B(m^{*2} + 1)\varepsilon \right\},$$

$$\mathfrak{S}_t(Z) \triangleq \left\{ \forall \boldsymbol{\theta}' \text{ s.t. } \|(\boldsymbol{\mu}^* - \boldsymbol{\theta}') \odot \mathbf{e}_Z\|_\infty \leq \varepsilon, \mathfrak{R}(\boldsymbol{\theta}' \odot \mathbf{e}_Z + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z^c}, Z) \text{ holds} \right\},$$

$$\mathfrak{T}_t(Z) \triangleq \left\{ \|(\boldsymbol{\mu}^* - \boldsymbol{\theta}_t) \odot \mathbf{e}_Z\|_\infty > \varepsilon \right\}.$$

Lemma 1.

$$\mathfrak{Z}_t, \neg \mathfrak{C}_t \Rightarrow \exists Z \subset \mathsf{T}(A^*), Z \neq \emptyset \text{ s.t. the event } \mathfrak{S}_t(Z) \wedge \mathfrak{T}_t(Z) \text{ holds.}$$

Lemma 2. *There are two constants c, c' such that*

$$\sum_{Z \subset \mathsf{T}(A^*), Z \neq \emptyset} \mathbb{E} \left[\sum_{t=1}^T \mathbb{I}\{\mathfrak{S}_t(Z), \mathfrak{T}_t(Z)\} \right] \leq \frac{cm^*}{p^* \varepsilon^2} \left(\frac{c'}{\varepsilon^4} \right)^{m^*}.$$

Proof of Lemma 1. It is sufficient to prove that

$$\mathfrak{Z}_t, \neg \mathfrak{C}_t \Rightarrow \exists Z \subset \mathsf{T}(A^*), Z \neq \emptyset \text{ s.t. } \mathfrak{S}_t(Z) \text{ holds,} \quad (4)$$

because $\neg \mathfrak{C}_t$ and $\mathfrak{S}_t(Z)$ together imply $\mathfrak{T}_t(Z)$. Then, we get (4) in a similar way as Lemma 2 from Huyuk and Tekin [2019], which is possible as their Assumption 3 is implied by our Assumption 2. More precisely, the only places where we use our Assumption 2 instead of their Assumption 3 are in cases 1a, 2a, ..., when we show that $\mathfrak{R}(\boldsymbol{\theta}' \odot \mathbf{e}_Z + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z^c}, Z)$ holds. The detailed proof is given in the following.

We first consider the choice $Z = Z_1 = \mathsf{T}(A^*)$. Two cases can be distinguished:

$$1a) \forall \boldsymbol{\theta}' \text{ s.t. } \|(\boldsymbol{\mu}^* - \boldsymbol{\theta}') \odot \mathbf{e}_{Z_1}\|_\infty \leq \varepsilon, \text{ we have } Z_1 \subset \mathsf{T}(\text{Oracle}(\boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c})).$$

$$1b) \exists \boldsymbol{\theta}' \text{ s.t. } \|(\boldsymbol{\mu}^* - \boldsymbol{\theta}') \odot \mathbf{e}_{Z_1}\|_\infty \leq \varepsilon \text{ such that } Z_1 \not\subset \mathsf{T}(\text{Oracle}(\boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c})).$$

1a) For the first case, consider any vector $\boldsymbol{\theta}'$ such that $\|(\boldsymbol{\mu}^* - \boldsymbol{\theta}') \odot \mathbf{e}_{Z_1}\|_\infty \stackrel{(5)}{\leq} \varepsilon$ and let $A \stackrel{(6)}{=} \text{Oracle}(\boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c})$. We can write

$$r(A, \boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}) \stackrel{(7)}{\geq} r(A^*, \boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}) \stackrel{(8)}{\geq} r(A^*, \boldsymbol{\mu}^*) - Bm^* \varepsilon,$$

where (7) is from (6), and (8) is from (5). This rewrites as

$$r(A, \boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}) \geq r(A^*, \boldsymbol{\mu}^*) - Bm^*\varepsilon > r(A^*, \boldsymbol{\mu}^*) - B(m^{*2} + 1)\varepsilon,$$

so $\mathfrak{R}_t(\boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}, Z_1)$ holds. Therefore, we have proved that $\mathfrak{S}_t(Z_1)$ holds.

1b) For the second case, we have some vector $\boldsymbol{\theta}'$ such that $\|(\boldsymbol{\mu}^* - \boldsymbol{\theta}') \odot \mathbf{e}_{Z_1}\|_\infty \stackrel{(9)}{\leq} \varepsilon$, and some action $A = \text{Oracle}(\boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c})$ such that $Z_1 \not\subset \mathsf{T}(A)$. We consider $Z_2 = Z_1 \cap \mathsf{T}(A) \neq Z_1$. We first prove that $Z_2 \neq \emptyset$ by showing that if an action A' is such that $Z_1 \cap \mathsf{T}(A') \stackrel{(10)}{=} \emptyset$, then $A' \neq A$ as it has a lower reward value than that of A^* :

$$\begin{aligned} r(A', \boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}) &\stackrel{(11)}{=} r(A', \boldsymbol{\theta}_t) \stackrel{(12)}{\leq} r(A_t, \boldsymbol{\theta}_t) \\ &\stackrel{(13)}{\leq} r(A^*, \boldsymbol{\mu}^*) - B(m^{*2} + 1)\varepsilon \\ &< r(A^*, \boldsymbol{\mu}^*) - Bm^*\varepsilon \\ &\stackrel{(14)}{\leq} r(A^*, \boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}), \end{aligned}$$

where (11) is from (10), (12) is from the definition of A_t , (13) is from $\neg \mathfrak{C}_t$ and (14) is from (9). Now, we again distinguish two cases:

- 2a) $\forall \boldsymbol{\theta}''$ s.t. $\|(\boldsymbol{\mu}^* - \boldsymbol{\theta}'') \odot \mathbf{e}_{Z_2}\|_\infty \leq \varepsilon$, we have $Z_2 \subset \mathsf{T}(\text{Oracle}(\boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}))$.
- 2b) $\exists \boldsymbol{\theta}''$ s.t. $\|(\boldsymbol{\mu}^* - \boldsymbol{\theta}'') \odot \mathbf{e}_{Z_2}\|_\infty \leq \varepsilon$ such that $Z_2 \not\subset \mathsf{T}(\text{Oracle}(\boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}))$.

Notice that when $\|(\boldsymbol{\mu}^* - \boldsymbol{\theta}'') \odot \mathbf{e}_{Z_2}\|_\infty \stackrel{(15)}{\leq} \varepsilon$, then

$$r(A, \boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}) \geq r(A, \boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}) - 2B(m^* - 1)\varepsilon. \quad (16)$$

Indeed, (16) is a consequence of

$$\begin{aligned} &\|(\boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c} - \boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} - \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}) \odot \mathbf{e}_{\mathsf{T}(A)}\|_1 \\ &= \|(\boldsymbol{\theta}' - \boldsymbol{\theta}'') \odot \mathbf{e}_{Z_2}\|_1 \\ &\leq \|(\boldsymbol{\mu}^* - \boldsymbol{\theta}') \odot \mathbf{e}_{Z_2}\|_1 + \|(\boldsymbol{\mu}^* - \boldsymbol{\theta}'') \odot \mathbf{e}_{Z_2}\|_1 \\ &\leq 2(m^* - 1)\varepsilon, \end{aligned}$$

where we used (15), (9) and that Z_2 is strictly included in Z_1 .

2a) For the first case, considering any vector $\boldsymbol{\theta}''$ such that $\|(\boldsymbol{\mu}^* - \boldsymbol{\theta}'') \odot \mathbf{e}_{Z_2}\|_\infty \leq \varepsilon$, we have with $\tilde{A} = \text{Oracle}(\boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c})$ that

$$\begin{aligned} r(\tilde{A}, \boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}) &\geq r(A, \boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}) \\ &\stackrel{(17)}{\geq} r(A, \boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}) - 2B(m^* - 1)\varepsilon \\ &\geq r(A^*, \boldsymbol{\theta}' \odot \mathbf{e}_{Z_1} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_1^c}) - 2B(m^* - 1)\varepsilon \\ &\stackrel{(18)}{\geq} r(A^*, \boldsymbol{\mu}^*) - Bm^*\varepsilon - 2B(m^* - 1)\varepsilon \\ &\geq r(A^*, \boldsymbol{\mu}^*) - B(m^{*2} + 1)\varepsilon, \end{aligned}$$

where (17) uses (16) and (18) uses (9). Therefore, $\mathfrak{R}_t(\boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}, Z_2)$ holds, and we proved that $\mathfrak{S}_t(Z_2)$ holds.

2b) For the second case, we have a vector $\boldsymbol{\theta}''$ such that $\|(\boldsymbol{\mu}^* - \boldsymbol{\theta}'') \odot \mathbf{e}_{Z_2}\|_\infty \leq \varepsilon$ and an action $\tilde{A} = \text{Oracle}(\boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c})$ such that $Z_2 \not\subset \mathsf{T}(\tilde{A})$. We consider $Z_3 =$

$Z_2 \cap \mathsf{T}(\tilde{A})$. Again, $Z_3 \neq \emptyset$ because for any A'' such that $\mathsf{T}(A'') \cap Z_2 = \emptyset$, we have $A'' \neq \text{Oracle}(\boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c})$:

$$\begin{aligned} r(A'', \boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}) &= r(A'', \boldsymbol{\theta}_t) \leq r(A_t, \boldsymbol{\theta}_t) \\ &\leq r(A^*, \boldsymbol{\mu}^*) - B(m^{*2} + 1)\varepsilon \\ &< r(A^*, \boldsymbol{\mu}^*) - Bm^*\varepsilon - 2B(m^* - 1)\varepsilon \\ &\leq r(A, \boldsymbol{\theta}'' \odot \mathbf{e}_{Z_2} + \boldsymbol{\theta}_t \odot \mathbf{e}_{Z_2^c}), \end{aligned}$$

where the last inequality is obtained in the same way as in inequalities from (17) to (18).

We could repeat the above argument and each time the size Z_i is decreased by at least 1. Thus, after at most $m^* - 1$ steps, since $m^* + 2(m^* - 1) + 2(m^* - 2) + \dots + 2 = m^{*2} < m^{*2} + 1$, we could reach the end and find a $Z_i \neq \emptyset$ such that $\mathfrak{S}_t(Z_i)$ holds. \square

Proof of Lemma 2. Let $Z = \{z_0, \dots, z_{|Z|-1}\} \subset \mathsf{T}(A^*)$, $Z \neq \emptyset$. If $|Z| = 1$, we let

$$\eta_{q,0} \triangleq \{t \geq 1, |\{t' \in [t-1], \mathfrak{S}_{t'}(Z) \wedge \neg \mathfrak{I}_{t'}(Z) \wedge \{z_0 \in S_{t'}\}\}| = q\}.$$

Else, for $t > 1$, we recursively define

$$c_{t+1} \triangleq c_t + \mathbb{I}\{\mathfrak{S}_t(Z) \wedge \neg \mathfrak{I}_t(Z) \wedge \{z_{c_t} \in S_t\}\} \pmod{|Z|},$$

with $c_1 \triangleq 0$ and let

$$\eta_{q,k} \triangleq \{t \geq 1, c_t = k, |t' \in [t-1], c_{t'} = k \neq c_{t'+1}| = q\}.$$

Notice that for $\tau \geq \inf \eta_{q,0}$, we have $N_{i,\tau-1} \geq q$ for all $i \in Z$. We have

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \mathbb{I}\{\mathfrak{S}_t(Z), \mathfrak{I}_t(Z)\} \right] &= \sum_{q \geq 0} \sum_{k=0}^{|Z|-1} \mathbb{E} \left[\sum_{t \in \eta_{q,k}} \mathbb{I}\{\mathfrak{S}_t(Z), \mathfrak{I}_t(Z)\} \right] \\ &\leq \frac{1}{p^*} \sum_{q \geq 0} \sum_{k=0}^{|Z|-1} \mathbb{E} \left[\sum_{t \in \eta_{q,k}} \mathbb{I}\{\mathfrak{S}_t(Z), \mathfrak{I}_t(Z), z_{c_t} \in S_t\} \right] \\ &\leq \frac{1}{p^*} \sum_{q \geq 0} \sum_{k=0}^{|Z|-1} \left(\mathbb{E} \left[\sup_{\tau \geq \inf \eta_{q,k}} \frac{1}{\mathbb{P}[\neg \mathfrak{I}_\tau(Z) | \mathcal{H}_\tau]} \right] - 1 \right) \\ &\leq \frac{|Z|}{p^*} \sum_{q \geq 0} \left(\mathbb{E} \left[\sup_{\tau \geq \inf \eta_{q,0}} \prod_{i \in Z} \frac{1}{\mathbb{P}[|\theta_{i,\tau} - \mu_i^*| \leq \varepsilon | \mathcal{H}_\tau]} \right] - 1 \right). \end{aligned}$$

From the initialization phase, we can assume that the event

$$\mathfrak{M}_t \triangleq \{\forall i \in [n], N_{i,t-1} \geq 1\}$$

holds (under the complementary event, we have the upper bound n). If there is no initialization, we can have $q = 0$ in the following, noticing that when $\theta_{i,t}$ is uniform on $[a, b]$, then the probability $\mathbb{P}[|\theta_{i,t} - \mu_i^*| \leq \varepsilon | \mathcal{H}_t]$ is equal to $2\varepsilon/(b-a)$. We are thus interested in bounding

$$\frac{|Z|}{p^*} \underbrace{\sum_{q \geq 1} \left(\mathbb{E} \left[\sup_{\tau \geq \tau_q} \prod_{i \in Z} \frac{1}{\mathbb{P}[|\theta_{i,\tau} - \mu_i^*| \leq \varepsilon | \mathcal{H}_\tau]} \right] - 1 \right)}_{(19)},$$

with $\tau_q \triangleq \inf \eta_{q,0}$. We have

$$(19) \leq \sum_{q \geq 1} \mathbb{E} \left[\sup_{\tau \geq \tau_q} \sum_{Z' \subset Z, Z' \neq \emptyset} \prod_{i \in Z'} \left(\frac{1}{\mathbb{P}[|\theta_{i,\tau} - \mu_i^*| \leq \varepsilon | \mathcal{H}_\tau]} - 1 \right) \right]$$

$$\leq \sum_{q \geq 1} \sum_{Z' \subset Z, Z' \neq \emptyset} \mathbb{E} \left[\underbrace{\sup_{\tau \geq \tau_q} \prod_{i \in Z'} \left(\frac{1}{\mathbb{P}[|\theta_{i,\tau} - \mu_i^*| \leq \varepsilon | \mathcal{H}_\tau]} - 1 \right)}_{(20)} \right]$$

Then, we can take a union bound on the counters:

$$(20) \leq \sum_{\mathbf{k} \in [q.. \infty]^{Z'}} \mathbb{E} \left[\sup_{\tau \geq \tau_q} \mathbb{I}\{\forall i \in Z', N_{i,\tau-1} = k_i\} \prod_{i \in Z'} \left(\frac{1}{\mathbb{P}[|\theta_{i,\tau} - \mu_i^*| \leq \varepsilon | \mathcal{H}_\tau]} - 1 \right) \right].$$

From this point, there are two distinct analysis depending on whether we consider Algorithm 1 or Algorithm 2.

For Algorithm 1:

For any arm $i \in [n]$, $k_i \in \mathbb{N}$, we define p_{i,k_i} as the probability of $|\tilde{\theta}_{i,k_i} - \mu_i^*| \leq \varepsilon$, where $\tilde{\theta}_{i,k_i}$ is a sample from the posterior of arm i when there are k_i observations of arm i (i.e., p_{i,k_i} is a random variable measurable with respect to those k_i independent draws of arm i). We have

$$\begin{aligned} \mathbb{E} \left[\sup_{\tau \geq \tau_q} \mathbb{I}\{\forall i \in Z', N_{i,\tau-1} = k_i\} \prod_{i \in Z'} \left(\frac{1}{\mathbb{P}[|\theta_{i,\tau} - \mu_i^*| \leq \varepsilon | \mathcal{H}_\tau]} - 1 \right) \right] &= \mathbb{E} \left[\prod_{i \in Z'} \left(\frac{1}{p_{i,k_i}} - 1 \right) \right], \\ &= \prod_{i \in Z'} \mathbb{E} \left[\left(\frac{1}{p_{i,k_i}} - 1 \right) \right]. \end{aligned}$$

From Lemma 5,6 in Wang and Chen [2018], we know that

$$\mathbb{E} \left[\frac{1}{p_{i,k_i}} \right] \leq \begin{cases} 4/\varepsilon^2 & \text{for every } k_i \geq 0 \\ 1 + 6c'' \cdot e^{-\varepsilon^2 k_i/2} \varepsilon^{-2} + \frac{2}{e^{\varepsilon^2 k_i/8} - 2} & \text{if } k_i > 8/\varepsilon^2, \end{cases}$$

for some universal constant c'' . There are thus two cases: If $q > 8/\varepsilon^2$, then some simple calculations show that $\sum_{Z' \subset Z, Z' \neq \emptyset} (20)$ is bounded by a term of the form $e^{-\varepsilon^2 q/8} (c' \varepsilon^{-4})^{|Z'|}$, where c' is a universal constant, and if $q \leq 8/\varepsilon^2$, then $\sum_{Z' \subset Z, Z' \neq \emptyset} (20)$ is bounded by $(c \varepsilon^{-4})^{|Z'|}$, where c is a universal constant. Summing over $q \geq 1$, we thus get the desired result.

For Algorithm 2:

One can notice that for all $i \in Z'$, all $k_i \geq q$, $\mathbb{I}\{N_{i,\tau-1} = k_i\} \left(\frac{1}{\mathbb{P}[|\theta_{i,\tau} - \mu_i^*| \leq \varepsilon | \mathcal{H}_\tau]} - 1 \right)$ is of the form $\mathbb{I}\{N_{i,\tau-1} = k_i\} g_i(|\bar{\mu}_{i,\tau-1} - \mu_i^*|)$, with g_i being an increasing function on \mathbb{R}_+ . Indeed, we see that the conditional distribution of $\theta_{i,\tau} - \bar{\mu}_{i,\tau-1}$ is $\mathcal{N}(0, \beta N_{i,\tau-1}^{-1}/4)$, which is symmetric, so we have

$$\mathbb{P}[|\theta_{i,\tau} - \mu_i^*| \leq \varepsilon | \mathcal{H}_\tau] = \mathbb{P}[|\theta_{i,\tau} - \bar{\mu}_{i,\tau-1} + |\bar{\mu}_{i,\tau-1} - \mu_i^*|| \leq \varepsilon | \mathcal{H}_\tau].$$

In addition, under $\mathbb{I}\{N_{i,\tau-1} = k_i\}$, the conditional distribution of $\theta_{i,\tau} - \bar{\mu}_{i,\tau-1}$ does not depend on the history, but only on k_i . Therefore, the above probability is a function of $|\bar{\mu}_{i,\tau-1} - \mu_i^*|$ and so the function g_i exists. It is increasing on \mathbb{R}_+ because for any fixed $\sigma > 0$,

$$\frac{\partial}{\partial x} \int_{x-\varepsilon}^{x+\varepsilon} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{u^2}{2\sigma^2}} du = \frac{1}{\sqrt{2\pi\sigma^2}} \left(e^{-\frac{(x+\varepsilon)^2}{2\sigma^2}} - e^{-\frac{(x-\varepsilon)^2}{2\sigma^2}} \right) < 0 \text{ for } x > 0.$$

In particular, we can consider the inverse function g_i^{-1} . We now want to use a stochastic dominance argument in order to treat the outcomes as if they were Gaussian: we have for any $\mathbf{k} \in [q.. \infty]^{Z'}$,

$$\begin{aligned} &\mathbb{E} \left[\sup_{\tau \geq \tau_q} \prod_{i \in Z'} \left(\mathbb{I}\{N_{i,\tau-1} = k_i\} g_i(|\bar{\mu}_{i,\tau-1} - \mu_i^*|) \right) \right] \\ &= \mathbb{E} \left[\sup_{\tau \geq \tau_q} \prod_{i \in Z'} \left(\mathbb{I}\{N_{i,\tau-1} = k_i\} \int_0^\infty \mathbb{I}\{g_i(|\bar{\mu}_{i,\tau-1} - \mu_i^*|) \geq u_i\} du_i \right) \right] \end{aligned}$$

$$\begin{aligned}
&\leq \int_{\mathbf{u} \in \mathbb{R}_+^{Z'}} \mathbb{E} \left[\sup_{\tau \geq \tau_q} \prod_{i \in Z'} \mathbb{I}\{N_{i,\tau-1} = k_i\} \mathbb{I}\{g_i(|\bar{\mu}_{i,\tau-1} - \mu_i^*|) \geq u_i\} \right] d\mathbf{u} \\
&= \int_{\mathbf{u} \in \mathbb{R}_+^{Z'}} \mathbb{E} \left[\prod_{i \in Z'} \mathbb{I}\{N_{i,\tau^*-1} = k_i\} \mathbb{I}\{g_i(|\bar{\mu}_{i,\tau^*-1} - \mu_i^*|) \geq u_i\} \right] d\mathbf{u}, \tag{21}
\end{aligned}$$

where τ^* is the first $\tau \geq \tau_q$ such that $\mathbb{I}\{\forall i \in Z', N_{i,\tau-1} = k_i \text{ and } g_i(|\bar{\mu}_{i,\tau-1} - \mu_i^*|) \geq u_i\}$ holds, and is ∞ if it never holds.

$$\begin{aligned}
(21) &= \int_{\mathbf{u} \in \mathbb{R}_+^{Z'}} \mathbb{E} \left[\prod_{i \in Z'} \mathbb{I}\{N_{i,\tau^*-1} = k_i\} \mathbb{I}\{g_i(|\bar{\mu}_{i,\tau^*-1} - \mu_i^*|) \geq u_i \vee g_i(0)\} \right] d\mathbf{u} \\
&= \int_{\mathbf{u} \in \mathbb{R}_+^{Z'}} \mathbb{E} \left[\prod_{i \in Z'} \mathbb{I}\{N_{i,\tau^*-1} = k_i\} \mathbb{I}\{|\bar{\mu}_{i,\tau^*-1} - \mu_i^*| \geq g_i^{-1}(u_i \vee g_i(0))\} \right] d\mathbf{u} \\
&= \int_{\mathbf{u} \in \mathbb{R}_+^{Z'}} \sum_{\mathbf{s} \in \{-1,1\}^{Z'}} \underbrace{\mathbb{E} \left[\prod_{i \in Z'} \mathbb{I}\{N_{i,\tau^*-1} = k_i\} \mathbb{I}\{s_i(\bar{\mu}_{i,\tau^*-1} - \mu_i^*) \geq g_i^{-1}(u_i \vee g_i(0))\} \right]}_{(22)} d\mathbf{u} \\
(22) &\leq \mathbb{P} \left[\frac{e^{\sum_{i \in Z'} N_{i,\tau^*-1} (4s_i g_i^{-1}(u_i \vee g_i(0)) (\bar{\mu}_{i,\tau^*-1} - \mu_i^*) - 2(g_i^{-1}(u_i \vee g_i(0)))^2)}}}{e^{\sum_{i \in Z'} 2(g_i^{-1}(u_i \vee g_i(0)))^2 k_i}} \geq 1, (N_{i,\tau^*-1})_{i \in Z'} = \mathbf{k} \right] \\
&\leq \mathbb{P} \left[\frac{e^{\sum_{i \in Z'} N_{i,\tau^*-1} (4s_i g_i^{-1}(u_i \vee g_i(0)) (\bar{\mu}_{i,\tau^*-1} - \mu_i^*) - 2(g_i^{-1}(u_i \vee g_i(0)))^2)}}}{e^{\sum_{i \in Z'} 2(g_i^{-1}(u_i \vee g_i(0)))^2 k_i}} \geq 1 \right] \\
&\leq \frac{\mathbb{E} \left[\exp \left(\sum_{i \in Z'} N_{i,\tau^*-1} (4s_i g_i^{-1}(u_i \vee g_i(0)) (\bar{\mu}_{i,\tau^*-1} - \mu_i^*) - 2(g_i^{-1}(u_i \vee g_i(0)))^2) \right) \right]}{\exp \left(\sum_{i \in Z'} 4(g_i^{-1}(u_i \vee g_i(0)))^2 k_i \right)} \\
&= \frac{\mathbb{E} \left[\exp \left(\sum_{t=1}^{\tau^*-1} \sum_{i \in Z' \cap S_t} (4s_i g_i^{-1}(u_i \vee g_i(0)) (X_{i,t} - \mu_i^*) - 2(g_i^{-1}(u_i \vee g_i(0)))^2) \right) \right]}{\exp \left(\sum_{i \in Z'} 2(g_i^{-1}(u_i \vee g_i(0)))^2 k_i \right)}.
\end{aligned}$$

From Assumption 3, and from the fact that either D_{trig} is independent from the outcomes, or the outcomes are mutually independent and each individual outcome is independent from the fact that it is triggered, we have that

$$M_\tau = \exp \left(\sum_{t=1}^{\tau-1} \sum_{i \in Z' \cap S_t} (4s_i g_i^{-1}(u_i \vee g_i(0)) (X_{i,t} - \mu_i^*) - 2(g_i^{-1}(u_i \vee g_i(0)))^2) \right)$$

is a supermartingale:

$$\begin{aligned}
\mathbb{E}[M_\tau | \mathcal{F}_{\tau-1}] &= M_{\tau-1} \mathbb{E} \left[e^{\sum_{i \in Z' \cap S_{\tau-1}} (4s_i g_i^{-1}(u_i \vee g_i(0)) (X_{i,\tau-1} - \mu_i^*) - 2(g_i^{-1}(u_i \vee g_i(0)))^2)} \middle| \mathcal{F}_{\tau-1} \right] \\
&\leq M_{\tau-1}.
\end{aligned}$$

Since τ^* is a stopping time with respect to \mathcal{F}_τ , we have from Doob's optional sampling theorem for non-negative supermartingales⁷ that $\mathbb{E}[M_{\tau^*}] \leq 1$. Therefore,

$$(22) \leq \exp \left(- \sum_{i \in Z'} 2(g_i^{-1}(u_i \vee g_i(0)))^2 k_i \right).$$

⁷We use the version that relies on Fatou's lemma (Durrett [2019], Theorem 5.7.6), so that it is not needed to have any additional condition on the stopping time τ^* .

Now, we want to use the following fact (see [Chang et al. \[2011\]](#)): if $\eta \sim \mathcal{N}(0, 1)$, then with $\beta > 1$,

$$\sqrt{\frac{2e}{\pi}} \frac{\sqrt{\beta-1}}{\beta} e^{-\beta x^2/2} \leq \mathbb{P}[|\eta| \geq x].$$

Indeed, this gives

$$\sqrt{\frac{2e}{\pi}} \frac{\sqrt{\beta-1}}{\beta} \exp\left(-2(g_i^{-1}(u_i \vee g_i(0)))^2 k_i\right) \leq \mathbb{P}\left[|\eta_i| \geq g_i^{-1}(u_i \vee g_i(0)) \sqrt{\frac{4k_i}{\beta}}\right],$$

where $\eta \sim \mathcal{N}(0, 1)^{\otimes Z'}$. Thus,

$$\begin{aligned} (21) &\leq \left(\sqrt{\frac{\pi}{2e}} \frac{2\beta}{\sqrt{\beta-1}}\right)^{|Z'|} \int_{\mathbf{u} \in \mathbb{R}_+^{Z'}} \prod_{i \in Z'} \mathbb{P}\left[\sqrt{\frac{\beta}{4k_i}} |\eta_i| \geq g_i^{-1}(u_i \vee g_i(0))\right] d\mathbf{u} \\ &= \left(\sqrt{\frac{\pi}{2e}} \frac{2\beta}{\sqrt{\beta-1}}\right)^{|Z'|} \int_{\mathbf{u} \in \mathbb{R}_+^{Z'}} \prod_{i \in Z'} \mathbb{P}\left[g_i\left(\sqrt{\frac{\beta}{4k_i}} |\eta_i|\right) \geq u_i \vee g_i(0)\right] d\mathbf{u} \\ &= \left(\sqrt{\frac{\pi}{2e}} \frac{2\beta}{\sqrt{\beta-1}}\right)^{|Z'|} \int_{\mathbf{u} \in \mathbb{R}_+^{Z'}} \prod_{i \in Z'} \mathbb{P}\left[g_i\left(\sqrt{\frac{\beta}{4k_i}} |\eta_i|\right) \geq u_i\right] d\mathbf{u} \\ &= \left(\sqrt{\frac{\pi}{2e}} \frac{2\beta}{\sqrt{\beta-1}}\right)^{|Z'|} \prod_{i \in Z'} \int_0^\infty \mathbb{P}\left[g_i\left(\sqrt{\frac{\beta}{4k_i}} |\eta_i|\right) \geq u_i\right] du_i \\ &= \left(\sqrt{\frac{\pi}{2e}} \frac{2\beta}{\sqrt{\beta-1}}\right)^{|Z'|} \prod_{i \in Z'} \mathbb{E}\left[g_i\left(\sqrt{\frac{\beta}{4k_i}} |\eta_i|\right)\right]. \end{aligned}$$

We now want to bound $\mathbb{E}\left[g_i\left(\sqrt{\frac{\beta}{4k_i}} |\eta_i|\right)\right]$. We define $\alpha = 2 - \sqrt{2}$, the unique solution in $(1/2, 1)$ of $\alpha - 1/2 = (\alpha - 1)^2/2$. Notice that $\alpha - 1/2 \geq 1/12$. Define $\varepsilon_i \triangleq \varepsilon \sqrt{\frac{4k_i}{\beta}}$. By definition, we have

$$\begin{aligned} \mathbb{E}\left[g_i\left(\sqrt{\frac{\beta}{4k_i}} |\eta_i|\right)\right] &= \int_{-\infty}^{+\infty} \frac{e^{-x^2/2}}{\int_{x-\varepsilon_i}^{x+\varepsilon_i} e^{-y^2/2} dy} dx - 1 \\ &= \underbrace{2 \int_{\alpha\varepsilon_i}^{+\infty} \frac{1}{\int_{x-\varepsilon_i}^{x+\varepsilon_i} e^{-\frac{y^2-x^2}{2}} dy} dx}_{A_1} + \underbrace{\int_{-\alpha\varepsilon_i}^{\alpha\varepsilon_i} \frac{e^{-x^2/2}}{\int_{x-\varepsilon_i}^{x+\varepsilon_i} e^{-y^2/2} dy} dx}_{A_2} - 1. \end{aligned}$$

We first bound A_1 . With the change of variable $u = y - x$, we get:

$$\begin{aligned} A_1 &= 2 \int_{\alpha\varepsilon_i}^{+\infty} \frac{1}{\int_{-\varepsilon_i}^{\varepsilon_i} e^{-u^2/2-ux} du} dx \\ &\leq 2 \int_{\alpha\varepsilon_i}^{+\infty} \frac{1}{\int_{-\varepsilon_i}^0 e^{-u^2/2-ux} du} dx \end{aligned}$$

Note that for $x \geq \alpha\varepsilon_i$ and $u \in [-\varepsilon_i, 0]$, $-u^2/2 - ux \geq -(1 - \frac{1}{2\alpha})ux$ and thus:

$$\begin{aligned} A_1 &\leq 2 \int_{\alpha\varepsilon_i}^{+\infty} \frac{1}{\int_{-\varepsilon_i}^0 e^{-(1-\frac{1}{2\alpha})ux} du} dx \\ &= 2 \int_{\alpha\varepsilon_i}^{+\infty} \frac{(1-\frac{1}{2\alpha})x}{e^{(1-\frac{1}{2\alpha})\varepsilon_i x} - 1} dx. \end{aligned} \tag{23}$$

We distinguish two regimes. First, if $\varepsilon_i^2 \geq 12$, then

$$(23) \leq \frac{2e^{(\alpha-\frac{1}{2})\varepsilon_i^2}}{e^{(\alpha-\frac{1}{2})\varepsilon_i^2} - 1} \int_{\alpha\varepsilon_i}^{+\infty} \left(1 - \frac{1}{2\alpha}\right) x e^{-(1-\frac{1}{2\alpha})\varepsilon_i x} dx$$

$$\begin{aligned}
&= \frac{2e^{(\alpha-\frac{1}{2})\varepsilon_i^2}}{e^{(\alpha-\frac{1}{2})\varepsilon_i^2}-1} \frac{1}{(1-\frac{1}{2\alpha})\varepsilon_i^2} \int_{(\alpha-\frac{1}{2})\varepsilon_i^2}^{+\infty} x e^{-x} dx \\
&= \frac{2e^{(\alpha-\frac{1}{2})\varepsilon_i^2}}{e^{(\alpha-\frac{1}{2})\varepsilon_i^2}-1} \frac{1}{(1-\frac{1}{2\alpha})\varepsilon_i^2} [-(x+1)e^{-x}]_{(\alpha-\frac{1}{2})\varepsilon_i^2}^{\infty} \\
&= \frac{2e^{(\alpha-\frac{1}{2})\varepsilon_i^2}}{e^{(\alpha-\frac{1}{2})\varepsilon_i^2}-1} \frac{1}{(1-\frac{1}{2\alpha})\varepsilon_i^2} \left(\left(\alpha - \frac{1}{2} \right) \varepsilon_i^2 + 1 \right) e^{-(\alpha-\frac{1}{2})\varepsilon_i^2} \\
&= \frac{2}{e^{(\alpha-\frac{1}{2})\varepsilon_i^2}-1} \left(\alpha + \frac{\alpha}{(\alpha-\frac{1}{2})\varepsilon_i^2} \right) \\
&\leq 4e^{-\varepsilon_i^2/12}.
\end{aligned}$$

Otherwise, we have

$$\begin{aligned}
(23) &= \frac{2(1-\frac{1}{2\alpha})}{\varepsilon_i^2} \int_{\alpha\varepsilon_i^2}^{\infty} \frac{u}{e^{(1-\frac{1}{2\alpha})u}-1} du \\
&\leq \frac{2(1-\frac{1}{2\alpha})}{\varepsilon_i^2} \int_0^{\infty} \frac{u}{e^{(1-\frac{1}{2\alpha})u}-1} du \\
&= \frac{2(1-\frac{1}{2\alpha})}{\varepsilon_i^2} \frac{\pi^2}{6(1-\frac{1}{2\alpha})^2} \\
&\leq \frac{6\beta}{\varepsilon^2}.
\end{aligned}$$

We now bound A_2 . As $x \in [-\alpha\varepsilon_i, \alpha\varepsilon_i]$, it comes that $[-(1-\alpha)\varepsilon_i, (1-\alpha)\varepsilon_i] \subset [x-\varepsilon_i, x+\varepsilon_i]$. This implies that

$$\begin{aligned}
A_2 &\leq \frac{\int_{-\alpha\varepsilon_i}^{\alpha\varepsilon_i} e^{-x^2/2} dx}{\int_{-(1-\alpha)\varepsilon_i}^{(1-\alpha)\varepsilon_i} e^{-x^2/2} dx} - 1 \\
&= \frac{2 \int_{(1-\alpha)\varepsilon_i}^{\alpha\varepsilon_i} e^{-x^2/2} dx}{\int_{-(1-\alpha)\varepsilon_i}^{(1-\alpha)\varepsilon_i} e^{-x^2/2} dx} \\
&\leq \frac{2 \int_{(1-\alpha)\varepsilon_i}^{\infty} e^{-x^2/2} dx}{\int_{-(1-\alpha)\varepsilon_i}^{(1-\alpha)\varepsilon_i} e^{-x^2/2} dx} \\
&\leq \frac{e^{-(1-\alpha)^2\varepsilon_i^2/2}}{1 - e^{-(1-\alpha)^2\varepsilon_i^2/2}} \leq \left(1 + \frac{12}{\varepsilon_i^2} \right) e^{-\varepsilon_i^2/12}.
\end{aligned}$$

The penultimate inequality relies on $\int_x^{\infty} e^{-u^2/2} du \leq \sqrt{\frac{\pi}{2}} e^{-x^2/2}$ (see [Jacobs and Wozenkraft \[1965\]](#), eq. (2.122)). We obtain again two regimes: $2e^{-\varepsilon_i^2/12}$ if $\varepsilon_i^2 \geq 12$, and $1 + \frac{3\beta}{\varepsilon^2}$ otherwise. To summarize, we proved that (21) is bounded by

$$\left(\sqrt{\frac{\pi}{2e}} \frac{2\beta}{\sqrt{\beta-1}} \right)^{|Z'|} \prod_{i \in Z'} \left(\mathbb{I} \left\{ \varepsilon^2 \frac{4k_i}{\beta} < 12 \right\} \left(1 + 9 \frac{\beta}{\varepsilon^2} \right) + \mathbb{I} \left\{ \varepsilon^2 \frac{4k_i}{\beta} \geq 12 \right\} 6e^{-\varepsilon^2 \frac{k_i}{3\beta}} \right).$$

After the summation on \mathbf{k} , on Z' , on q , and on Z , we obtain that there exists two constants C, C' such that

$$\sum_{Z \subset \tau(A^*), Z \neq \emptyset} \sum_{q \geq 1} \sum_{Z' \subset Z, Z' \neq \emptyset} \sum_{\mathbf{k} \in \{q, \infty\}^{Z'}} (21) \leq (C\varepsilon^{-2}\beta) \left(\frac{C'\beta}{\sqrt{\beta-1}} \varepsilon^{-4}\beta^2 \right)^{m^*}.$$

□

Step 2: bound under $\mathfrak{Z}_t \wedge \mathfrak{B}_t$ The filtered regret bound

$$\mathbb{E} \left[\sum_{t=1}^T \Delta_t \mathbb{I}\{\mathfrak{Z}_t \wedge \mathfrak{B}_t\} \right] \leq \frac{n\Delta_{\max}}{p^*} \left(1 + \left(\frac{\Delta_{\min}}{2nB} - \frac{(m^{*2} + 1)\varepsilon}{n} \right)^{-2} \right)$$

is obtained as follows. Let $t \geq 1$. First, note that \mathfrak{B}_t implies

$$\left\{ \exists i \in \mathsf{T}(A_t) \text{ s.t. } nB_i |\bar{\mu}_{i,t-1} - \mu_i^*| > \Delta_{\min}/2 - B(m^{*2} + 1)\varepsilon \right\}.$$

Then, fixing $i \in [n]$, we can ensure that $i \in S_t$ in the event, using that $p_i(A_t) = \mathbb{P}[i \in S_t | \mathcal{F}_t]$:

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \mathbb{I}\left\{ i \in \mathsf{T}(A_t), nB_i |\bar{\mu}_{i,t-1} - \mu_i^*| > \Delta_{\min}/2 - B(m^{*2} + 1)\varepsilon \right\} \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \frac{p_i(A_t)}{p_i(A_t)} \mathbb{I}\left\{ i \in \mathsf{T}(A_t), nB_i |\bar{\mu}_{i,t-1} - \mu_i^*| > \Delta_{\min}/2 - B(m^{*2} + 1)\varepsilon \right\} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\sum_{t=1}^T \frac{1}{p_i(A_t)} \mathbb{I}\left\{ i \in S_t, nB_i |\bar{\mu}_{i,t-1} - \mu_i^*| > \Delta_{\min}/2 - B(m^{*2} + 1)\varepsilon \right\} \middle| \mathcal{F}_t \right] \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \frac{1}{p_i(A_t)} \mathbb{I}\left\{ i \in S_t, nB_i |\bar{\mu}_{i,t-1} - \mu_i^*| > \Delta_{\min}/2 - B(m^{*2} + 1)\varepsilon \right\} \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^T \frac{1}{p^*} \mathbb{I}\left\{ i \in S_t, nB_i |\bar{\mu}_{i,t-1} - \mu_i^*| > \Delta_{\min}/2 - B(m^{*2} + 1)\varepsilon \right\} \right] \\ &\leq \frac{1}{p^*} \sum_{t \geq 0} 1 \wedge \left(2 \exp \left(-2t \left(\frac{\Delta_{\min}}{2nB_i} - \frac{B(m^{*2} + 1)\varepsilon}{nB_i} \right)^2 \right) \right) \\ &\leq \frac{1}{p^*} \left(1 + \frac{2 \exp \left(-2 \left(\frac{\Delta_{\min}}{2nB_i} - \frac{B(m^{*2} + 1)\varepsilon}{nB_i} \right)^2 \right)}{1 - \exp \left(-2 \left(\frac{\Delta_{\min}}{2nB_i} - \frac{B(m^{*2} + 1)\varepsilon}{nB_i} \right)^2 \right)} \right) \\ &\leq \frac{1 + \left(\frac{\Delta_{\min}}{2nB_i} - \frac{B(m^{*2} + 1)\varepsilon}{nB_i} \right)^{-2}}{p^*} \\ &\leq \frac{1 + \left(\frac{\Delta_{\min}}{2nB} - \frac{(m^{*2} + 1)\varepsilon}{n} \right)^{-2}}{p^*}. \end{aligned}$$

Step 3: bound under $\mathfrak{Z}_t \wedge \mathfrak{D}_t$ The filtered regret bound

$$\mathbb{E} \left[\sum_{t=1}^T \Delta_t \mathbb{I}\{\mathfrak{Z}_t \wedge \mathfrak{D}_t\} \right] \leq \Delta_{\max} \sum_{t \in [T]} \mathbb{E}[\mathbb{P}[\mathfrak{D}_t | \mathcal{H}_t]] \leq \Delta_{\max} \sum_{t \in [T]} 1/T = \Delta_{\max}$$

follows from the following Lemma 3.

Lemma 3. *In Algorithm 1 and 2, for all round $t \geq 1$, we have that the probability*

$$\mathbb{P} \left[\left\| \mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1}) \right\|_1 \geq \sqrt{2 \log(2^n (1 + \lceil \log_2(1/p^*) \rceil)^n T) \sum_{i \in [n]} \frac{p_i(A_t)^2 B_i^2 \beta}{N_{i,t-1}}} \middle| \mathcal{H}_t \right].$$

is lower than $1/T$.

Proof. We rely on the fact that conditionally on the history, the sample $\boldsymbol{\theta}_t$ is either a Gaussian random vector of mean $\bar{\boldsymbol{\mu}}_{t-1}$ and of diagonal covariance given by $\beta N_{i,t-1}^{-1}/4$ (for Algorithm 2), or a product

of Beta random variable, that is sub-Gaussian with the same covariance matrix [Marchal et al., 2017] (for Algorithm 1). We thus define the functions

$$\alpha_t(A) \triangleq \sqrt{2\log(2^n(1 + \lceil \log_2(1/p^*) \rceil)^n T) \sum_{i \in [n]} \frac{p_i(A)^2 B_i^2 \beta}{N_{i,t-1}}},$$

$$\lambda_t(A) \triangleq \frac{2\alpha_t(A)}{\sum_{i \in A} \beta B_i^2 p_i(A)^2 / N_{i,t-1}},$$

we have, with $Q \triangleq (\{0\} \cup \{2^{-k}, k \in [\lceil \log_2(1/p^*) \rceil]\})^n$,

$$\begin{aligned} & \mathbb{P}[\|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1 \geq \alpha_t(\mathbf{p}(A_t)) | \mathcal{H}_t] \\ & \leq \sum_{\mathbf{q} \in Q} \mathbb{P}[\mathbf{q} \leq \mathbf{p}(A_t) \leq 2\mathbf{q}, \|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1 \geq \alpha_t(\mathbf{p}(A_t)) | \mathcal{H}_t] \\ & \leq \sum_{\mathbf{q} \in Q} \mathbb{P}[\|\mathbf{q} \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1 \geq \alpha_t(\mathbf{q})/2 | \mathcal{H}_t] \\ & \leq \sum_{\mathbf{q} \in Q} e^{-\lambda_t(\mathbf{q})\alpha_t(\mathbf{q})/2} \mathbb{E}\left[e^{\lambda_t(\mathbf{q})\|\mathbf{q} \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1} | \mathcal{H}_t\right] \\ & \leq \sum_{\mathbf{q} \in Q} e^{-\lambda_t(\mathbf{q})\alpha_t(\mathbf{q})/2} \prod_{i \in [n]} \mathbb{E}\left[e^{\lambda_t(\mathbf{q})q_i B_i |\theta_{i,t} - \bar{\mu}_{i,t-1}|} | \mathcal{H}_t\right] \\ & \leq \sum_{\mathbf{q} \in Q} e^{-\lambda_t(\mathbf{q})\alpha_t(\mathbf{q})/2} \prod_{i \in [n]} \mathbb{E}\left[e^{\lambda_t(\mathbf{q})q_i B_i (\theta_{i,t} - \bar{\mu}_{i,t-1})} + e^{\lambda_t(\mathbf{q})q_i B_i (\bar{\mu}_{i,t-1} - \theta_{i,t})} | \mathcal{H}_t\right] \\ & \leq \sum_{\mathbf{q} \in Q} 2^n e^{-\lambda_t(\mathbf{q})\alpha_t(\mathbf{q})/2} e^{\lambda_t(\mathbf{q})^2 \sum_{i \in A} \beta B_i^2 q_i^2 / (8N_{i,t-1})} \leq 1/T. \end{aligned}$$

□

Step 4: bound under $\mathfrak{Z}_t \wedge \mathfrak{C}_t \wedge \neg \mathfrak{B}_t \wedge \neg \mathfrak{D}_t$ We get that $\mathbb{E}\left[\sum_{t=1}^T \Delta_t \mathbb{I}\{\mathfrak{Z}_t \wedge \mathfrak{C}_t \wedge \neg \mathfrak{B}_t \wedge \neg \mathfrak{D}_t\}\right]$ is bounded by

$$n\Delta_{\max} + \sum_{i \in [n]} \frac{8(3 + \log(m))\beta B_i^2 \log\left(2^n \left(1 + \lceil \log_2\left(\frac{1}{p^*}\right) \rceil\right)^n T\right)}{\min_{A \in \mathcal{A}, p_i(A) > 0, \Delta(A) > 0} \Delta(A)/p_i(A)}$$

from the following derivations. Let $t \geq 1$. Under $\mathfrak{Z}_t \wedge \mathfrak{C}_t \wedge \neg \mathfrak{B}_t \wedge \neg \mathfrak{D}_t$, we have

$$\begin{aligned} \Delta_t & \leq \|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \boldsymbol{\mu}^*)\|_1 + B(m^{*2} + 1)\varepsilon & \mathfrak{C}_t \\ & \leq \|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1 + \|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\bar{\boldsymbol{\mu}}_{t-1} - \boldsymbol{\mu}^*)\|_1 + B(m^{*2} + 1)\varepsilon \\ & \leq \|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1 + \Delta_{\min}/2 - B(m^{*2} + 1)\varepsilon + B(m^{*2} + 1)\varepsilon & \neg \mathfrak{B}_t \\ & \leq \|\mathbf{p}(A_t) \odot \mathbf{B} \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1 + \Delta_t/2 & \mathfrak{Z}_t \\ & \leq \sqrt{2\log(2^n(1 + \lceil \log_2(1/p^*) \rceil)^n T) \sum_{i \in A_t} \frac{p_i(A_t)^2 B_i^2 \beta}{N_{i,t-1}}} + \Delta_t/2. & \neg \mathfrak{D}_t \end{aligned}$$

Thus, the following event holds

$$\mathfrak{A}_t \triangleq \left\{ \Delta_t \leq \sqrt{4\log(2^n(1 + \lceil \log_2(1/p^*) \rceil)^n T) \sum_{i \in A_t} \frac{p_i(A_t)^2 B_i^2 \beta}{N_{i,t-1}}} \right\},$$

and we can apply Lemma 4 to get the bound

$$\mathbb{E}\left[\sum_{t=1}^T \Delta_t \mathbb{I}\{\mathfrak{A}_t\}\right] \leq n\Delta_{\max} + \sum_{i \in [n]} \frac{(24 + 8\log(m))\beta B_i^2 \log(2^n(1 + \lceil \log_2(1/p^*) \rceil)^n T)}{\min_{A \in \mathcal{A}, p_i(A) > 0, \Delta(A) > 0} \Delta(A)/p_i(A)}.$$

Lemma 4 (Adapted from Wang and Chen [2018]). *Let's fix the time horizon T . For all $i \in [n]$, let $\beta_{i,T} \in \mathbb{R}_+$. For $t \geq 1$, consider the event*

$$\mathfrak{A}_t \triangleq \left\{ \Delta_t \leq \sqrt{\frac{1}{2} \sum_{i \in [n]} \frac{p_i(A_t)^2 \beta_{i,T}}{N_{i,t-1}}} \right\}.$$

Let

$$\delta_{i,\min} \triangleq \min_{A \in \mathcal{A}, p_i(A) > 0, \Delta(A) > 0} \Delta(A)/p_i(A).$$

Then, we have

$$\mathbb{E} \left[\sum_{t=1}^T \Delta_t \mathbb{I}\{\mathfrak{A}_t\} \right] \leq n \Delta_{\max} + \sum_{i \in [n]} \frac{(3 + \log(m)) \beta_{i,T}}{\delta_{i,\min}}.$$

Proof. We use the regret allocation method from Wang and Chen [2018]. Specifically, we want to prove that for any time step t where \mathfrak{A}_t holds, we have the following allocation of the regret to each arm $i \in [n]$:

$$\Delta_t \leq \sum_{i \in [n]} g_i(N_{i,t-1}), \quad (24)$$

where the allocation functions are defined for all $i \in [n]$ as

$$g_i(t) \triangleq \mathbb{I}\{t = 0\} \Delta_{\max} + \mathbb{I}\{0 < t \leq L_{i,2}\} \frac{p_i(A_t) \beta_{i,T}^{1/2}}{t^{1/2}} + \mathbb{I}\{L_{i,2} < t \leq L_{i,1}\} \frac{p_i(A_t) \beta_{i,T}}{t \delta_{i,\min}},$$

$$L_{i,1} \triangleq \frac{m \beta_{i,T}}{\delta_{i,\min}^2}, \quad L_{i,2} \triangleq \frac{\beta_{i,T}}{\delta_{i,\min}^2}.$$

Indeed, we can already see that such an allocation produces the bound we are looking for. Notably the following derivation uses, for $i \in [n]$, the equality $p_i(A_t) = \mathbb{E}[\mathbb{I}\{N_{i,t} = N_{i,t-1} + 1\} | \mathcal{F}_t]$, so that for all function f , a sum of the form $\sum_{t \geq 1} p_i(A_t) f(N_{i,t-1})$ is equal in expectation to $\sum_{t \geq 0} f(t)$.

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \mathbb{I} \left\{ \Delta_t \leq \sum_{i \in [n]} g_i(N_{i,t-1}) \right\} \Delta_t \right] \\ & \leq \mathbb{E} \left[\sum_{t \in [T]} \sum_{i \in [n]} g_i(N_{i,t-1}) \right] \\ & \leq n \Delta_{\max} + \mathbb{E} \left[\sum_{i \in [n]} \sum_{t=1}^{L_{i,2}} \frac{\beta_{i,T}^{1/2}}{t^{1/2}} \right] + \mathbb{E} \left[\sum_{i \in [n]} \sum_{t=L_{i,2}+1}^{L_{i,1}} \frac{\beta_{i,T}}{t \delta_{i,\min}} \right] \\ & \leq n \Delta_{\max} + \sum_{i \in [n]} 2 \sqrt{L_{i,2} \beta_{i,T}} + \sum_{i \in [n]} \left(1 + \log \left(\frac{L_{i,1}}{L_{i,2}} \right) \right) \frac{\beta_{i,T}}{\delta_{i,\min}} \\ & = n \Delta_{\max} + \sum_{i \in [n]} \frac{2 \beta_{i,T}}{\delta_{i,\min}} + \sum_{i \in [n]} (1 + \log(m)) \frac{\beta_{i,T}}{\delta_{i,\min}}. \end{aligned}$$

Therefore, we prove the lemma if we show that \mathfrak{A}_t implies (24). Let $t \geq 1$ be such that \mathfrak{A}_t holds. We can assume that $\Delta_t > 0$ (otherwise, the inequality (24) is trivial). We first observe that

$$\begin{aligned} \sum_{i \in [n], N_{i,t-1} > L_{i,1}} \frac{p_i(A_t)^2 \beta_{i,T}}{\Delta_t N_{i,t-1}} & \leq \sum_{i \in [n], N_{i,t-1} > L_{i,1}} \frac{p_i(A_t)^2 \beta_{i,T}}{\Delta_t L_{i,1}} \\ & = \sum_{i \in [n], N_{i,t-1} > L_{i,1}} \left(\frac{\delta_{i,\min}}{\Delta_t / p_i(A_t)} \right)^2 \frac{\Delta_t}{m} \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{i \in [n], N_{i,t-1} > L_{i,1}} \frac{\Delta_t}{m} \\
&\leq \Delta_t.
\end{aligned} \tag{25}$$

This will be useful to prove that the allocation on all arms i such that $N_{i,t-1} > L_{i,1}$ can be 0. Then we distinguish the following cases:

- If there exists $i_0 \in A_t$ such that $N_{i_0,t-1} = 0$, then

$$\Delta_t \leq \Delta_{\max} = g_{i_0}(N_{i_0,t-1}) \leq \sum_{i \in [n]} g_i(N_{i,t-1}).$$

- If there exists $i_0 \in A_t$ such that $0 < N_{i_0,t-1} \leq p_{i_0}(A_t)^2 \beta_{i_0,T} / \Delta_t^2$, then $N_{i_0,t-1} \leq L_{i_0,2}$ and we have

$$\Delta_t \leq \frac{p_{i_0}(A_t) \beta_{i_0,T}^{1/2}}{N_{i_0,t-1}^{1/2}} = g_{i_0}(N_{i_0,t-1}) \leq \sum_{i \in [n]} g_i(N_{i,t-1}).$$

- If for all $i \in [n]$, $N_{i,t-1} > p_i(A_t)^2 \beta_{i,T} / \Delta_t^2$, then,

$$\begin{aligned}
\sum_{i \in [n], L_{i,2} \geq N_{i,t-1}} \frac{p_i(A_t)^2 \beta_{i,T}}{\Delta_t N_{i,t-1}} &= \sum_{i \in [n], L_{i,2} \geq N_{i,t-1}} g_i(N_{i,t-1}) \frac{p_i(A_t) \beta_{i,T}^{1/2}}{\Delta_t N_{i,t-1}^{1/2}} \\
&\leq \sum_{i \in [n], L_{i,2} \geq N_{i,t-1}} g_i(N_{i,t-1}).
\end{aligned} \tag{26}$$

On the other hand, using the event \mathfrak{A}_t , we have

$$\begin{aligned}
\Delta_t &\leq \sum_{i \in [n]} \frac{p_i(A_t)^2 \beta_{i,T}}{2\Delta_t N_{i,t-1}} \\
&= \sum_{i \in [n], L_{i,1} \geq N_{i,t-1}} \frac{p_i(A_t)^2 \beta_{i,T}}{2\Delta_t N_{i,t-1}} + \sum_{i \in [n], N_{i,t-1} > L_{i,1}} \frac{p_i(A_t)^2 \beta_{i,T}}{2\Delta_t N_{i,t-1}}.
\end{aligned}$$

Now, using (25), we get

$$\Delta_t \leq \sum_{i \in [n], L_{i,1} \geq N_{i,t-1}} \frac{p_i(A_t)^2 \beta_{i,T}}{2\Delta_t N_{i,t-1}} + \frac{\Delta_t}{2}.$$

We can therefore end the proof in the following way, using (26),

$$\begin{aligned}
\Delta_t &\leq \sum_{i \in [n], L_{i,1} \geq N_{i,t-1}} \frac{p_i(A_t)^2 \beta_{i,T}}{\Delta_t N_{i,t-1}} \\
&= \sum_{i \in [n], L_{i,2} \geq N_{i,t-1}} \frac{p_i(A_t)^2 \beta_{i,T}}{\Delta_t N_{i,t-1}} + \sum_{i \in [n], L_{i,1} \geq N_{i,t-1} > L_{i,2}} \frac{p_i(A_t)^2 \beta_{i,T}}{\Delta_t N_{i,t-1}} \\
&\leq \sum_{i \in [n], L_{i,2} \geq N_{i,t-1}} g_i(N_{i,t-1}) + \sum_{i \in [n], L_{i,1} \geq N_{i,t-1} > L_{i,2}} \frac{p_i(A_t) \beta_{i,T}}{\delta_{i,\min} N_{i,t-1}} \\
&= \sum_{i \in [n], L_{i,2} \geq N_{i,t-1}} g_i(N_{i,t-1}) + \sum_{i \in [n], L_{i,1} \geq N_{i,t-1} > L_{i,2}} g_i(N_{i,t-1}) \\
&= \sum_{i \in [n]} g_i(N_{i,t-1}).
\end{aligned}$$

□

B Proof of Theorem 2

As for Theorem 1, in the following proof, we will treat both algorithms at the same time. Let's first remark that we can apply steps 1,2 and 3 from Theorem 1, but taking r_1 instead of the true reward function. Indeed, letting $E_{j,t} \in \arg \max_{E \in \mathcal{E}_{j,t}} r_j(E, \boldsymbol{\theta}_t)$ and $E_{j,t}^* \in \arg \max_{E \in \mathcal{E}_{j,t}} r_j(E, \boldsymbol{\mu}^*)$ for all $t \geq 1, j \in [\ell]$, we see that the sub-policy playing $E_{1,t}$ at round t is actually minimizing the regret with respect to r_1 using CTS with an exact oracle. Since r_1 satisfies the assumptions required for Theorem 1, we can apply the steps. Although we still get Δ_{\max} in these bounds (because the suffered regret remains Δ_t), we notice that Δ_{\min} is replaced by the minimal gap with respect to r_1 . To get around this issue, when applying steps 1,2 and 3, we place ourselves under the event that the gap with respect to r_1 is greater than $\Delta_{\min}/(2 \sum_{j \in [\ell]} c_j)$. We can thus replace the minimal gap with respect to r_1 by this quantity in the bounds. To summarize, we can either place ourselves under the events (for r_1) of step 4, obtaining in parallel 3 constant terms, or place ourselves under the event that the gap with respect to r_1 is lower than $\Delta_{\min}/(2 \sum_{j \in [\ell]} c_j)$. Our goal now is to do the same for the other reward functions r_j . However, since $\mathcal{E}_{j,t}$ can depend on $E_{1,t}, \dots, E_{j-1,t}$, we define the following filtration

$$(\mathcal{G}_0, \mathcal{G}_1, \dots, \mathcal{G}_{\ell-1}) = (\mathcal{H}_t, \sigma(\mathcal{H}_t, E_{1,t}), \sigma(\mathcal{H}_t, E_{1,t}, E_{2,t}), \dots, \sigma(\mathcal{H}_t, E_{1,t}, \dots, E_{\ell-1,t})).$$

Let us suppose that we have treated the r_1, \dots, r_{j-1} cases, then, we have at our disposal a filtered approximation regret against events $\mathfrak{Y}_{1,t}, \dots, \mathfrak{Y}_{j-1,t}$ that r_1, \dots, r_{j-1} are either in the situation of the step 4 or such the corresponding gap is lower than $\Delta_{\min}/(2 \sum_{j \in [\ell]} c_j)$, respectively. We can write the filtered approximation regret in the following way, by conditioning the expectation with this filtration to get rid of the randomness carried by $\mathcal{E}_{j,t}$:

$$\mathbb{E} \left[\mathbb{E} \left[\sum_{t \in [T]} \Delta_t \middle| \mathcal{G}_{j-1} \right] \mathbb{I}\{\mathfrak{Y}_{1,t}, \dots, \mathfrak{Y}_{j-1,t}\} \right].$$

We can now apply the same procedure as described above with the reward function r_j , on the inner conditional expectation, thus obtaining 3 additional T -independent terms and the new filtered approximation regret

$$\mathbb{E} \left[\mathbb{E} \left[\sum_{t \in [T]} \Delta_t \mathbb{I}\{\mathfrak{Y}_{j,t}\} \middle| \mathcal{G}_{j-1} \right] \mathbb{I}\{\mathfrak{Y}_{1,t}, \dots, \mathfrak{Y}_{j-1,t}\} \right] = \mathbb{E} \left[\sum_{t \in [T]} \Delta_t \mathbb{I}\{\mathfrak{Y}_{1,t}, \dots, \mathfrak{Y}_{j,t}\} \right].$$

Therefore, in the end, we have $3j$ constant terms and a filtered approximation regret where all reward functions r_j are in the situation of step 4 or with a corresponding gap lower than $\Delta_{\min}/(2 \sum_{j \in [\ell]} c_j)$. Now, we place ourselves under this event to derive the dominant term of the bound on the approximation regret. We let J be the indices j such that r_j are in the situation of step 4. The derivation of step 4 applied to a function r_j for $j \in J$ gives

$$r_j(E_{j,t}^*, \boldsymbol{\mu}^*) - r_j(E_{j,t}, \boldsymbol{\mu}^*) \leq 2 \|\mathbf{p}(A_t) \odot \mathbf{B}_j \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1.$$

We can therefore use equation (3) to obtain

$$\begin{aligned} \Delta_t &\leq \sum_{j \in [\ell]} (r_j(E_{j,t}^*, \boldsymbol{\mu}^*) - r_j(E_{j,t}, \boldsymbol{\mu}^*)) \cdot c_j \\ &\leq 2 \sum_{j \in [\ell]} \left(r_j(E_{j,t}^*, \boldsymbol{\mu}^*) - r_j(E_{j,t}, \boldsymbol{\mu}^*) - \Delta_{\min}/(2 \sum_{j \in [\ell]} c_j) \right) \cdot c_j \\ &\leq 2 \sum_{j \in J} \left(r_j(E_{j,t}^*, \boldsymbol{\mu}^*) - r_j(E_{j,t}, \boldsymbol{\mu}^*) - \Delta_{\min}/(2 \sum_{j \in [\ell]} c_j) \right) \cdot c_j \\ &\leq 4 \sum_{j \in J} c_j \|\mathbf{p}(A_t) \odot \mathbf{B}_j \odot (\boldsymbol{\theta}_t - \bar{\boldsymbol{\mu}}_{t-1})\|_1. \end{aligned}$$

From there, we can repeat the end of the proof of Theorem 1 with the weight in front of each arm i being $\sum_{j \in [\ell]} B_{ij} c_j$.

C Travelling salesman problem (TSP)

Here we give another example of a problem that falls into the REDUCE2EXACT setting, namely the *Travelling salesman problem* (TSP). We will see that unlike the examples mentioned before, here we need to modify the algorithm a bit to fully fall within the REDUCE2EXACT setting.

Given a complete undirected weighted graph $G = (V, E)$ whose distances $\mu^* \triangleq (d(u, v))_{(u, v) \in E}$ have to satisfy the triangle inequality, the goal is to find an Hamiltonian cycle $A \in \mathcal{A} \triangleq \{ \{(v_0, v_1), \dots, (v_{|V|-1}, v_{|V|})\} : \{v_1, \dots, v_{|V|} = v_0\} = V \}$ of minimum cost $\sum_{i \in [|V|]} d(v_{i-1}, v_i)$. We consider the following oracle from the [Christofides \[1976\]](#) algorithm ($\alpha = 2/3$).

- Oracle₁(μ) : The algorithm of Christofides combines $\mathcal{E}_1 \triangleq \{\text{spanning trees of } G\}$ and $\mathcal{E}_2 \triangleq \{\text{perfect matchings of the subgraph } G' \text{ of } G \text{ induced by the vertices of odd order in } E_1\}$, with r_1 being the weight of the spanning tree and r_2 the weight of the perfect matching.
- Oracle₂(E_1, E_2) : Oracle₂ combines the edges of E_1 and E_2 to form a connected multigraph $\tilde{G} = (\tilde{V}, \tilde{E})$ in which all vertices have even degree (so it is Eulerian), forms an Eulerian circuit in this multigraph, and finally, outputs the Hamiltonian cycle obtained by skipping repeated vertices (shortcutting).

Thanks to the triangle inequality, shortcutting does not increase the weight, so we have

$$\mathbf{e}_A^\top \mu^* \leq r_1(E_1, \mu^*) + r_2(E_2, \mu^*).$$

Let's now deal with A^* . Removing an edge from A^* produces a spanning tree, so $\mathbf{e}_{A^*}^\top \mu^* \geq r_1(E_1^*, \mu^*)$. On the other hand, by the triangle inequality, the weight of the optimal TSP solution for G' is lower than $\mathbf{e}_{A^*}^\top \mu^*$ (visiting more nodes does not, in any case, reduce the total cost). Taking every second edge of this cycle (which is of even length since all graphs have an even number of vertices of odd degree) we obtain a matching that has a weight less than half the weight of the cycle (if this is not the case we can take the complementary), so $\mathbf{e}_{A^*}^\top \mu^* / 2 \geq r_2(E_2^*, \mu^*)$. To summarize, we have

$$\frac{2}{3} \mathbf{e}_A^\top \mu^* - \mathbf{e}_{A^*}^\top \mu^* \leq \frac{2}{3} (r_1(E_1, \mu^*) - r_1(E_1^*, \mu^*) + r_2(E_2, \mu^*) - r_2(E_2^*, \mu^*)).$$

From the above, we can see that all the criteria of REDUCE2EXACT are satisfied, except Assumption 2. Indeed, we assume that we receive feedback from A , while we would need feedback from the set \tilde{E} . We could probably have foreseen that the TSP would pose a difficulty in our assumption: indeed, among the many operations performed by the oracle to build the final solution, shortcutting is the one that does not imply an optimization, but rather makes the solution feasible (so it does not represent a sub-problem as we have defined it in this paper). In other words, if we allowed the tour to pass over the same vertex several times, then the TSP would belong to REDUCE2EXACT by skipping the shortcutting step. Yet, we note that a workaround is possible by taking a closer look at the shortcutting step: in this step, we have an Eulerian circuit that we follow by skipping some edges. Even if the skipped edges are replaced by new ones such that the distance traveled decreases, it is precisely the absence of feedback on the skipped edges that poses an issue. Therefore, for a given edge of the Eulerian circuit, it would be helpful to be able to guarantee that some feedback is obtained on this edge, i.e., that it has a chance to belong to the final Hamiltonian cycle. The trick is then to notice that the shortcutting step depends on the edge from which we start. In particular, this edge is guaranteed to be in the final Hamiltonian cycle. This choice of the first edge is generally presupposed to be arbitrary and has no influence on the guarantees obtained previously, but for us, it can be used to force an edge to produce a feedback. More precisely, to choose this first edge, we can use a uniform randomization on \tilde{E} , meaning that for each edge $e \in \tilde{E}$, the probability q_e that e belongs to A is such that $q_e \geq \frac{1}{|\tilde{E}|}$. Notice that since \tilde{E} is a multigraph edge set, for a "true" edge $e \in E$, we can have two edges $e_1, e_2 \in \tilde{E}$ representing it and in that case $q_{e_1} + q_{e_2} = p_e(A)$. We thus get our Assumption 4:

$$\begin{aligned} |r_1(E_1, \mu) + r_2(E_2, \mu) - (r_1(E_1, \mu') + r_2(E_2, \mu'))| &\leq \sum_{e \in \tilde{E}} |\mu_e - \mu'_e| \leq \left| \tilde{E} \right| \sum_{e \in \tilde{E}} q_e |\mu_e - \mu'_e| \\ &= \left| \tilde{E} \right| \sum_{e \in E} p_e(A) |\mu_e - \mu'_e|. \end{aligned}$$