

Logarithmic Regret Bounds for Bandits with Knapsacks*

Arthur Flajolet

Operations Research Center
Massachusetts Institute of Technology
Cambridge, MA 02139
flajolet@mit.edu

Patrick Jaillet

Department of Electrical Engineering and Computer Science
Laboratory for Information and Decision Systems
Operations Research Center
Massachusetts Institute of Technology
Cambridge, MA 02139
jaillet@mit.edu

Optimal regret bounds for Multi-Armed Bandit problems are now well documented. They can be classified into two categories based on the growth rate with respect to the time horizon T : (i) small, distribution-dependent, bounds of order of magnitude $\ln(T)$ and (ii) robust, distribution-free, bounds of order of magnitude \sqrt{T} . The Bandits with Knapsacks model, an extension to the framework allowing to model resource consumption, lacks this clear-cut distinction. While several algorithms have been shown to achieve asymptotically optimal distribution-free bounds on regret, there has been little progress toward the development of small distribution-dependent regret bounds. We partially bridge the gap by designing a general-purpose algorithm with distribution-dependent regret bounds that are logarithmic in the initial endowments of resources in several important cases that cover many practical applications, including dynamic pricing with limited supply, bid optimization in online advertisement auctions, and dynamic procurement.

1. Introduction.

1.1. Motivation. Multi-Armed Bandit (MAB) is a benchmark model for repeated decision making in stochastic environments with very limited feedback on the outcomes of alternatives. In these circumstances, a decision maker must strive to find an overall optimal sequence of decisions while making as few suboptimal ones as possible when exploring the decision space in order to generate as much revenue as possible, a trade-off coined *exploration-exploitation*. The original problem, first formulated in its predominant version in Robbins [21], has spurred a new line of research that aims at introducing additional constraints that reflect more accurately the reality of the decision making process. *Bandits with Knapsacks* (BwK), a model formulated in its most general form in Badanidiyuru et al. [8], fits into this framework and is characterized by the consumption of a limited supply of resources (e.g. time, money, and natural resources) that comes with every decision. This extension is motivated by a number of applications in electronic markets such as dynamic pricing with limited supply, see Besbes and Zeevi [11] and Babaioff et al. [6], online advertising, see Slivkins [22], online bid optimization for sponsored search auctions, see Tran-Thanh et al. [26], and crowdsourcing, see Badanidiyuru et al. [7]. A unifying paradigm of online learning is to evaluate algorithms based on their regret performance. In the BwK theory, this performance criterion is expressed as the gap between the total payoff of an optimal oracle algorithm aware of how the rewards and the amounts of resource consumption are generated and the total payoff of the algorithm. Many approaches have been proposed to tackle the original MAB problem, where time is the only limited resource with a prescribed time

* Research funded in part by ONR, grants N00014-12-1-0033 and N00014-15-1-2083.

horizon T , and the optimal regret bounds are now well documented. They can be classified into two categories with qualitatively different asymptotic growth rates. Many algorithms, such as UCB1, see Auer et al. [4], Thompson sampling, see Agrawal and Goyal [3], and ϵ -greedy, see Auer et al. [4], achieve distribution-dependent, i.e. with constant factors that depend on the underlying unobserved distributions, asymptotic bounds on regret of order $\Theta(\ln(T))$, which is shown to be optimal in Lai and Robbins [19]. While these results prove very satisfying in many settings, the downside is that the bounds can get arbitrarily large if a malicious opponent was to select the underlying distributions in an adversarial fashion. In contrast, algorithms such as Exp3, designed in Auer et al. [5], achieve distribution-free bounds that can be computed in an online fashion, at the price of a less attractive growth rate $\Theta(\sqrt{T})$. The BwK theory lacks this clear-cut distinction. While provably optimal distribution-free bounds have recently been established, see Agrawal and Devanur [1] and Badanidiyuru et al. [8], there has been little progress toward the development of asymptotically optimal distribution-dependent regret bounds. To bridge the gap, we introduce a template algorithm with proven regret bounds which are asymptotically logarithmic in the initial supply of each resource, in four important cases that cover a wide range of applications:

- Case 1, where there is a single limited resource other than time, which is not limited, and the amount of resource consumed as a result of making a decision is stochastic. Applications in online advertising, see Tran-Thanh et al. [23], fit in this framework;
- Case 2, where there are arbitrarily many resources and the amounts of resources consumed as a result of making a decision are deterministic. Applications to network revenue management of perishable goods, see Besbes and Zeevi [11], shelf optimization of perishable goods, see Graczová and Jacko [15], and wireless sensor networks, see Tran-Thanh et al. [25], fit in this framework;
- Case 3, where there are two limited resources, one of which is assumed to be time while the consumption of the other is stochastic, under a nondegeneracy condition. Typical applications include online bid optimization in sponsored search auctions, see Tran-Thanh et al. [26], dynamic pricing with limited supply, see Babaioff et al. [6], and dynamic procurement, see Badanidiyuru et al. [7];
- Case 4, where there are arbitrarily many resources, under a stronger nondegeneracy condition than for Case 3. Typical applications include dynamic ad allocation, see Slivkins [22], dynamic pricing of multiple products, see Badanidiyuru et al. [8], and network revenue management, see Besbes and Zeevi [11].

In terms of applicability and significance of the results, Case 3 is the most important case. Case 4 is the most general one but the analysis is more involved and requires stronger assumptions which makes it less attractive from a practical standpoint. The analysis is easier for Cases 1 and 2 but their modeling power is more limited.

In fast-paced environments, such as in ad auctions, the stochastic assumptions at the core of the BwK model are only valid for a short period of time but there are typically a large number of actions to be performed per second (e.g. submit a bid for a new ad auction). In these situations, the initial endowments of resources are thus typically large and logarithmic regret bounds can be significantly more attractive than distribution-free ones.

1.2. Problem statement and contributions. At each time period $t \in \mathbb{N}$, a decision needs to be made among a predefined finite set of actions, represented by arms and labeled $k = 1, \dots, K$. We denote by a_t the arm pulled at time t . Pulling arm k at time t yields a random reward $r_{k,t} \in [0, 1]$ (after scaling) and incurs the consumption of $C \in \mathbb{N}$ different resource types by random amounts $c_{k,t}(1), \dots, c_{k,t}(C) \in [0, 1]^C$ (after scaling). Note that time itself may or may not be a limited resource. At any time t and for any arm k , the vector $(r_{k,t}, c_{k,t}(1), \dots, c_{k,t}(C))$ is jointly drawn from a fixed probability distribution ν_k independently from the past. The rewards and the amounts of resource consumption can be arbitrarily correlated across arms. We denote by $(\mathcal{F}_t)_{t \in \mathbb{N}}$ the natural filtration generated by the rewards and the amounts of resource consumption revealed to the decision maker, i.e. $((r_{a_t,t}, c_{a_t,t}(1), \dots, c_{a_t,t}(C)))_{t \in \mathbb{N}}$. The consumption of any resource $i \in \{1, \dots, C\}$ is constrained by an initial budget $B(i) \in \mathbb{R}_+$. As a result, the decision maker can

keep pulling arms only so long as he does not run out of any of the C resources and the game ends at time period τ^* , defined as:

$$\tau^* = \min\{t \in \mathbb{N} \mid \exists i \in \{1, \dots, C\}, \sum_{\tau=1}^t c_{a_\tau, \tau}(i) > B(i)\}. \quad (1)$$

Note that τ^* is a stopping time with respect to $(\mathcal{F}_t)_{t \geq 1}$. When it comes to choosing which arm to pull next, the difficulty for the decision maker lies in the fact that none of the underlying distributions, i.e. $(\nu_k)_{k=1, \dots, K}$, are initially known. Furthermore, the only feedback provided to the decision maker upon pulling arm a_t (but prior to selecting a_{t+1}) is $(r_{a_t, t}, c_{a_t, t}(1), \dots, c_{a_t, t}(C))$, i.e. the decision maker does not observe the rewards that would have been obtained and the amounts of resources that would have been consumed as a result of pulling a different arm. The goal is to design a non-anticipating algorithm that, at any time t , selects a_t based on the information acquired in the past so as to keep the pseudo regret defined as:

$$R_{B(1), \dots, B(C)} = \text{ER}_{\text{OPT}}(B(1), \dots, B(C)) - \mathbb{E}\left[\sum_{t=1}^{\tau^*-1} r_{a_t, t}\right], \quad (2)$$

as small as possible, where $\text{ER}_{\text{OPT}}(B(1), \dots, B(C))$ is the maximum expected sum of rewards that can be obtained by a non-anticipating oracle algorithm that has knowledge of the underlying distributions. Here, an algorithm is said to be non-anticipating if the decision to pull a given arm does not depend on the future observations. We develop algorithms and establish distribution-dependent regret bounds, that hold for any choice of the unobserved underlying distributions $(\nu_k)_{k=1, \dots, K}$, as well as distribution-independent regret bounds. This entails studying the asymptotic behavior of $R_{B(1), \dots, B(C)}$ when all the budgets $(B(i))_{i=1, \dots, C}$ go to infinity. In order to simplify the analysis, it is convenient to assume that the ratios $(B(i)/B(C))_{i=1, \dots, C}$ are constants independent of any other relevant quantities and to denote $B(C)$ by B .

ASSUMPTION 1. *For any resource $i \in \{1, \dots, C\}$, we have $B(i) = b(i) \cdot B$ for some fixed constant $b(i) \in (0, 1]$. Hence $b = \min_{i=1, \dots, C} b(i)$ is a positive quantity.*

When time is a limited resource, we use the notation T in place of B . Assumption 1 is widely used in the dynamic pricing literature where the inventory scales linearly with the time horizon, see Besbes and Zeevi [11] and Johnson et al. [17]. Assumption 1 will only prove useful when deriving distribution-dependent regret bounds and it can largely be relaxed, see Section A of the Appendix.

As the mean turns out to be an important statistics, we denote the mean reward and amounts of resource consumption by $\mu_k^r, \mu_k^c(1), \dots, \mu_k^c(C)$ and their respective empirical estimates by $\bar{r}_{k, t}, \bar{c}_{k, t}(1), \dots, \bar{c}_{k, t}(C)$. These estimates depend on the number of times each arm has been pulled by the decision maker up to, but not including, time t , which we write $n_{k, t}$. We end with a general assumption, which we use throughout the paper, meant to have the game end in finite time.

ASSUMPTION 2. *For any arm $k \in \{1, \dots, K\}$, we have $\max_{i=1, \dots, C} \mu_k^c(i) > 0$.*

Note that Assumption 2 is automatically satisfied if time is a limited resource.

Contributions. We design an algorithm that runs in time polynomial in K for which we establish $O(K^C \cdot \ln(B)/\Delta)$ (resp. $\sqrt{K^C \cdot B \cdot \ln(B)}$) distribution-dependent (resp. distribution-free) regret bounds, where Δ is a parameter that generalizes the optimality gap for the standard MAB problem. We establish these regret bounds in four cases of increasing difficulty making additional technical assumptions that become stronger as we make progress towards tackling the general case. We choose to present these intermediate cases since: (i) we get improved constant factors under weaker assumptions and (ii) they subsume many practical applications. Note that our distribution-dependent regret bounds scale as a polynomial function of K of degree C , which may be unacceptable when the number of resources is large. We provide evidence that suggests that a linear dependence on K can be achieved by tweaking the algorithm, at least in

some particular cases of interest. Finally, we point out that the constant factors hidden in the O notations are not scale-free, in the sense that jointly scaling down the amounts of resources consumed at each round along with their respective initial endowments worsens the bounds. As a consequence, initially scaling down the amounts of resource consumption in order to guarantee that they lie in $[0, 1]$ should be done with caution: the scaling factors should be as small as possible.

1.3. Literature review. The Bandits with Knapsacks framework was first introduced in its full generality in Badanidiyuru et al. [8], but special cases had been studied before, see for example Tran-Thanh et al. [23], Ding et al. [13], and Babaioff et al. [6]. Since the standard MAB problem fits in the BwK framework, with time being the only scarce resource, the results listed in the introduction tend to suggest that regret bounds with logarithmic growth with respect to the budgets may be possible for BwK problems but very few such results are documented. When there are arbitrarily many resources and a time horizon, Badanidiyuru et al. [8] and Agrawal and Devanur [1] obtain $\tilde{O}(\sqrt{K \cdot T})$ distribution-free bounds on regret that hold on average as well as with high probability, where the \tilde{O} notation hides logarithmic factors. These results were later extended to the contextual version of the problem in Badanidiyuru et al. [9] and Agrawal et al. [2]. Johnson et al. [17] extend Thompson sampling to tackle the general BwK problem and obtain distribution-dependent bounds on regret of order $\tilde{O}(\sqrt{T})$ (with an unspecified dependence on K) when time is a limited resource, under a nondegeneracy condition. Trovo et al. [27] develop algorithms for BwK problems with a continuum of arms and a single limited resource constrained by a budget B and obtain $o(B)$ regret bounds. Combes et al. [12] consider a closely related framework that allows to model any history-dependent constraint on the number of times any arm can be pulled and obtain $O(K \cdot \ln(T))$ regret bounds when time is a limited resource. However, the benchmark oracle algorithm used in Combes et al. [12] to define the regret is significantly weaker than the one considered here as it only has knowledge of the distributions of the rewards, as opposed to the joint distributions of the rewards and the amounts of resource consumption. Babaioff et al. [6] establish a $\Omega(\sqrt{T})$ distribution-dependent lower bound on regret for a dynamic pricing problem which can be cast as a BwK problem with a time horizon, a resource whose consumption is stochastic, and a continuum of arms. This lower bound does not apply here as we are considering finitely many arms and it is well known that the minimax regret can be exponentially smaller when we move from finitely many arms to uncountably many arms for the standard MAB problem, see Kleinberg and Leighton [18]. Tran-Thanh et al. [24] tackles BwK problems with a single limited resource whose consumption is deterministic and constrained by a global budget B and obtain $O(K \cdot \ln(B))$ regret bounds. This result was later extended to the case of a stochastic resource in Xia et al. [30]. Wu et al. [29] study a contextual version of the BwK problem when there are two limited resources, one of which is assumed to be time while the consumption of the other is deterministic, and obtain $O(K \cdot \ln(T))$ regret bounds under a nondegeneracy condition. Logarithmic regret bounds are also derived in Slivkins [22] for a dynamic ad allocation problem that can be cast as a BwK problem.

Organization. The remainder of the paper is organized as follows. We present applications of the BwK model in Section 2. We expose the algorithmic ideas underlying our approach in Section 3 and apply these ideas to Cases (1), (2), (3), and (4) in Sections 4, 5, 6, and 7 respectively. We choose to discuss each case separately in a self-contained fashion so that readers can delve into the setting they are most interested in. This comes at the price of some overlap in the analysis. We relax some of the assumptions made in the course of proving the regret bounds and discuss extensions in Section A of the Appendix. To provide as much intuition as possible, the ideas and key technical steps are all included in the main body, sometimes through proof sketches, while the technical details are deferred to the Appendix.

Notations. For a set S , $|S|$ denotes the cardinality of S while $\mathbb{1}_S$ is the indicator function of S . For a vector $x \in \mathbb{R}^n$ and S a subset of $\{1, \dots, n\}$, x_S refers to the subvector $(x_i)_{i \in S}$. For a square matrix A , $\det(A)$ is the determinant of A while $\text{adj}(A)$ denotes its adjugate. For $x \in \mathbb{R}$, x_+ is the positive part of x . We use standard asymptotic notations such as $O(\cdot)$, $o(\cdot)$, $\Omega(\cdot)$, and $\Theta(\cdot)$.

2. Applications. A number of applications of the BwK framework are documented in the literature. For the purpose of being self-contained, we review a few popular ones that satisfy our technical assumptions.

2.1. Online advertising.

Bid optimization in repeated second-price auctions. Consider a bidder participating in sealed second-price auctions who is willing to spend a budget B . This budget may be allocated only for a period of time (for the next T auctions) or until it is completely exhausted. Rounds, indexed by $t \in \mathbb{N}$, correspond to auctions the bidder participates in. If the bid submitted by the bidder for auction t is larger than the highest bid submitted by the competitors, denoted by m_t , the bidder wins the auction, derives a private utility $v_t \in [0, 1]$ (whose monetary value is typically difficult to assess), and is charged m_t . Otherwise, m_t is not revealed to the bidder and v_t cannot be assessed. We consider a stochastic setting where the environment and the competitors are not fully adversarial: $((v_t, m_t))_{t \in \mathbb{N}}$ is assumed to be an i.i.d. stochastic process. The goal for the bidder is to design a strategy to maximize the expected total utility derived given that the bidder has selected a grid of bids to choose from (b_1, \dots, b_K) (e.g. $b_1 = \$0.10, \dots, b_K = \1). This is a BwK problem with two resources: time and money. Pulling arm k at round t corresponds to bidding b_k in auction t , costs $c_{k,t} = m_t \cdot \mathbb{1}_{b_k \geq m_t}$, and yields a reward $r_{k,t} = v_t \cdot \mathbb{1}_{b_k \geq m_t}$. Weed et al. [28] design bidding strategies for a variant of this problem where the bidder is not limited by a budget and $r_{k,t} = (v_t - m_t) \cdot \mathbb{1}_{b_k \geq m_t}$.

This model was first formalized in Tran-Thanh et al. [26] in the context of sponsored search auctions. In sponsored search auctions, advertisers can bid on keywords to have ads (typically in the form of a link followed by a text description) displayed alongside the search results of a web search engine. When a user types a search query, a set of relevant ads are selected and an auction is run in order to determine which ones will be displayed. The winning ads are allocated to ad slots based on the outcome of the auction and, in the prevailing cost-per-click pricing scheme, their owners get charged only if the user clicks on their ads. Because the auction is often a variant of a sealed second-price auction (e.g. a generalized second-price auction), very limited feedback is provided to the advertiser if the auction is lost. In addition, both the demand and the supply cannot be predicted ahead of time and are thus commonly modeled as random variables, see Ghosh et al. [14]. For these reasons, bidding repeatedly on a keyword can be formulated as a BwK problem. In particular, when the search engine has a single ad slot per query, this problem can be modeled as above: B is the budget the advertiser is willing to spend on a predetermined keyword and rounds correspond to ad auctions the advertiser has been selected to participate in. If the advertiser wins the auction, his or her ad gets displayed and he or she derives a utility $v_t = \mathbb{1}_{A_t}$, where A_t is the event that the ad gets clicked on. The goal is to maximize the expected total number of clicks given the budget constraint. The advertiser may also be interested in optimizing the ad to be displayed, which will affect the probability of a click. In this case, the modeling is similar but arms correspond to pairs of bid values and ads.

Dynamic ad allocation. This problem was first modeled in the BwK framework in Slivkins [22]. A publisher, i.e. the owner of a collection of websites where ads can be displayed, has previously agreed with K advertisers, indexed by $k \in \{1, \dots, K\}$, on a predetermined cost-per-click p_k . Additionally, advertiser k is not willing to spend more than a prescribed budget, B_k , for a predetermined period of time (which corresponds to the next T visits or rounds). Denote by A_t^k the event that the ad provided by advertiser k gets clicked on at round t . We consider a stochastic setting where the visitors are not fully adversarial: $(\mathbb{1}_{A_t^k})_{t \in \mathbb{N}}$ is assumed to be an i.i.d. stochastic process for any advertiser k . The goal for the publisher is to maximize the total expected revenues by choosing which ad to display at every round, i.e. every time somebody visits one of the websites. This situation can be modeled as a BwK problem with $K + 1$ resources: time and money for each of the K advertisers. Pulling arm $k \in \{1, \dots, K\}$ at round t corresponds to displaying the ad owned by advertiser k , incurs the costs $c_{k,t}(i) = p_k \cdot \mathbb{1}_{A_t^k} \cdot \mathbb{1}_{i=k}$ to advertiser $i \in \{1, \dots, K\}$, and yields a revenue $r_{k,t} = p_k \cdot \mathbb{1}_{A_t^k}$.

2.2. Revenue management.

Dynamic pricing with limited supply. This BwK model was first proposed in Babaioff et al. [6]. An agent has B identical items to sell to T potential customers that arrive sequentially. Customer $t \in \{1, \dots, T\}$ is offered a take-it-or-leave-it price p_t and purchases the item only if p_t is no larger than his or her own valuation v_t , which is never disclosed. Customers are assumed to be non-strategic in the sense that their valuations are assumed to be drawn i.i.d. from a distribution unknown to the agent. The goal for the agent is to maximize the total expected revenues by offering prices among a predetermined list (p_1, \dots, p_K) . This is a BwK problem with two resources: time and item inventory. Pulling arm $k \in \{1, \dots, K\}$ at round t corresponds to offering the price p_k , depletes the inventory of $c_{k,t} = \mathbb{1}_{p_k \leq v_t}$ unit, and generates a revenue $r_{k,t} = p_k \cdot \mathbb{1}_{p_k \leq v_t}$. Badanidiyuru et al. [8] propose an extension where multiple units of M different products may be offered to a customer, which then buys as many as needed of each kind in order to maximize his or her own utility function. In this case, the modeling is similar but arms correspond to vectors of dimension $2M$ specifying the number of items offered along with the price tag for each product and there are $M + 1$ resources: time and item inventory for each of the M products.

Network revenue management.

Non-perishable goods. This is an extension of the dynamic pricing problem developed in Besbes and Zeevi [11] which is particularly suited for applications in the online retailer industry, e.g. the online fashion sample sales industry, see Johnson et al. [17]. Each product $m = 1, \dots, M$ is produced from a finite amount of C different kinds of raw materials (which may be products themselves). Producing one unit of product $m \in \{1, \dots, M\}$ consumes a deterministic amount of resource $i \in \{1, \dots, C\}$ denoted by $c_m(i)$. Customer $t \in \{1, \dots, T\}$ is offered a product $m_t \in \{1, \dots, M\}$ along with a take-it-or-leave-it price $p_t^{m_t}$ and purchases it if his or her valuation $v_t^{m_t}$ is larger than $p_t^{m_t}$. Products are manufactured online as customers order them. We assume that $((v_t^1, \dots, v_t^M))_{t \in \mathbb{N}}$ is an i.i.d. stochastic process with distribution unknown to the agent. This is a BwK problem with $C + 1$ resources: time and the initial endowment of each resource. Given a predetermined list of arms $((m_k, p_k))_{k=1, \dots, K}$, pulling arm k at round t corresponds to offering product m_k at the price p_k , incurs the consumption of resource i by an amount $c_{k,t}(i) = c_{m_k}(i) \cdot \mathbb{1}_{p_k \leq v_t^{m_k}}$, and generates a revenue $r_{k,t} = p_k \cdot \mathbb{1}_{p_k \leq v_t^{m_k}}$.

Perishable goods. This is a variant of the last model developed for perishable goods, with applications in the food retail industry and the newspaper industry. At each time period $t \in \{1, \dots, T\}$, a retailer chooses how many units $\lambda_t^m \in \mathbb{N}$ of product $m \in \{1, \dots, M\}$ to manufacture along with a price offer for it p_t^m . At time t , the demand for product m sold at the price p is a random quantity denoted by $d_t^m(p)$. We assume that customers are non-strategic: for any vector of prices (p_1, \dots, p_M) , $((d_t^1(p_1), \dots, d_t^M(p_M)))_{t \in \mathbb{N}}$ is an i.i.d. stochastic process with distribution unknown to the agent. Products perish at the end of each round irrespective of whether they have been purchased. Given a predetermined list of arms $((\lambda_k^1, p_k^1, \dots, \lambda_k^M, p_k^M))_{k=1, \dots, K}$, pulling arm k at round t corresponds to offering λ_k^m units of product m at the price p_k^m for any $m \in \{1, \dots, M\}$, incurs the consumption of resource i by a deterministic amount $c_{k,t}(i) = \sum_{m=1}^M \lambda_k^m \cdot c_m(i)$ (where $c_m(i)$ is defined in the previous paragraph), and generates a revenue $r_{k,t} = \sum_{m=1}^M p_k^m \cdot \min(d_t^m(p_k^m), \lambda_k^m)$.

Shelf optimization for perishable goods. This is a variant of the model introduced in Graczová and Jacko [15]. Consider a retailer who has an unlimited supply of M different types of products. At each time period t , the retailer has to decide how many units, λ_t^m , of each product, $m \in \{1, \dots, M\}$, to allocate to a promotion space given that at most N items fit in the limited promotion space. Moreover, the retailer also has to decide on a price tag p_t^m for each product m . All units of product $m \in \{1, \dots, M\}$ perish by time period T_m and the retailer is planning the allocation for the next T time periods. At round t , the demand for product m is a random quantity denoted by $d_t^m(p)$. Customers are non-strategic: for any vector of prices (p_1, \dots, p_M) , $((d_t^1(p_1), \dots, d_t^M(p_M)))_{t \in \mathbb{N}}$ is an i.i.d. stochastic process with distribution unknown to the agent. This is a BwK problem with $M + 1$ resources: time horizon and time after which each product perishes. Given a predetermined list of arms $((\lambda_k^1, p_k^1, \dots, \lambda_k^M, p_k^M))_{k=1, \dots, K}$ satisfying $\sum_{m=1}^M \lambda_k^m \leq K$ for any $k \in \{1, \dots, K\}$, pulling arm k at round t corresponds to allocating λ_k^m units of product m to

the promotion space for every $m \in \{1, \dots, M\}$ with the respective price tags (p_k^1, \dots, p_k^M) , incurs the consumption of resource i by a deterministic amount $c_{k,t}(i) = 1$, and generates a revenue $r_{k,t} = \sum_{m=1}^M p_k^m \cdot \min(d_t^m(p_k^m), \lambda_k^m)$.

2.3. Dynamic procurement. This problem was first studied in Badanidiyuru et al. [7]. Consider a buyer with a budget B facing T agents arriving sequentially, each interested in selling one good. Agent $t \in \{1, \dots, T\}$ is offered a take-it-or-leave-it price, p_t , and makes a sell only if the value he or she attributes to the item, v_t , is no larger than p_t . We consider a stochastic setting where the sellers are not fully adversarial: $(v_t)_{t \in \mathbb{N}}$ is an i.i.d. stochastic process with distribution unknown to the buyer. The goal for the buyer is to maximize the total expected number of goods purchased by offering prices among a predetermined list (p_1, \dots, p_K) . This is a BwK problem with two resources: time and money. Pulling arm k at round t corresponds to offering the price p_k , incurs a cost $c_{k,t} = p_k \cdot \mathbb{1}_{p_k \geq v_t}$, and yields a reward $r_{k,t} = \mathbb{1}_{p_k \geq v_t}$. It is also possible to model situations where the agents are selling multiple types of products and/or multiple units, in which case arms correspond to vectors specifying the number of units of each product required along with their respective prices, see Badanidiyuru et al. [8].

Applications of this model to crowdsourcing platforms are described in Badanidiyuru et al. [7] and Badanidiyuru et al. [8]. In this setting, agents correspond to workers that are willing to carry out microtasks which are submitted by buyers (called “requesters”) using a posted-price mechanism. Requesters are typically submitting large batches of jobs and can thus adjust the posted prices as they learn about the pool of workers.

2.4. Wireless sensor networks. This is a variant of the model introduced in Tran-Thanh et al. [25]. Consider an agent collecting information using a network of wireless sensors powered by batteries. Activating sensor $k \in \{1, \dots, K\}$ consumes some amount of energy, c_k , which is depleted from the sensor’s initial battery level, B_k , and triggers a measurement providing a random amount of information (measured in bits), $r_{k,t}$, which is transmitted back to the agent. Sensors cannot harvest energy and the goal for the agent is to maximize the total expected amount of information collected over T actions. This is a BwK problem with $K + 1$ resources: time and the energy stored in the battery of each sensor. Pulling arm $k \in \{1, \dots, K\}$ corresponds to activating sensor k , incurs the consumption of resource $i \in \{1, \dots, K\}$ by a deterministic amount $c_{k,t} = c_k \cdot \mathbb{1}_{k=i}$, and yields a random reward $r_{k,t}$.

3. Algorithmic ideas.

3.1. Preliminaries. To handle the exploration-exploitation trade-off, an approach that has proved to be particularly successful hinges on the *optimism in the face of uncertainty* paradigm. The idea is to consider all plausible scenarios consistent with the information collected so far and to select the decision that yields the most revenue among all the scenarios identified. Concentration inequalities are intrinsic to the paradigm as they enable the development of systematic closed form confidence intervals on the quantities of interest, which together define a set of plausible scenarios. We make repeated use of the following result.

LEMMA 1. *Hoeffding’s inequality* Consider X_1, \dots, X_n n random variables with support in $[0, 1]$. If $\forall t \leq n \mathbb{E}[X_t | X_1, \dots, X_{t-1}] \leq \mu$, then $\mathbb{P}[X_1 + \dots + X_n \geq n\mu + a] \leq \exp(-\frac{2a^2}{n}) \quad \forall a \geq 0$.
If $\forall t \leq n \mathbb{E}[X_t | X_1, \dots, X_{t-1}] \geq \mu$, then $\mathbb{P}[X_1 + \dots + X_n \leq n\mu - a] \leq \exp(-\frac{2a^2}{n}) \quad \forall a \geq 0$.

Auer et al. [4] follow the *optimism in the face of uncertainty* paradigm to develop the Upper Confidence Bound algorithm (UCB1). UCB1 is based on the following observations: (i) the optimal strategy always consists in pulling the arm with the highest mean reward when time is the only limited resource, (ii) informally, Lemma 1 shows that $\mu_k^* \in [\bar{r}_{k,t} - \epsilon_{k,t}, \bar{r}_{k,t} + \epsilon_{k,t}]$ at time t with probability at least $1 - 2/t^3$ for $\epsilon_{k,t} = \sqrt{2 \ln(t) / n_{k,t}}$, irrespective of the number of times arm k has been pulled. Based on these observations, UCB1 always selects the arm with highest UCB index, i.e. $a_t \in \arg \max_{k=1, \dots, K} I_{k,t}$, where the UCB

index of arm k at time t is defined as $I_{k,t} = \bar{r}_{k,t} + \epsilon_{k,t}$. The first term can be interpreted as an exploitation term, the ultimate goal being to maximize revenue, while the second term is an exploration term, the smaller $n_{k,t}$, the bigger it is. This fruitful paradigm goes well beyond this special case and many extensions of UCB1 have been designed to tackle variants of the MAB problem, see for example Slivkins [22]. Agrawal and Devanur [1] embrace the same ideas to tackle BwK problems. The situation is more complex in this all-encompassing framework as the optimal oracle algorithm involves pulling several arms. In fact, finding the optimal pulling strategy given the knowledge of the underlying distributions is already a challenge in its own, see Papadimitriou and Tsitsiklis [20] for a study of the computational complexity of similar problems. This raises the question of how to evaluate $\text{ER}_{\text{OPT}}(B(1), \dots, B(C))$ in (2). To overcome this issue, Badanidiyuru et al. [8] upper bound the total expected payoff of any non-anticipating algorithm by the optimal value of a linear program, which is easier to compute.

LEMMA 2. *Adapted from Badanidiyuru et al. [8]*

The total expected payoff of any non-anticipating algorithm is no greater than B times the optimal value of the linear program:

$$\begin{aligned} & \sup_{(\xi_k)_{k=1, \dots, K}} \sum_{k=1}^K \mu_k^r \cdot \xi_k \\ & \text{subject to} \quad \sum_{k=1}^K \mu_k^c(i) \cdot \xi_k \leq b(i), \quad i = 1, \dots, C \\ & \quad \quad \quad \xi_k \geq 0, \quad k = 1, \dots, K. \end{aligned} \tag{3}$$

plus the constant term $\max_{\substack{k=1, \dots, K \\ i=1, \dots, C \\ \text{with } \mu_k^c(i) > 0}} \frac{\mu_k^r}{\mu_k^c(i)}$.

The optimization problem (3) can be interpreted as follows. For any arm k , $B \cdot \xi_k$ corresponds to the expected number of times arm k is pulled by the optimal algorithm. Hence, assuming we introduce a dummy arm 0 which is equivalent to skipping the current round, ξ_k can be interpreted as the probability of pulling arm k at any round when there is a time horizon T . Observe that the constraints restrict the feasible set of expected number of pulls by imposing that the amounts of resources consumed are no greater than their respective budgets in expectations, as opposed to almost surely which would be a more stringent constraint. This explains why the optimal value of (3) is larger than the maximum achievable payoff. In this paper, we use standard linear programming notions such as the concept of a basis and a basic feasible solution. We refer to Bertsimas and Tsitsiklis [10] for an introduction to linear programming. A pseudo-basis x is described by two subsets $\mathcal{K}_x \subset \{1, \dots, K\}$ and $\mathcal{C}_x \subset \{1, \dots, C\}$ such that $|\mathcal{K}_x| = |\mathcal{C}_x|$. A pseudo-basis x is a basis for (3) if the matrix $A_x = (\mu_k^c(i))_{(i,k) \in \mathcal{C}_x \times \mathcal{K}_x}$ is invertible. Furthermore, x is said to be a feasible basis for (3) if the corresponding basic solution, denoted by $(\xi_k^x)_{k=1, \dots, K}$ and determined by $\xi_k^x = 0$ for $k \notin \mathcal{K}_x$ and $A_x \xi_{\mathcal{K}_x}^x = b_{\mathcal{C}_x}$ (where $b_{\mathcal{C}_x}$ is the subvector $(b(i))_{i \in \mathcal{C}_x}$), is feasible for (3). When x is a feasible basis for (3), we denote by $\text{obj}_x = \sum_{k=1}^K \mu_k^r \cdot \xi_k^x$ its objective function. From Lemma 2, we derive:

$$R_{B(1), \dots, B(C)} \leq B \cdot \text{obj}_{x^*} - \mathbb{E} \left[\sum_{t=1}^{\tau^*} r_{a_t, t} \right] + O(1), \tag{4}$$

where x^* is an optimal feasible basis for (3). For mathematical convenience, we consider that the game carries on even if one of the resources is already exhausted so that a_t is well defined for any $t \in \mathbb{N}$. Of course, the rewards obtained for $t \geq \tau^*$ are not taken into account in the decision maker's payoff when establishing regret bounds.

3.2. Solution methodology. Lemma 2 also provides insight into designing algorithms. The idea is to incorporate confidence intervals on the mean rewards and the mean amounts of resource consumption into the offline optimization problem (3) and to base the decision upon the resulting optimal solution. There are several ways to carry out this task, each leading to a different algorithm. When there is a time horizon T , Agrawal and Devanur [1] use high-probability lower (resp. upper) bounds on the mean amounts of resource consumption (resp. rewards) in place of the unknown mean values in (3) and pull an arm at random according to the resulting optimal distribution. Specifically, at any round t , the authors suggest to compute an optimal solution $(\xi_{k,t}^*)_{k=1,\dots,K}$ to the linear program:

$$\begin{aligned} & \sup_{(\xi_k)_{k=1,\dots,K}} \sum_{k=1}^K (\bar{r}_{k,t} + \epsilon_{k,t}) \cdot \xi_k \\ & \text{subject to} \quad \sum_{k=1}^K (\bar{c}_{k,t}(i) - \epsilon_{k,t}) \cdot \xi_k \leq (1 - \gamma) \cdot b(i), \quad i = 1, \dots, C - 1 \\ & \quad \quad \quad \sum_{k=1}^K \xi_k \leq 1 \\ & \quad \quad \quad \xi_k \geq 0, \quad k = 1, \dots, K, \end{aligned} \tag{5}$$

for a well-chosen $\gamma \in (0, 1)$, and then to pull arm k with probability $\xi_{k,t}^*$ or skip the round with probability $1 - \sum_{k=1}^K \xi_{k,t}^*$. If we relate this approach to UCB1, the intuition is clear: the idea is to be optimistic about both the rewards and the amounts of resource consumption. We argue that this approach cannot yield logarithmic regret bounds. First, because γ has to be of order $1/\sqrt{T}$. Second, because, even if we were given an optimal solution to (3), $(\xi_k^*)_{k=1,\dots,K}$, before starting the game, consistently choosing which arm to pull at random according to this distribution at every round would incur regret $\Omega(\sqrt{T})$, as we next show.

LEMMA 3. *For all the cases treated in this paper, pulling arm k with probability $\xi_k^{x^*}$ at any round t yields a regret of order $\Omega(\sqrt{T})$ unless pulling any arm in $\{k \in \{1, \dots, K\} \mid \xi_k^{x^*} > 0\}$ incurs the same deterministic amount of resource consumption for all resources in \mathcal{C}_{x^*} and for all rounds $t \in \mathbb{N}$.*

Proof. We use the shorthand notation $\mathcal{K}_* = \{k \in \{1, \dots, K\} \mid \xi_k^{x^*} > 0\}$. Observe that, for any resource $i \in \{1, \dots, C - 1\}$, we have:

$$\begin{aligned} T \cdot \text{obj}_{x^*} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t,t}\right] &= (T - \mathbb{E}[\tau^*]) \cdot \text{obj}_{x^*} \\ &\geq \mathbb{E}\left[\left(\sum_{t=\tau^*}^T c_{a_t,t}(i) + \sum_{t=1}^{\tau^*-1} c_{a_t,t}(i) - B(i)\right)_+\right] \cdot \text{obj}_{x^*} + O(1) \\ &= \mathbb{E}\left[\left(\sum_{t=1}^T \{c_{a_t,t}(i) - b(i)\}\right)_+\right] \cdot \text{obj}_{x^*} + O(1), \end{aligned}$$

where the inequality is derived using $c_{a_t,t}(i) \leq 1$ for all rounds t and $\sum_{t=1}^{\tau^*-1} c_{a_t,t}(i) \leq B(i)$. Since, for $i \in \mathcal{C}_{x^*}$, $(c_{a_t,t}(i))_{t \in \mathbb{N}}$ is an i.i.d. bounded stochastic process with mean $b(i)$, we have:

$$\mathbb{E}\left[\left(\sum_{t=1}^T \{c_{a_t,t}(i) - b(i)\}\right)_+\right] = \Omega(\sqrt{T}), \tag{6}$$

provided that $c_{a_t,t}(i)$ has positive variance, which is true if there exists at least one arm $k \in \mathcal{K}_*$ such that $c_{k,t}(i)$ has positive variance or if there exist two arms $k, l \in \mathcal{K}_*$ such that $c_{k,t}(i)$ and $c_{l,t}(i)$ are not almost surely equal to the same deterministic value. Strictly speaking, this is not enough to conclude that $R_{B(1), \dots, B(C-1), T} = \Omega(\sqrt{T})$ as $T \cdot \text{obj}_{x^*}$ is only an upper bound on the maximum total expected payoff.

However, in Sections 4, 5, 6, and 7, we show that there exists an algorithm that satisfies $T \cdot \text{obj}_{x^*} - \mathbb{E}[\sum_{t=1}^* r_{a_t,t}] = O(\ln(T))$ for all the cases considered in this paper. This result, together with (6) and Lemma 2, implies that the regret incurred when pulling arm k with probability $\xi_k^{x^*}$ at any round is $\Omega(\sqrt{T})$.

The fundamental shortcoming of this approach is that it systematically leads us to plan to consume the same average amount of resource i per round $b(i)$, for any resource $i = 1, \dots, C-1$, irrespective of whether we have significantly over- or under-consumed in the past. Based on this observation, a natural idea is to solve the linear program:

$$\begin{aligned} & \sup_{(\xi_k)_{k=1,\dots,K}} \sum_{k=1}^K (\bar{r}_{k,t} + \epsilon_{k,t}) \cdot \xi_k \\ & \text{subject to} \quad \sum_{k=1}^K (\bar{c}_{k,t}(i) - \epsilon_{k,t}) \cdot \xi_k \leq (1 - \gamma) \cdot b_t(i), \quad i = 1, \dots, C-1 \\ & \quad \quad \quad \sum_{k=1}^K \xi_k \leq 1 \\ & \quad \quad \quad \xi_k \geq 0, \quad k = 1, \dots, K, \end{aligned} \tag{7}$$

instead of (5), where $b_t(i)$ denotes the ratio of the remaining amount of resource i at time t to the remaining time horizon, i.e. $T - t + 1$. Bounding the regret incurred by this adaptive algorithm is, however, difficult from a theoretical standpoint. To address this issue, we propose the following family of algorithms, whose behaviors are similar to the adaptive algorithm but lend themselves to an easier analysis.

Algorithm: UCB-Simplex

Take $\lambda \geq 1$ and $(\eta_i)_{i=1,\dots,C} \geq 0$ (these quantities will need to be carefully chosen). The algorithm is preceded by an initialization phase which consists in pulling each arm a given number of times, to be specified. For each subsequent time period t , proceed as follows.

Step-Simplex: Find an optimal basis x_t to the linear program:

$$\begin{aligned} & \sup_{(\xi_k)_{k=1,\dots,K}} \sum_{k=1}^K (\bar{r}_{k,t} + \lambda \cdot \epsilon_{k,t}) \cdot \xi_k \\ & \text{subject to} \quad \sum_{k=1}^K (\bar{c}_{k,t}(i) - \eta_i \cdot \epsilon_{k,t}) \cdot \xi_k \leq b(i), \quad i = 1, \dots, C \\ & \quad \quad \quad \xi_k \geq 0, \quad k = 1, \dots, K \end{aligned} \tag{8}$$

Adapting the notations, x_t is described by two subsets $\mathcal{K}_{x_t} \subset \{1, \dots, K\}$ and $\mathcal{C}_{x_t} \subset \{1, \dots, C\}$ such that $|\mathcal{K}_{x_t}| = |\mathcal{C}_{x_t}|$, the matrix $\bar{A}_{x_t,t} = (\bar{c}_{k,t}(i) - \eta_i \cdot \epsilon_{k,t})_{(i,k) \in \mathcal{C}_{x_t} \times \mathcal{K}_{x_t}}$, and the corresponding basic feasible solution $(\xi_{k,t}^{x_t})_{k=1,\dots,K}$ determined by $\xi_{k,t}^{x_t} = 0$ for $k \notin \mathcal{K}_{x_t}$ and $\bar{A}_{x_t,t} \xi_{\mathcal{K}_{x_t},t}^{x_t} = b_{\mathcal{C}_{x_t}}$.

Step-Load-Balance: Identify the arms involved in the optimal basis, i.e. \mathcal{K}_{x_t} . There are at most $\min(K, C)$ such arms. Use a load balancing algorithm \mathcal{A}_{x_t} , to be specified, to determine which of these arms to pull.

For all the cases considered in this paper, (8) is always bounded and Step-Simplex is well defined. The Simplex algorithm is an obvious choice to carry out Step-Simplex, especially when $\eta_i = 0$ for any resource $i \in \{1, \dots, C\}$, because, in this case, we only have to update one column of the constraint matrix per round which makes warm-starting properties attractive. However, note that this can also be done in time polynomial in K and C , see Grötschel et al. [16]. If we compare (8) with (5), the idea remains to be overly optimistic but, as we will see, more about the rewards than the amounts of resource consumption through the exploration factor λ which will typically be larger than η_i , thus transferring most of the burden of exploration from the constraints to the objective function. The details of Step-Load-Balance are purposefully left out

and will be specified for each of the cases treated in this paper. When there is a time horizon T , the general idea is to determine, at any time period t and for each resource $i = 1, \dots, C$, whether we have over- or under-consumed in the past and to perturb the probability distribution $(\xi_{k,t}^{x,t})_{k=1, \dots, K}$ accordingly to get back on track.

The algorithm we propose is intrinsically tied to the existence of basic feasible optimal solutions to (3) and (8). We denote by \mathcal{B} (resp. \mathcal{B}_t) the subset of bases of (3) (resp. (8)) that are feasible for (3) (resp. (8)). Step-Simplex can be interpreted as an extension of the index-based decision rule of UCB1. Indeed, Step-Simplex assigns an index $I_{x,t}$ to each basis $x \in \mathcal{B}_t$ and outputs $x_t \in \arg \max_{x \in \mathcal{B}_t} I_{x,t}$, where $I_{x,t} = \text{obj}_{x,t} + E_{x,t}$ with a clear separation (at least when $\eta_i = 0$ for any resource i) between the exploitation term, $\text{obj}_{x,t} = \sum_{k=1}^K \xi_{k,t}^x \cdot \bar{r}_{k,t}$, and the exploration term, $E_{x,t} = \lambda \cdot \sum_{k=1}^K \xi_{k,t}^x \cdot \epsilon_{k,t}$. Observe that, for $x \in \mathcal{B}_t$ that is also feasible for (3), $(\xi_{k,t}^x)_{k=1, \dots, K}$ and $\text{obj}_{x,t}$ are plug-in estimates of $(\xi_k^x)_{k=1, \dots, K}$ and obj_x when $\eta_i = 0$ for any resource i . Also note that when $\lambda = 1$ and $\eta_i = 0$ for any resource i and when time is the only limited resource, UCB-Simplex is identical to UCB1 as Step-Load-Balance is unambiguous in this special case, each basis involving a single arm. For any $x \in \mathcal{B}$, we define $\Delta_x = \text{obj}_{x^*} - \text{obj}_x \geq 0$ as the optimality gap. A feasible basis x is said to be suboptimal if $\Delta_x > 0$. At any time t , $n_{x,t}$ denotes the number of times basis x has been selected at Step-Simplex up to time t while $n_{k,t}^x$ denotes the number of times arm k has been pulled up to time t when selecting x at Step-Simplex. For all the cases treated in this paper, we will show that, under a nondegeneracy assumption, Step-Simplex guarantees that a suboptimal basis cannot be selected more than $O(\ln(B))$ times on average, a result reminiscent of the regret analysis of UCB1 carried out in Auer et al. [4]. However, in stark contrast with the situation of a single limited resource, this is merely a prerequisite to establish a $O(\ln(B))$ bound on regret. Indeed, a low regret algorithm must also balance the load between the arms as closely as possible to optimality. Hence, the choice of the load balancing algorithms \mathcal{A}_x is crucial to obtain logarithmic regret bounds.

4. A single limited resource. In this section, we tackle the case of a single resource whose consumption is limited by a global budget B , i.e. $C = 1$ and $b(1) = 1$. To simplify the notations, we omit the indices identifying the resources as there is only one, i.e. we write μ_k^c , $c_{k,t}$, $\bar{c}_{k,t}$, and η as opposed to $\mu_k^c(1)$, $c_{k,t}(1)$, $\bar{c}_{k,t}(1)$, and η_1 . We also use the shorthand $\epsilon = \min_{k=1, \dots, K} \mu_k^c$. Recall that, under Assumption 2, ϵ is positive and a priori unknown to the decision maker. In order to derive logarithmic bounds, we will also need to assume that the decision maker knows an upper bound on the optimal value of (3).

ASSUMPTION 3. *The decision maker knows $\kappa \geq \max_{k=1, \dots, K} \frac{\mu_k^r}{\mu_k^c}$ ahead of round 1.*

Assumption 3 is natural in repeated second-price auctions, as detailed in the last paragraph of this section. Moreover, note that if ϵ happens to be known ahead of round 1 we can take $\kappa = 1/\epsilon$.

Specification of the algorithm. We implement UCB-Simplex with $\lambda = 1 + \kappa$ and $\eta = 0$. The initialization step consists in pulling each arm until the amount of resource consumed as a result of pulling that arm is non-zero. The purpose of this step is to have $\bar{c}_{k,t} > 0$ for all periods to come and for all arms. Step-Load-Balance is unambiguous here as basic feasible solutions involve a single arm. Hence, we identify a basis x such that $\mathcal{K}_x = \{k\}$ and $\mathcal{C}_x = \{1\}$ with the corresponding arm and write $x = k$ to simplify the notations. In particular, $k^* \in \{1, \dots, K\}$ identifies an optimal arm in the sense defined in Section 3. For any arm k , the exploration and exploitation terms defined in Section 3 specialize to:

$$\text{obj}_{k,t} = \frac{\bar{r}_{k,t}}{\bar{c}_{k,t}} \text{ and } E_{k,t} = (1 + \kappa) \cdot \frac{\epsilon_{k,t}}{\bar{c}_{k,t}},$$

while $\text{obj}_k = \mu_k^r / \mu_k^c$, so that:

$$k^* \in \arg \max_{k=1, \dots, K} \frac{\mu_k^r}{\mu_k^c}, \quad a_t \in \arg \max_{k=1, \dots, K} \frac{\bar{r}_{k,t} + (1 + \kappa) \cdot \epsilon_{k,t}}{\bar{c}_{k,t}}, \quad \text{and } \Delta_k = \frac{\mu_{k^*}^r}{\mu_{k^*}^c} - \frac{\mu_k^r}{\mu_k^c}.$$

We point out that, for the particular setting considered in this section, UCB-Simplex is almost identical to the fractional KUBE algorithm proposed in Tran-Thanh et al. [24] to tackle the case of a single resource whose consumption is deterministic. It only differs by the presence of the scaling factor $1 + \kappa$ to favor exploration over exploitation, which becomes unnecessary when the amounts of resource consumed are deterministic, see Section A of the Appendix.

Regret analysis. We omit the initialization step in the theoretical analysis because the amount of resource consumed is $O(1)$ and the reward obtained is non-negative and not taken into account in the decision maker's total payoff. Moreover, the initialization step ends in finite time almost surely as a result of Assumption 2. First observe that (4) specializes to:

$$R_B \leq B \cdot \frac{\mu_{k^*}^r}{\mu_{k^*}^c} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t,t}\right] + O(1). \quad (9)$$

To bound the right-hand side, we start by estimating the expected time horizon.

LEMMA 4. *For any non-anticipating algorithm, we have: $\mathbb{E}[\tau^*] \leq \frac{B+1}{\epsilon}$.*

Sketch of proof. By definition of τ^* , we have $\sum_{t=1}^{\tau^*-1} c_{a_t,t} \leq B$. Taking expectations on both sides yields $B \geq \mathbb{E}[\sum_{t=1}^{\tau^*} \mu_{a_t}^c] - 1 \geq \mathbb{E}[\tau^*] \cdot \epsilon - 1$ by Assumption 2. Rearranging this last inequality yields the claim.

The next result is crucial. Used in combination with Lemma 4, it shows that any suboptimal arm is pulled at most $O(\ln(B))$ times in expectations, a well-known result for UCB1, see Auer et al. [4]. The proof is along the same lines as for UCB1, namely we assume that arm k has already been pulled more than $\Theta(\ln(\tau^*)/(\Delta_k)^2)$ times and conclude that arm k cannot be pulled more than a few more times, with the additional difficulty of having to deal with the random stopping time and the fact that the amount of resource consumed at each step is stochastic.

LEMMA 5. *For any suboptimal arm k , we have:*

$$\mathbb{E}[n_{k,\tau^*}] \leq 2^6 \left(\frac{\lambda}{\mu_k^c}\right)^2 \cdot \frac{\mathbb{E}[\ln(\tau^*)]}{(\Delta_k)^2} + \frac{4\pi^2}{3\epsilon^2}.$$

Sketch of proof. We use the shorthand notation $\beta_k = 2^5(\lambda/\mu_k^c)^2 \cdot (1/\Delta_k)^2$. First observe that if we want to bound $\mathbb{E}[n_{k,\tau^*}]$, we may assume, without loss of generality, that arm k has been pulled at least $\beta_k \cdot \ln(t)$ times at any time t up to an additive term of $2\beta_k \cdot \mathbb{E}[\ln(\tau^*)]$ in the final inequality. We then just have to bound by a constant the probability that k is selected at any time t given that $n_{k,t} \geq \beta_k \cdot \ln(t)$. If k is selected at time t , it must be that k is optimal for (8), which, in particular, implies that $\text{obj}_{k,t} + E_{k,t} \geq \text{obj}_{k^*,t} + E_{k^*,t}$. This can only happen if either: (i) $\text{obj}_{k,t} \geq \text{obj}_k + E_{k,t}$, i.e. the objective value of k is overly optimistic, (ii) $\text{obj}_{k^*,t} \leq \text{obj}_{k^*} - E_{k^*,t}$, i.e. the objective value of k^* is overly pessimistic, or (iii) $\text{obj}_{k^*} < \text{obj}_k + 2E_{k,t}$, i.e. the optimality gap of arm k is small compared to its exploration factor. The probability of events (i) and (ii) can be bounded by $\sim 1/t^2$ in the same fashion, irrespective of how many times these arms have been pulled in the past. For example for event (i), this is because if $\bar{r}_{k,t}/\bar{c}_{k,t} = \text{obj}_{k,t} \geq \text{obj}_k + E_{k,t} = \mu_k^r/\mu_k^c + E_{k,t}$, then either (a) $\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}$ or (b) $\bar{c}_{k,t} \leq \mu_k^c - \epsilon_{k,t}$ and both of these events have probability at most $\sim 1/t^2$ by Lemma 1. Indeed, if (a) and (b) do not hold, we have:

$$\begin{aligned} \frac{\bar{r}_{k,t}}{\bar{c}_{k,t}} - \frac{\mu_k^r}{\mu_k^c} &= \frac{(\bar{r}_{k,t} - \mu_k^r)\mu_k^c + (\mu_k^c - \bar{c}_{k,t})\mu_k^r}{\bar{c}_{k,t} \cdot \mu_k^c} \\ &< \frac{\epsilon_{k,t}}{\bar{c}_{k,t}} + \frac{\epsilon_{k,t}}{\bar{c}_{k,t}} \cdot \frac{\mu_k^r}{\mu_k^c} \leq (1 + \kappa) \cdot \frac{\epsilon_{k,t}}{\bar{c}_{k,t}} = E_{k,t}. \end{aligned}$$

As for event (iii), observe that if $\text{obj}_{k^*} < \text{obj}_k + 2E_{k,t}$ and $n_{k,t} \geq \beta_k \cdot \ln(t)$ then we have $\bar{c}_{k,t} \leq \mu_k^c/2$, which happens with probability at most $\sim 1/t^2$ by Lemma 1 given that arm k has already been pulled at least $\sim \ln(t)/(\mu_k^c)^2$ times.

Building on the last two results, we derive a distribution-dependent regret bound which improves upon the one derived in Xia et al. [30]: the decision maker is only assumed to know κ , as opposed to a lower bound on ϵ , ahead of round 1. This is more natural in bidding applications as detailed in the last paragraph of this section. This bound generalizes the one obtained by Auer et al. [4] when time is the only scarce resource.

THEOREM 1. *We have:*

$$R_B \leq 2^6 \lambda^2 \cdot \left(\sum_{k \in \{1, \dots, K\} \mid \Delta_k > 0} \frac{1}{\mu_k^c \cdot \Delta_k} \right) \cdot \ln\left(\frac{B+1}{\epsilon}\right) + O(1).$$

Sketch of proof. We build upon (9):

$$\begin{aligned} R_B &\leq B \cdot \frac{\mu_{k^*}^r}{\mu_{k^*}^c} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t, t}\right] + O(1) \\ &= B \cdot \frac{\mu_{k^*}^r}{\mu_{k^*}^c} - \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[n_{k, \tau^*}] + O(1) \\ &= \frac{\mu_{k^*}^r}{\mu_{k^*}^c} \cdot \left(B - \sum_{k \mid \Delta_k = 0} \mu_k^c \cdot \mathbb{E}[n_{k, \tau^*}]\right) - \sum_{k \mid \Delta_k > 0} \mu_k^r \cdot \mathbb{E}[n_{k, \tau^*}] + O(1). \end{aligned}$$

By definition of τ^* , the resource is exhausted at time τ^* , i.e. $B \leq \sum_{t=1}^{\tau^*} c_{a_t, t}$. Taking expectations on both sides yields $B \leq \sum_{k=1}^K \mu_k^c \cdot \mathbb{E}[n_{k, \tau^*}]$. Plugging this last inequality back into the regret bound, we get:

$$R_B \leq \sum_{k \mid \Delta_k > 0} \mu_k^c \cdot \Delta_k \cdot \mathbb{E}[n_{k, \tau^*}] + O(1).$$

Using the upper bound of Lemma 4, the concavity of the logarithmic function, and Lemma 5, we derive:

$$R_B \leq 2^6 \lambda^2 \cdot \left(\sum_{k \mid \Delta_k > 0} \frac{1}{\mu_k^c \cdot \Delta_k} \right) \cdot \ln\left(\frac{B+1}{\epsilon}\right) + \frac{4\pi^2}{3\epsilon^2} \cdot \left(\sum_{k \mid \Delta_k > 0} \mu_k^c \cdot \Delta_k \right) + O(1)$$

which yields the claim since $\Delta_k \leq \mu_{k^*}^r / \mu_{k^*}^c \leq \kappa$ and $\mu_k^c \leq 1$ for any arm k .

Observe that the set of optimal arms, namely $\arg \max_{k=1, \dots, K} \mu_k^r / \mu_k^c$, does not depend on B and that Δ_k is a constant independent of B for any suboptimal arm. We conclude that $R_B = O(K \cdot \ln(B)/\Delta)$ with $\Delta = \min_{k \in \{1, \dots, K\} \mid \Delta_k > 0} \Delta_k$. Interestingly, the algorithm we propose does not rely on B to achieve this regret bound, much like what happens for UCB1 with the time horizon, see Auer et al. [4]. This result is optimal up to constant factors as the standard MAB problem is a special case of the framework considered in this section, see Lai and Robbins [19] for a proof of a lower bound in this context. It is possible to improve the constant factors when the consumption of the resource is deterministic as we can take $\lambda = 1$ in this scenario and the resulting regret bound is scale-free, see Section A of the Appendix. Building on Theorem 1, we can also derive a near-optimal distribution-free regret bound in the same fashion as for UCB1.

THEOREM 2. *We have:*

$$R_B \leq 8\lambda \cdot \sqrt{K \cdot \frac{B+1}{\epsilon} \cdot \ln\left(\frac{B+1}{\epsilon}\right)} + O(1).$$

Proof To get the distribution-free bound, we start from the penultimate inequality derived in the proof sketch of Theorem 1 and apply Lemma 5 only if Δ_k is big enough, noting that:

$$\sum_{k=1}^K \mathbb{E}[n_{k, \tau^*}] = \mathbb{E}[\tau^*] \leq (B+1)/\epsilon.$$

Specifically, we have:

$$\begin{aligned}
R_B &\leq \sup_{\substack{(n_1, \dots, n_K) \geq 0 \\ \sum_{k=1}^K n_k \leq \frac{B+1}{\epsilon}}} \left\{ \sum_{k \mid \Delta_k > 0} \min(\mu_k^c \cdot \Delta_k \cdot n_k, 2^6 \lambda^2 \cdot \frac{\ln(\frac{B+1}{\epsilon})}{\mu_k^c \cdot \Delta_k} + \frac{4\pi^2}{3\epsilon^2} \cdot \mu_k^c \cdot \Delta_k) \right\} + O(1) \\
&\leq \sup_{\substack{(n_1, \dots, n_K) \geq 0 \\ \sum_{k=1}^K n_k \leq \frac{B+1}{\epsilon}}} \left\{ \sum_{k \mid \Delta_k > 0} \min(\mu_k^c \cdot \Delta_k \cdot n_k, 2^6 \lambda^2 \cdot \frac{\ln(\frac{B+1}{\epsilon})}{\mu_k^c \cdot \Delta_k}) \right\} + K \cdot \frac{4\pi^2 \kappa}{3\epsilon^2} + O(1) \\
&\leq \sup_{\substack{(n_1, \dots, n_K) \geq 0 \\ \sum_{k=1}^K n_k \leq \frac{B+1}{\epsilon}}} \left\{ \sum_{k \mid \Delta_k > 0} \sqrt{2^6 \lambda^2 \cdot n_k \cdot \ln(\frac{B+1}{\epsilon})} \right\} + O(1) \\
&\leq 8\lambda \cdot \sqrt{K \cdot \frac{B+1}{\epsilon} \cdot \ln(\frac{B+1}{\epsilon})} + O(1),
\end{aligned}$$

where the second inequality is obtained with $\Delta_k \leq \mu_{k^*}^r / \mu_{k^*}^c \leq \kappa$ and $\mu_k^c \leq 1$, the third inequality is derived by maximizing on $(\mu_k^c \cdot \Delta_k) \geq 0$ for all arms k , and the last inequality is obtained with the Cauchy–Schwarz inequality.

We conclude that $R_B = O(\sqrt{K \cdot B \cdot \ln(B)})$, where the hidden constant factors are independent of the underlying distributions $(\nu_k)_{k=1, \dots, K}$.

Applications. Assumption 3 is natural for bidding in repeated second-price auctions when the auctioneer sets a reserve price R (this is common practice in sponsored search auctions). Indeed, then we have:

$$\begin{aligned}
\mathbb{E}[c_{k,t}] &= \mathbb{E}[m_t \cdot \mathbb{1}_{b_k \geq m_t}] \\
&\geq R \cdot \mathbb{E}[\mathbb{1}_{b_k \geq m_t}] \\
&\geq R \cdot \mathbb{E}[v_t \cdot \mathbb{1}_{b_k \geq m_t}] = R \cdot \mathbb{E}[r_{k,t}],
\end{aligned}$$

for any arm $k \in \{1, \dots, K\}$ and Assumption 3 is satisfied with $\kappa = 1/R$.

5. Arbitrarily many limited resources whose consumptions are deterministic. In this section, we study the case of multiple limited resources when the amounts of resources consumed as a result of pulling an arm are deterministic and globally constrained by prescribed budgets $(B(i))_{i=1, \dots, C}$, where C is the number of resources. Note that time need not be a constraint. Because the amounts of resources consumed are deterministic, we can substitute the notation $\mu_k^c(i)$ with $c_k(i)$ for any arm $k \in \{1, \dots, K\}$ and any resource $i \in \{1, \dots, C\}$. We point out that the stopping time need not be deterministic as the decision to select an arm at any round is based on the past realizations of the rewards. We define $\rho \leq \min(C, K)$ as the rank of the matrix $(c_k(i))_{1 \leq k \leq K, 1 \leq i \leq C}$.

Specification of the algorithm. We implement UCB-Simplex with an initialization step which consists in pulling each arm ρ times. The motivation behind this step is mainly technical and is simply meant to have:

$$n_{k,t} \geq \rho + \sum_{x \in \mathcal{B} \mid k \in \mathcal{K}_x} n_{k,t}^x \quad \forall t \in \mathbb{N}, \forall k \in \{1, \dots, K\}. \quad (10)$$

Compared to Section 4, we choose to take $\lambda = 1$ and $\eta_i = 0$ for any $i \in \{1, \dots, C\}$. As a result and since the amounts of resource consumption are deterministic, the exploration (resp. exploitation) terms defined in Section 3 specialize to $\text{obj}_{x,t} = \sum_{k=1}^K \zeta_k^x \cdot \bar{r}_{k,t}$ (resp. $E_{x,t} = \sum_{k=1}^K \zeta_k^x \cdot \epsilon_{k,t}$). Compared to the case of a single resource, we are required to specify the load balancing algorithms involved in Step-Load-Balance of UCB-Simplex as a feasible basis selected at Step-Simplex may involve several arms. Although Step-Simplex will also need to be specified in Sections 6 and 7, designing good load balancing algorithms is arguably easier here as the optimal load balance is known for each basis from the start. Nonetheless, one challenge remains:

we can never identify the (possibly many) optimal bases of (3) with absolute certainty. As a result, any basis selected at Step-Simplex should be treated as potentially optimal when balancing the load between the arms involved in this basis, but this inevitably causes some interference issues as an arm may be involved in several bases, and worst, possibly several optimal bases. Therefore, one point that will appear to be of particular importance in the analysis is the use of load balancing algorithms that are decoupled from one another, in the sense that they do not rely on what happened when selecting other bases. More specifically, we use the following class of load balancing algorithms.

Algorithm: Load balancing algorithm \mathcal{A}_x for a feasible basis $x \in \mathcal{B}$

If basis x is selected at time t , pull any arm $k \in \mathcal{K}_x$ such that $n_{k,t}^x \leq n_{x,t} \cdot \frac{\xi_k^x}{\sum_{l=1}^K \xi_l^x}$.

The load balancing algorithms \mathcal{A}_x thus defined are decoupled because, for each basis, the number of times an arm has been pulled when selecting any other basis is not taken into account. The following lemma shows that \mathcal{A}_x is always well defined and guarantees that the ratios $(n_{k,t}^x/n_{l,t}^x)_{k,l \in \mathcal{K}_x}$ remain close to the optimal ones $(\xi_k^x/\xi_l^x)_{k,l \in \mathcal{K}_x}$ at all times.

LEMMA 6. *For any basis $x \in \mathcal{B}$, \mathcal{A}_x is well defined and moreover, at any time t and for any arm $k \in \mathcal{K}_x$, we have:*

$$n_{x,t} \cdot \frac{\xi_k^x}{\sum_{l=1}^K \xi_l^x} - \rho \leq n_{k,t}^x \leq n_{x,t} \cdot \frac{\xi_k^x}{\sum_{l=1}^K \xi_l^x} + 1,$$

while $n_{k,t}^x = 0$ for any arm $k \notin \mathcal{K}_x$.

Proof. We need to show that there always exists an arm $k \in \mathcal{K}_x$ such that $n_{k,t}^x \leq n_{x,t} \cdot \xi_k^x / \sum_{l=1}^K \xi_l^x$. Suppose there is none, we have:

$$n_{x,t} = \sum_{k \in \mathcal{K}_x} n_{k,t}^x > \sum_{k \in \mathcal{K}_x} n_{x,t} \cdot \frac{\xi_k^x}{\sum_{l=1}^K \xi_l^x} = n_{x,t},$$

a contradiction. Moreover, we have, at any time t and for any arm $k \in \mathcal{K}_x$, $n_{k,t}^x \leq n_{x,t} \cdot \xi_k^x / \sum_{l=1}^K \xi_l^x + 1$. Indeed, suppose otherwise and define $t^* \leq t$ as the last time arm k was pulled, we have:

$$n_{k,t^*}^x = n_{k,t}^x - 1 > n_{x,t} \cdot \frac{\xi_k^x}{\sum_{l=1}^K \xi_l^x} \geq n_{x,t^*} \cdot \frac{\xi_k^x}{\sum_{l=1}^K \xi_l^x},$$

which shows, by definition, that arm k could not have been pulled at time t^* . We also derive as a byproduct that, at any time t and for any arm $k \in \mathcal{K}_x$, $n_{x,t} \cdot \xi_k^x / \sum_{l=1}^K \xi_l^x - \rho \leq n_{k,t}^x$ since $n_{x,t} = \sum_{k \in \mathcal{K}_x} n_{k,t}^x$ and since a basis involves at most ρ arms.

Observe that the load balancing algorithms \mathcal{A}_x run in time $O(1)$ but may require a memory storage capacity exponential in C and polynomial in K , although always bounded by $O(B)$ (because we do not need to keep track of $n_{k,t}^x$ if x has never been selected). In practice, only a few bases will be selected at Step-Simplex, so that a hash table is an appropriate data structure to store the sequences $(n_{k,t}^x)_{k \in \mathcal{K}_x}$. In Section A of the Appendix, we introduce another class of load balancing algorithms that is both time and memory efficient while still guaranteeing logarithmic regret bounds (under an additional assumption) but no distribution-free regret bounds.

Regret Analysis. We use the shorthand notation:

$$\epsilon = \min_{\substack{k=1, \dots, K \\ i=1, \dots, C \\ \text{with } c_k(i) > 0}} c_k(i).$$

Note that $\epsilon < \infty$ under Assumption 2. We discard the initialization step in the theoretical study because the amounts of resources consumed are bounded by a constant and the total reward obtained is non-negative and not taken into account in the decision maker's total payoff. We again start by estimating the expected time horizon.

LEMMA 7. *For any non-anticipating algorithm, we have: $\mathbb{E}[\tau^*] \leq \frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1$.*

Proof. By definition of τ^* , we have $\sum_{t=1}^{\tau^*-1} c_{a_t, t}(i) \leq B(i)$ for any resource $i \in \{1, \dots, C\}$. Summing up these inequalities and using Assumption 2 and the fact that $(c_{k,t}(i))_{t=1, \dots, T}$ are deterministic, we get $(\tau^* - 1) \cdot \epsilon \leq \sum_{i=1}^C B(i)$. Taking expectations on both sides and using Assumption 1 yields the result.

We follow by bounding the number of times any suboptimal basis can be selected at Step-Simplex in the same spirit as in Section 4.

LEMMA 8. *For any suboptimal basis $x \in \mathcal{B}$, we have:*

$$\mathbb{E}[n_{x, \tau^*}] \leq 16\rho \cdot \left(\sum_{k=1}^K \xi_k^x \right)^2 \cdot \frac{\mathbb{E}[\ln(\tau^*)]}{(\Delta_x)^2} + \rho \cdot \frac{\pi^2}{3}.$$

Sketch of proof. We use the shorthand notation $\beta_x = 8\rho \cdot (\sum_{k=1}^K \xi_k^x / \Delta_x)^2$. The proof is along the same lines as for Lemma 5. First note that we may assume, without loss of generality, that x has been selected at least $\beta_x \cdot \ln(t)$ times at any time t up to an additive term of $2\beta_x \cdot \mathbb{E}[\ln(\tau^*)]$ in the final inequality. We then just have to bound by a constant the probability that x is selected at any time t given that $n_{x,t} \geq \beta_x \cdot \ln(t)$. If x is selected at time t , x is an optimal basis to (8). Since the amounts of resources consumed are deterministic, x^* is feasible to (8) at time t , which implies that $\text{obj}_{x,t} + E_{x,t} \geq \text{obj}_{x^*,t} + E_{x^*,t}$. This can only happen if either: (i) $\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}$, (ii) $\text{obj}_{x^*,t} \leq \text{obj}_{x^*} - E_{x^*,t}$, or (iii) $\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}$. First note that (iii) is impossible because, assuming this is the case, we would have:

$$\begin{aligned} \frac{\Delta_x}{2} &< \sum_{k \in \mathcal{K}_x} \xi_k^x \cdot \sqrt{\frac{2 \ln(t)}{n_{k,t}}} \\ &\leq \sum_{k \in \mathcal{K}_x} \xi_k^x \cdot \sqrt{\frac{2 \ln(t)}{\rho + n_{k,t}^x}} \\ &\leq \sqrt{\sum_{k \in \mathcal{K}_x} \xi_k^x} \cdot \sum_{k \in \mathcal{K}_x} \sqrt{\xi_k^x} \cdot \sqrt{\frac{2 \ln(t)}{n_{x,t}}} \\ &\leq \sqrt{\rho} \cdot \sum_{k \in \mathcal{K}_x} \xi_k^x \cdot \sqrt{\frac{2}{\beta_x}} = \frac{\Delta_x}{2}, \end{aligned}$$

where we use (10), Lemma 6 for each $k \in \mathcal{K}_x$, the Cauchy–Schwarz inequality, and the fact that a basis involves at most ρ arms. Along the same lines as for Lemma 5, the probability of events (i) and (ii) can be bounded by $\sim \rho/t^2$ in the same fashion, irrespective of how many times x and x^* have been selected in the past. For example for event (i), this is because if $\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}$, then there must exist $k \in \mathcal{K}_x$ such that $\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}$, but any of these events have individual probability at most $\sim 1/t^2$ by Lemma 1. Indeed otherwise, if $\bar{r}_{k,t} < \mu_k^r + \epsilon_{k,t}$ for all $k \in \mathcal{K}_x$, we would have:

$$\begin{aligned} \text{obj}_{x,t} - \text{obj}_x &= \sum_{k \in \mathcal{K}_x} (\bar{r}_{k,t} - \mu_k^r) \cdot \xi_k^x \\ &< \sum_{k \in \mathcal{K}_x} \epsilon_{k,t} \cdot \xi_k^x = E_{x,t}, \end{aligned}$$

a contradiction

Lemma 8 used in combination with Lemma 7 shows that a suboptimal basis is selected at most $O(\ln(B))$ times. To establish the regret bound, what remains to be done is to lower bound the expected total payoff derived when selecting any of the optimal bases. This is more involved than in the case of a single limited resource because the load balancing step comes into play at this stage.

THEOREM 3. *We have:*

$$R_{B(1), \dots, B(C)} \leq 16 \frac{\rho \cdot \sum_{i=1}^C b(i)}{\epsilon} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln \left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1 \right) + O(1).$$

Sketch of proof. The proof proceeds along the same lines as for Theorem 1. We build upon (4):

$$\begin{aligned} R_{B(1), \dots, B(C)} &\leq B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E} \left[\sum_{t=1}^{\tau^*} r_{a_t, t} \right] + O(1) \\ &= B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \sum_{k=1}^K \mu_k^r \cdot \mathbb{E} [n_{k, \tau^*}^x] + O(1). \end{aligned}$$

Using the properties of the load balancing algorithm established in Lemma 6, we derive:

$$\begin{aligned} R_{B(1), \dots, B(C)} &\leq B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \left\{ \frac{\mathbb{E} [n_{x, \tau^*}]}{\sum_{k=1}^K \xi_k^x} \cdot \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x \right) \right\} + O(1) \\ &= \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \right) \cdot \left(B - \sum_{x \in \mathcal{B} \mid \Delta_x = 0} \frac{\mathbb{E} [n_{x, \tau^*}]}{\sum_{k=1}^K \xi_k^x} \right) \\ &\quad - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \left\{ \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x \right) \cdot \frac{\mathbb{E} [n_{x, \tau^*}]}{\sum_{k=1}^K \xi_k^x} \right\} + O(1). \end{aligned}$$

Now observe that, by definition, at least one resource is exhausted at time τ^* . Hence, there exists $i \in \{1, \dots, C\}$ such that:

$$\begin{aligned} B(i) &\leq \sum_{x \in \mathcal{B}} \sum_{k \in \mathcal{K}_x} c_k(i) \cdot n_{k, \tau^*}^x \\ &\leq O(1) + \sum_{x \in \mathcal{B}} \frac{n_{x, \tau^*}}{\sum_{k=1}^K \xi_k^x} \cdot \sum_{k \in \mathcal{K}_x} c_k(i) \cdot \xi_k^x \\ &\leq O(1) + b(i) \cdot \sum_{x \in \mathcal{B}} \frac{n_{x, \tau^*}}{\sum_{k=1}^K \xi_k^x}, \end{aligned}$$

where we use Lemma 6 again and the fact that any basis $x \in \mathcal{B}$ satisfies all the constraints of (3). We conclude that:

$$\sum_{x \in \mathcal{B} \mid \Delta_x = 0} \frac{n_{x, \tau^*}}{\sum_{k=1}^K \xi_k^x} \geq B - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{n_{x, \tau^*}}{\sum_{k=1}^K \xi_k^x} + O(1).$$

Taking expectations on both sides and plugging the result back into the regret bound yields:

$$R_{B(1), \dots, B(C)} \leq \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{\Delta_x}{\sum_{k=1}^K \xi_k^x} \cdot \mathbb{E} [n_{x, \tau^*}] + O(1).$$

Using Lemma 7, Lemma 8, and the concavity of the logarithmic function, we obtain:

$$\begin{aligned} R_{B(1), \dots, B(C)} &\leq 16 \rho \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{\sum_{k=1}^K \xi_k^x}{\Delta_x} \right) \cdot \ln \left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1 \right) \\ &\quad + \frac{\pi^2}{3} \rho \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{\Delta_x}{\sum_{k=1}^K \xi_k^x} \right) + O(1) \end{aligned}$$

which yields the claim since $\Delta_x \leq \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \leq \sum_{i=1}^C b(i) / \epsilon$, $\sum_{k=1}^K \xi_k^x \geq b$, and $\sum_{k=1}^K \xi_k^x \leq \sum_{i=1}^C b(i) / \epsilon$.

We point out that, if time is a limited resource with a time horizon T , we can also derive the (possibly better) regret bound:

$$R_{B(1), \dots, B(C)} \leq 16\rho \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln(T) + O(1).$$

Since the number of feasible bases to (3) is at most $\binom{K+\rho}{K} \leq 2K^\rho$, we get the distribution-dependent regret bound $O(\rho \cdot K^\rho \cdot \ln(B)/\Delta)$ where $\Delta = \min_{x \in \mathcal{B} \mid \Delta_x > 0} \Delta_x$. In Section A of the Appendix, we introduce an alternative class of load balancing algorithms which yields a better dependence on K and C with a regret bound of order $O(\rho^3 \cdot K \cdot \ln(B)/\Delta^2)$ provided that there is a unique optimal basis to (3). Along the same lines as in Section 4, the distribution-dependent bound of Theorem 3 almost immediately implies a distribution-free one.

THEOREM 4. *We have:*

$$R_{B(1), \dots, B(C)} \leq 4\sqrt{\rho \cdot |\mathcal{B}| \cdot \left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1 \right) \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1 \right)} + O(1).$$

Sketch of proof. The proof is along the same lines as for the case of a single limited resource, we start from the penultimate inequality derived in the proof sketch of Theorem 3 and apply Lemma 8 only if Δ_x is big enough, taking into account the fact that $\sum_{x \in \mathcal{B}} \mathbb{E}[n_{x, \tau^*}] \leq \mathbb{E}[\tau^*] \leq \sum_{i=1}^C b(i) \cdot B/\epsilon + 1$.

We conclude that $R_{B(1), \dots, B(C)} = O(\sqrt{\rho \cdot K^\rho \cdot B \cdot \ln(B)})$, where the hidden constant factors are independent of the underlying distributions $(\nu_k)_{k=1, \dots, K}$. If time is a limited resource, we can also derive the (possibly better) regret bound:

$$R_{B(1), \dots, B(C)} \leq 4\sqrt{\rho \cdot |\mathcal{B}| \cdot T \cdot \ln(T)} + O(1).$$

In any case, we stress that the dependence on K and C is not optimal since Badanidiyuru et al. [8] and Agrawal and Devanur [1] obtain a $\tilde{O}(\sqrt{K \cdot B})$ bound on regret, where the \tilde{O} notation hides factors logarithmic in B . Observe that the regret bounds derived in this section do not vanish with b . This can be remedied by strengthening Assumption 2, additionally assuming that $c_{k,t}(i) > 0$ for any arm $k \in \{1, \dots, K\}$ and resource $i \in \{1, \dots, C\}$. In this situation, we can refine the analysis and substitute $\sum_{i=1}^C b(i)$ with b in the regret bounds of Propositions 3 and 4 which become scale-free.

Applications. BwK problems where the amounts of resources consumed as a result of pulling an arm are deterministic find applications in network revenue management of perishable goods, shelf optimization of perishable goods, and wireless sensor networks, as detailed in Section 2.

6. A time horizon and another limited resource. In this section, we investigate the case of two limited resources, one of which is assumed to be time, with a time horizon T , while the consumption of the other is stochastic and constrained by a global budget B . To simplify the notations, we omit the indices identifying the resources since the second limited resource is time and we write μ_k^c , $c_{k,t}$, $\bar{c}_{k,t}$, B , and T as opposed to $\mu_k^c(1)$, $c_{k,t}(1)$, $\bar{c}_{k,t}(1)$, $B(1)$, and $B(2)$. Moreover, we refer to resource $i = 1$ as “the” resource. Observe that, in the particular setting considered in this section, $\tau^* = \min(\tau(B), T + 1)$ with $\tau(B) = \min\{t \in \mathbb{N} \mid \sum_{\tau=1}^t c_{a_\tau, \tau} > B\}$. Note that the budget constraint is not limiting if $B \geq T$, in which case the problem reduces to the standard MAB problem. Hence, without loss of generality under Assumption 1, we assume that the budget scales linearly with time, i.e. $B = b \cdot T$ for a fixed constant $b \in (0, 1)$, and we study the asymptotic regime $T \rightarrow \infty$. Motivated by technical considerations, we make two additional assumptions for the particular setting considered in this section that are perfectly reasonable in many applications, such as in repeated second-price auctions, dynamic procurement, and dynamic pricing, as detailed in the last paragraph of this section.

ASSUMPTION 4. *There exists $\sigma > 0$ such that $\mu_k^r \leq \sigma \cdot \mu_k^c$ for any arm $k \in \{1, \dots, K\}$.*

ASSUMPTION 5. *The decision maker knows $\kappa > 0$ such that:*

$$|\mu_k^r - \mu_l^r| \leq \kappa \cdot |\mu_k^c - \mu_l^c| \quad \forall (k, l) \in \{1, \dots, K\}^2,$$

ahead of round 1.

Note that σ , as opposed to κ , is not assumed to be known to the decision maker. Assumption 4 is relatively weak and is always satisfied in practical applications. In particular, note that if $\mu_k^c > 0$ for all arms $k \in \{1, \dots, K\}$, we can take $\sigma = 1 / \min_{k=1, \dots, K} \mu_k^c$.

Specification of the algorithm. We implement UCB-Simplex with $\lambda = 1 + 2\kappa$, $\eta_1 = 1$, $\eta_2 = 0$, and an initialization step which consists in pulling each arm once. Because the amount of resource consumed at each round is a random variable, a feasible basis for (8) may not be feasible for (3) and conversely. In particular, x^* may not be feasible for (8), thus effectively preventing it from being selected at Step-Simplex, and an infeasible basis for (3) may be selected instead. This is in contrast to the situation studied in Section 5 and this motivates the choice $\eta_1 > 0$ to guarantee that any feasible solution to (3) will be feasible to (8) with high probability at any round t .

Just like in Section 5, we need to specify Step-Load-Balance because a basis selected at Step-Simplex may involve up to two arms. To simplify the presentation, we introduce a dummy arm $k = 0$ with reward 0 and resource consumption 0 (pulling this arm corresponds to skipping the round) and K dummy arms $k = K + 1, \dots, 2K$ with reward identical to arm $K - k$ but resource consumption 1 so that any basis involving a single arm can be mapped to an ‘‘artificial’’ one involving two arms. Note, however, that we do not introduce new variables ξ_k in (8) for these arms as they are only used for mathematical convenience in Step-Load-Balance once a basis has been selected at Step-Simplex. Specifically, if a basis x_t involving a single arm determined by $\mathcal{K}_{x_t} = \{k_t\}$ and $\mathcal{C}_{x_t} = \{1\}$ (resp. $\mathcal{C}_{x_t} = \{2\}$) is selected at Step-Simplex, we map it to the basis x'_t determined by $\mathcal{K}_{x'_t} = \{0, k_t\}$ (resp. $\{k_t, K + k_t\}$) and $\mathcal{C}_{x'_t} = \{1, 2\}$. We then use a load balancing algorithm specific to this basis, denoted by \mathcal{A}_{x_t} , to determine which of the two arms in $\mathcal{K}_{x'_t}$ to pull. Similarly as in Section 5, using load balancing algorithms that are decoupled from one another is crucial because the decision maker can never identify the optimal bases with absolute certainty. This implies that each basis should be treated as potentially optimal when balancing the load between the arms, but this inevitably causes interference issues as an arm may be involved in several bases. Compared to Section 5, we face an additional challenge when designing the load balancing algorithms: the optimal load balances are initially unknown to the decision maker. It turns out that we can still approximately achieve the unknown optimal load balances by enforcing that, at any round t , the total amount of resource consumed remains close to the pacing target $b \cdot t$ with high probability, as precisely described below.

Algorithm: Load balancing algorithm \mathcal{A}_x for any basis x

For any time period t , define $b_{x,t}$ as the total amount of resource consumed when selecting x in the past $t - 1$ rounds. Suppose that x is selected at time t . Without loss of generality, write $\mathcal{K}_x = \{k, l\}$ with $\bar{c}_{k,t} - \epsilon_{k,t} \geq \bar{c}_{l,t} - \epsilon_{l,t}$. Pull arm k if $b_{x,t} \leq n_{x,t} \cdot b$ and pull arm l otherwise.

Observe that a basis x with $\mathcal{K}_x = \{k, l\}$ is feasible for (3) if either $\mu_k^c > b > \mu_l^c$ or $\mu_l^c > b > \mu_k^c$. Assuming we are in the first situation, the exploration and exploitation terms defined in Section 3 specialize to:

$$\text{obj}_{x,t} = \xi_{l,t}^x \cdot \bar{r}_{l,t} + \xi_{k,t}^x \cdot \bar{r}_{k,t} \text{ and } E_{x,t} = \lambda \cdot (\xi_{l,t}^x \cdot \epsilon_{l,t} + \xi_{k,t}^x \cdot \epsilon_{k,t})$$

with:

$$\xi_{l,t}^x = \frac{(\bar{c}_{k,t} - \epsilon_{k,t}) - b}{(\bar{c}_{k,t} - \epsilon_{k,t}) - (\bar{c}_{l,t} - \epsilon_{l,t})} \text{ and } \xi_{k,t}^x = \frac{b - (\bar{c}_{l,t} - \epsilon_{l,t})}{(\bar{c}_{k,t} - \epsilon_{k,t}) - (\bar{c}_{l,t} - \epsilon_{l,t})},$$

provided that $\bar{c}_{k,t} - \epsilon_{k,t} > b > \bar{c}_{l,t} - \epsilon_{l,t}$. Moreover, their offline counterparts are given by:

$$\text{obj}_x = \xi_l^x \cdot \mu_l^r + \xi_k^x \cdot \mu_k^r, \quad \xi_l^x = \frac{\mu_k^c - b}{\mu_k^c - \mu_l^c}, \text{ and } \xi_k^x = \frac{b - \mu_l^c}{\mu_k^c - \mu_l^c}.$$

Regret Analysis. We start by pointing out that, in degenerate scenarios, using the linear relaxation (3) as an upper bound on $\text{ER}_{\text{OPT}}(B, T)$ already dooms us to $\Omega(\sqrt{T})$ regret bounds. Precisely, if there exists a unique optimal basis x^* to (3) that happens to be degenerate, i.e. $\mathcal{K}_{x^*} = \{k^*\}$ (pre-mapping) with $\mu_{k^*}^c = b$, then, in most cases, $T \cdot \text{obj}_{x^*} \geq \text{ER}_{\text{OPT}}(B, T) + \Omega(\sqrt{T})$ as shown below.

LEMMA 9. *If there exists $k^* \in \{1, \dots, K\}$ such that: (i) the i.i.d. process $(c_{k^*, t})_{t \in \mathbb{N}}$ has positive variance, (ii) $\mu_{k^*}^c = b$, and (iii) $(\xi_k)_{k=1, \dots, K}$ determined by $\xi_{k^*} = 1$ and $\xi_k = 0$ for $k \neq k^*$ is the unique optimal solution to (3), then there exists a subsequence of $(\frac{T \cdot \text{obj}_{x^*} - \text{ER}_{\text{OPT}}(B, T)}{\sqrt{T}})_{T \in \mathbb{N}}$ that does not converge to 0.*

Sketch of proof. For any time horizon $T \in \mathbb{N}$ and any arm $k \in \{1, \dots, K\}$, we denote by $n_{k, T}^{\text{opt}}$ the expected number of times arm k is pulled by the optimal non-anticipating algorithm when the time horizon is T and the budget is $B = b \cdot T$. We expect that consistently pulling arm k^* is near-optimal. Unfortunately, this is also nothing more than an i.i.d. strategy which implies, along the same lines as in Lemma 3, that $\mathbb{E}[\tau^*] = T - \Omega(\sqrt{T})$ so that the total expected payoff is $\mathbb{E}[\tau^*] \cdot \mu_{k^*}^c = T \cdot \text{obj}_{x^*} - \Omega(\sqrt{T})$. To formalize these ideas, we study two cases: $T - n_{k^*, T}^{\text{opt}} = \Omega(\sqrt{T})$ (Case A) and $T - n_{k^*, T}^{\text{opt}} = o(\sqrt{T})$ (Case B) and we show that $\text{ER}_{\text{OPT}}(B, T) = T \cdot \text{obj}_{x^*} - \Omega(\sqrt{T})$ in both cases. In Case A, this is because the optimal value of (3) remains an upper bound on the maximum total expected payoff if we add the constraint $\xi_{k^*} \leq n_{k^*, T}^{\text{opt}}/T$ to the linear program (3) by definition of $n_{k^*, T}^{\text{opt}}$. Since the constraint $\xi_{k^*} \leq 1$ is binding for (3), the optimal value of this new linear program can be shown to be smaller than $\text{obj}_{x^*} - \Omega((T - n_{k^*, T}^{\text{opt}})/T)$ (by strong duality and strict complementary slackness). In Case B, up to an additive term of order $o(\sqrt{T})$ in the final bound, the optimal non-anticipating algorithm is equivalent to consistently pulling arm k^* , which is an i.i.d. strategy so the study is very similar to that of Lemma 3.

Dealing with these degenerate scenarios thus calls for a completely different approach than the one taken on in the BwK literature and we choose instead to rule them out in such a way that there can be no degenerate optimal basis to (3).

ASSUMPTION 6. *We have $|\mu_k^c - b| > 0$ for any arm $k \in \{1, \dots, K\}$.*

We use the shorthand notation $\epsilon = \min_{k=1, \dots, K} |\mu_k^c - b|$. Assumption 6 is equivalent to assuming that any basis for (3) is non-degenerate. This assumption can be relaxed to some extent at the price of more technicalities. However, in light of Lemma 9, the minimal assumption is that there is no degenerate optimal basis to (3). As a final remark, we stress that Assumption 6 is only necessary to carry out the analysis but Step-Simplex can be implemented in any case as ϵ is not assumed to be known to the decision maker.

We are now ready to establish regret bounds. Without loss of generality, we can assume that any pseudo-basis for (3) involves two arms, one of which may be a dummy arm introduced in the specification of the algorithm detailed above. As stressed at the beginning of this section, UCB-Simplex may sometimes select an infeasible basis or even a pseudo-basis x with $\det(A_x) = 0$ (i.e. such that $\mu_k^c = \mu_l^c$ assuming $\mathcal{K}_x = \{k, l\}$). Interestingly the load balancing algorithm plays a crucial role to guarantee that this does not happen very often.

LEMMA 10. *For any basis $x \notin \mathcal{B}$, we have:*

$$\mathbb{E}[n_{x, T}] \leq \frac{2^6}{\epsilon^3} \cdot \ln(T) + \frac{10\pi^2}{3\epsilon^2}.$$

The same inequality holds if x is a pseudo-basis but not a basis for (3).

Proof. We use the shorthand notation $\beta_x = 2^5/\epsilon^3$. Without loss of generality, we can assume that $\mathcal{K}_x = \{k, l\}$ with $\mu_k^c, \mu_l^c > b$ (the situation is symmetric if the reverse inequality holds). Along the same lines as in Lemma 5, we only have to bound by a constant the probability that x is selected at any round t given that x has already been selected at least $\beta_x \cdot \ln(t)$ times. If x is selected at round t and $n_{x, t} \geq \beta_x \cdot \ln(t)$, then $b_{x, t}$ must be larger than $n_{x, t} \cdot b$ by at least a margin of $\sim 1/\epsilon^2 \cdot \ln(t)$ with high probability given that $\mu_k^c, \mu_l^c > b$. Moreover, at least one arm, say k , has been pulled at least $\sim 1/\epsilon^3 \cdot \ln(t)$ times and, as a result,

$\bar{c}_{k,\tau} - \epsilon_{k,\tau} \geq b$ with high probability for the last $s \sim 1/\epsilon^2 \cdot \ln(t)$ rounds $\tau = \tau_1, \dots, \tau_s$ where x was selected. This implies that arm l must have been pulled at least $\sim 1/\epsilon^2 \cdot \ln(t)$ times already by definition of the load balancing algorithm but then we have $\bar{c}_{l,t} - \epsilon_{l,t} \geq b$ with high probability and x cannot be feasible for (8) at time t with high probability.

What remains to be done is to: (i) show that suboptimal bases are selected at most $O(\ln(T))$ times and (ii) lower bound the expected total payoff derived when selecting any of the optimal bases. The major difficulty lies in the fact that the amounts of resource consumed, the rewards obtained, and the stopping time are correlated in a non-trivial way through the budget constraint and the decisions made in the past. This makes it difficult to study the expected total payoff derived when selecting optimal bases independently from the amounts of resource consumed and the rewards obtained when selecting suboptimal ones. However, a key point is that, by design, the pulling decision made at Step-Load-Balance is based solely on the past history associated with the basis selected at Step-Simplex because the load balancing algorithms are decoupled. For this reason, the analysis proceeds in two steps irrespective of the number of optimal bases. In a first step, we show that, for any basis x for (3), the amount of resource consumed per round when selecting x remains close to the pacing target b with high probability. This enables us to show that the ratios $(\mathbb{E}[n_{k,T}^x]/\mathbb{E}[n_{l,T}^x])_{k,l \in \mathcal{K}_x}$ are close to the optimal ones $(\xi_k^x/\xi_l^x)_{k,l \in \mathcal{K}_x}$, as precisely stated below.

LEMMA 11. *For any basis $x \in \mathcal{B}$ and time period t , we have:*

$$\mathbb{P}[|b_{x,t} - n_{x,t} \cdot b| \geq u + (\frac{4}{\epsilon})^2 \cdot \ln(t)] \leq \frac{4}{\epsilon^2} \cdot \exp(-\epsilon^2 \cdot u) + \frac{8}{\epsilon^2 \cdot t^2} \quad \forall u \geq 1,$$

which, in particular, implies that:

$$\begin{aligned} \mathbb{E}[n_{k,T}^x] &\geq \xi_k^x \cdot \mathbb{E}[n_{x,T}] - 13/\epsilon^5 - 16/\epsilon^3 \cdot \ln(T) \\ \mathbb{E}[n_{l,T}^x] &\geq \xi_l^x \cdot \mathbb{E}[n_{x,T}] - 13/\epsilon^5 - 16/\epsilon^3 \cdot \ln(T). \end{aligned} \tag{11}$$

Sketch of proof. Without loss of generality, we can assume that $\mathcal{K}_x = \{k, l\}$ with $\mu_k^c > b > \mu_l^c$. Observe that, if the decision maker knew that $\mu_k^c > b > \mu_l^c$ ahead of round 1, he would always pull the ‘‘correct’’ arm in order not to deviate from the pacing target $n_{x,t} \cdot b$ and $|b_{x,t} - n_{x,t} \cdot b|$ would remain small with high probability given Assumption 6. However, because this information is not available ahead of round 1, the decision maker is led to pull the incorrect arm when arm k and l are swapped, in the sense that $\bar{c}_{k,t} - \epsilon_{k,t} \leq \bar{c}_{l,t} - \epsilon_{l,t}$. Fortunately, at any time t , there could have been at most $1/\epsilon^2 \cdot \ln(t)$ swaps with probability at least $\sim 1 - 1/t^2$ given Assumption 6. To derive (11), we use: $|\mathbb{E}[b_{x,T} - n_{x,T} \cdot b]| \leq \int_0^T \mathbb{P}[|b_{x,t} - n_{x,t} \cdot b| \geq u] du$ and $\mathbb{E}[b_{x,T}] = \mu_k^c \cdot \mathbb{E}[n_{k,T}^x] + \mu_l^c \cdot \mathbb{E}[n_{l,T}^x]$.

The next step is to show, just like in Section 5, that any suboptimal feasible basis is selected at most $O(\ln(T))$ times on average. Interestingly, the choice of the load balancing algorithm plays a minor role in the proof. Any load balancing algorithm that pulls each arm involved in a basis at least a constant fraction of the time this basis is selected does enforce this property.

LEMMA 12. *For any suboptimal basis $x \in \mathcal{B}$, we have:*

$$\mathbb{E}[n_{x,T}] \leq 2^9 \frac{\lambda^2}{\epsilon^3} \cdot \frac{\ln(T)}{(\Delta_x)^2} + \frac{10\pi^2}{\epsilon^2}.$$

Sketch of proof. We use the shorthand notation $\beta_x = 2^8/\epsilon^3 \cdot (\lambda/\Delta_x)^2$. Without loss of generality, we can assume that $\mathcal{K}_{x^*} = \{k^*, l^*\}$ with $\mu_{k^*}^c > b > \mu_{l^*}^c$ and $\mathcal{K}_x = \{k, l\}$ with $\mu_k^c > b > \mu_l^c$. Along the same lines as in Lemma 5, we only have to bound by a constant the probability that x is selected at any round t given that x has already been selected at least $\beta_x \cdot \ln(t)$ times. If x is selected at time t , x is optimal for (8). Observe that $(\xi_k^{x^*})_{k=1, \dots, K}$ is a feasible solution to (8) when $\bar{c}_{k^*,t} - \epsilon_{k^*,t} \leq \mu_{k^*}^c$ and $\bar{c}_{l^*,t} - \epsilon_{l^*,t} \leq \mu_{l^*}^c$, which happens with probability at least $\sim 1 - 1/t^2$. As a result, $\text{obj}_{x,t} + E_{x,t} \geq \text{obj}_{x^*}$ when additionally $\bar{r}_{k^*,t} + \epsilon_{k^*,t} \geq \mu_{k^*}^r$

and $\bar{r}_{l^*,t} + \epsilon_{l^*,t} \geq \mu_{l^*}^r$, which also happens with probability at least $\sim 1 - 1/t^2$. If $\text{obj}_{x,t} + E_{x,t} \geq \text{obj}_{x^*}$ then we have either (i) $\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}$ or (ii) $\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}$. Observe that (ii) can only happen with probability at most $\sim 1/t^2$ given that $n_{x,t} \geq \beta_x \cdot \ln(t)$ because (ii) implies that either $n_{l,t} \leq 8(\lambda/\Delta_x)^2 \cdot \ln(t)$ or $n_{k,t} \leq 8(\lambda/\Delta_x)^2 \cdot \ln(t)$ but the load balancing algorithm guarantees that each arm is pulled a fraction of the time x is selected (using Lemma 11). As for (i), if $\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}$, then, using Assumption 4, either $\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}$, $\bar{c}_{k,t} \notin [\mu_k^c - \epsilon_{k,t}, \mu_k^c + \epsilon_{k,t}]$, $\bar{r}_{l,t} \geq \mu_l^r + \epsilon_{l,t}$, or $\bar{c}_{l,t} \notin [\mu_l^c - \epsilon_{l,t}, \mu_l^c + \epsilon_{l,t}]$ but all of these events have individual probability at most $\sim 1/t^2$ by Lemma 1.

In a last step, we show, using Lemma 11, that, at the cost of an additive logarithmic term in the regret bound, we may assume that the game lasts exactly T rounds. This enables us to combine Lemmas 10, 11, and 12 to establish a distribution-dependent regret bound.

THEOREM 5. *We have:*

$$R_{B,T} \leq 2^9 \frac{\lambda^2}{\epsilon^3} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln(T) + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \cdot \ln(T)\right),$$

where the O notation hides universal constant factors.

Sketch of proof. We build upon (4):

$$\begin{aligned} R_{B,T} &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t,t}\right] + O(1) \\ &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^T r_{a_t,t}\right] + \sigma \cdot \mathbb{E}\left[\left(\sum_{t=1}^T c_{a_t,t} - B\right)_+\right] + O(1), \end{aligned}$$

where we use Assumption 4 for the second inequality. Moreover:

$$\begin{aligned} \mathbb{E}\left[\left(\sum_{t=1}^T c_{a_t,t} - B\right)_+\right] &\leq \sum_{x \in \mathcal{B}} \mathbb{E}[|b_{x,T} - n_{x,T} \cdot b|] + \sum_{x \notin \mathcal{B}} \mathbb{E}[n_{x,T}] \\ &\quad + \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x)=0}} \mathbb{E}[n_{x,T}] = O\left(\frac{K^2}{\epsilon^3} \ln(T)\right), \end{aligned}$$

using Lemmas 10 and 11. Plugging this last inequality back into the regret bound yields:

$$\begin{aligned} R_{B,T} &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^T r_{a_t,t}\right] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[n_{k,T}^x] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x,T}] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &= \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \cdot \left(T - \sum_{x \in \mathcal{B} \mid \Delta_x = 0} \mathbb{E}[n_{x,T}]\right) - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x,T}] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &= \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \mathbb{E}[n_{x,T}] + \sum_{x \notin \mathcal{B}} \mathbb{E}[n_{x,T}] + \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x)=0}} \mathbb{E}[n_{x,T}]\right) \\ &\quad - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x,T}] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \end{aligned}$$

$$\begin{aligned} &\leq \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \Delta_x \cdot \mathbb{E}[n_{x,T}] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &\leq 2^9 \frac{\lambda^2}{\epsilon^3} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln(T) + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right), \end{aligned}$$

where we use Lemma 11 for the third inequality, Lemma 10 along with $\sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \leq \sum_{k=1}^K \xi_k^{x^*} \leq 1$ for the fourth inequality, and Lemma 12 for the last inequality.

Since there are at most $2K^2$ feasible bases, we get the regret bound $O(K^2 \cdot (1/\Delta + \sigma/\epsilon^3) \cdot \ln(T))$, where $\Delta = \min_{x \in \mathcal{B} \mid \Delta_x > 0} \Delta_x$. Along the same lines as in Sections 4 and 5, pushing the analysis further almost immediately yields a distribution-free regret bound.

THEOREM 6. *We have:*

$$R_{B,T} \leq 2^5 \frac{\lambda}{\epsilon^{3/2}} \cdot \sqrt{|\mathcal{B}| \cdot T \cdot \ln(T)} + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right),$$

where the O notation hides universal constant factors.

Sketch of proof. The proof is along the same lines as for Theorems 2 and 4, we start from the penultimate inequality derived in the proof sketch of Theorem 5 and apply Lemma 12 only if Δ_x is big enough, taking into account the fact that $\sum_{x \in \mathcal{B}} \mathbb{E}[n_{x,T}] \leq T$.

We conclude that $R_{B,T} = O(\sqrt{K^2 \cdot T \cdot \ln(T)})$, where the hidden factors are independent of the underlying distributions $(\nu_k)_{k=1, \dots, K}$. Just like in Section 5, we stress that the dependence on K is not optimal since Badanidiyuru et al. [8] and Agrawal and Devanur [1] obtain a $\tilde{O}(\sqrt{K \cdot T})$ bound on regret, where the \tilde{O} notation hides factors logarithmic in T . Observe that the regret bounds derived in Theorems 5 and 6 do not vanish with b , which is not the expected behavior. This is a shortcoming of the analysis that can easily be remedied when $\min_{k=1, \dots, K} \mu_k^c > 0$ provided that instead of pulling the dummy arm 0 we always pull the other arm involved in the basis (i.e. we never skip rounds). Note that not skipping rounds can only improve the regret bounds derived in Theorems 5 and 6: arm 0 was introduced only in order to harmonize the notations for mathematical convenience.

THEOREM 7. *Relax Assumption 6 and redefine $\epsilon = \min_{k=1, \dots, K} \mu_k^c$. Suppose that $b \leq \epsilon/2$ and that we never skip rounds, then we have:*

$$R_{B,T} \leq 2^{12} \frac{\lambda^2}{\epsilon^3} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln\left(\frac{B+1}{\epsilon}\right) + O\left(\frac{K^2 \cdot \kappa}{\epsilon^3} \cdot \ln\left(\frac{B+1}{\epsilon}\right)\right)$$

and

$$R_{B,T} \leq 2^6 \frac{\lambda}{\epsilon^{3/2}} \cdot \sqrt{K \cdot \frac{B+1}{\epsilon} \cdot \ln\left(\frac{B+1}{\epsilon}\right)} + O\left(\frac{K^2 \cdot \kappa}{\epsilon^3} \cdot \ln\left(\frac{B+1}{\epsilon}\right)\right),$$

where the O notations hide universal constant factors.

Applications. Similarly, as in the case of a single resource, Assumptions 4 and 5 are natural when bidding in repeated second-price auctions if the auctioneer sets a reserve price R (which is common practice in sponsored search auctions). Indeed, we have:

$$\begin{aligned} |\mathbb{E}[c_{k,t}] - \mathbb{E}[c_{l,t}]| &= \mathbb{E}[m_t \cdot \mathbb{1}_{b_k \geq m_t > b_l}] \\ &\geq R \cdot \mathbb{E}[\mathbb{1}_{b_k \geq m_t > b_l}] \\ &\geq R \cdot \mathbb{E}[v_t \cdot \mathbb{1}_{b_k \geq m_t > b_l}] = R \cdot |\mathbb{E}[r_{k,t}] - \mathbb{E}[r_{l,t}]|, \end{aligned}$$

for any pair of arms $(k, l) \in \{1, \dots, K\}$ with $b_k \geq b_l$. Hence, Assumption 4 (resp. 5) is satisfied with $\sigma = 1/R$ (resp. $\kappa = 1/R$).

In dynamic procurement, Assumptions 4 and 5 are satisfied provided that the agents are not willing to sell their goods for less than a known price P . Indeed, in this case, pulling any arm k associated with a price $p_k \leq P$ is always suboptimal and we have:

$$\begin{aligned} |\mathbb{E}[c_{k,t}] - \mathbb{E}[c_{l,t}]| &= p_k \cdot \mathbb{P}[p_k \geq v_t] - p_l \cdot \mathbb{P}[p_l \geq v_t] \\ &\geq p_k \cdot \mathbb{P}[p_k \geq v_t > p_l] \\ &\geq P \cdot \mathbb{P}[p_k \geq v_t > p_l] = P \cdot |\mathbb{E}[r_{k,t}] - \mathbb{E}[r_{l,t}]|, \end{aligned}$$

for any pair of arms $(k, l) \in \{1, \dots, K\}$ with $p_k \geq p_l \geq P$. Hence, Assumption 4 (resp. 5) is satisfied with $\sigma = 1/P$ (resp. $\kappa = 1/P$).

In dynamic pricing, Assumptions 4 and 5 are satisfied if the distribution of valuations has a positive probability density function $f(\cdot)$. Indeed, in this case, we have:

$$\begin{aligned} |\mathbb{E}[r_{k,t}] - \mathbb{E}[r_{l,t}]| &= |p_l \cdot \mathbb{P}[p_l \leq v_t] - p_k \cdot \mathbb{P}[p_k \leq v_t]| \\ &= |p_l \cdot \mathbb{P}[p_l \leq v_t < p_k] + (p_l - p_k) \cdot \mathbb{P}[p_k \leq v_t]| \\ &\leq \max_{r=1, \dots, K} p_r \cdot \mathbb{P}[p_l \leq v_t < p_k] + |p_k - p_l| \\ &\leq \left(\max_{r=1, \dots, K} p_r + \frac{1}{\inf f(\cdot)} \right) \cdot \mathbb{P}[p_l \leq v_t < p_k] \\ &= \left(\max_{r=1, \dots, K} p_r + \frac{1}{\inf f(\cdot)} \right) \cdot |\mathbb{E}[r_{k,t}] - \mathbb{E}[r_{l,t}]|, \end{aligned}$$

for any pair of arms $(k, l) \in \{1, \dots, K\}$ with $k \geq l$. Hence, Assumption 4 (resp. 5) is satisfied with $\sigma = \max_{k=1, \dots, K} p_k + 1/\inf f(\cdot)$ (resp. $\kappa = \max_{k=1, \dots, K} p_k + 1/\inf f(\cdot)$).

7. Arbitrarily many limited resources. In this section, we tackle the general case of arbitrarily many limited resources. Additionally, we assume that one of them is time, with index $i = C$, but this assumption is almost without loss of generality, as detailed at the end of this section. To simplify the presentation, we consider the regime $K \geq C$, which is the most common in applications. This implies that $|\mathcal{K}_x| = |\mathcal{C}_x| \leq C$ for any pseudo-basis x . We also use the shorthand notation $\bar{A}_t = (\bar{c}_{k,t}(i))_{(i,k) \in \{1, \dots, C\} \times \{1, \dots, K\}}$ at any round t . For similar reasons as in Section 6, we are led to make two additional assumptions which are discussed in the last paragraph of this section.

ASSUMPTION 7. *There exists $\sigma > 0$ such that $r_{k,t} \leq \sigma \cdot \min_{i=1, \dots, C} c_{k,t}(i)$ for any arm $k \in \{1, \dots, K\}$ and for any round $t \in \mathbb{N}$.*

Note that Assumption 7 is stronger than Assumption 4 given that the amounts of resources consumed at each round have to dominate the rewards almost surely, as opposed to on average. Assumption 7 is not necessarily satisfied in all applications but it simplifies the analysis and can be relaxed at the price of an additive term of order $O(\ln^2(T))$ in the final regret bounds, see the last paragraph of this section.

ASSUMPTION 8. *There exists $\epsilon > 0$, known to the decision maker ahead of round 1, such that every basis x for (3) is ϵ -non-degenerate for (3) and satisfy $|\det(A_x)| \geq \epsilon$.*

Without loss of generality, we assume that $\epsilon \leq 1$. Observe that Assumption 8 generalizes Assumption 6 but is more restrictive because ϵ is assumed to be known to the decision maker initially. Just like in Section 6, this assumption can be relaxed to a large extent. For instance, if ϵ is initially unknown, taking ϵ as a vanishing function of T yields the same asymptotic regret bounds. However, note that Lemma 9 carries over to this more general setting and, as a result, the minimal assumption we need to get logarithmic rates is that any optimal basis for (3) is non-degenerate.

Specification of the algorithm. We implement UCB-Simplex with $\lambda = 1 + 2(C + 1)!^2/\epsilon$, $\eta_i = 0$ for any $i \in \{1, \dots, C\}$, and an initialization step which consists in pulling each arm $2^8(C + 2)!^4/\epsilon^6 \cdot \ln(T)$ times in order to get i.i.d. samples. Hence, Step-Simplex is run for the first time after round $t_{\text{ini}} = K \cdot 2^8(C + 2)!^4/\epsilon^6 \cdot \ln(T)$. Compared to Section 6, the initialization step plays as a substitute for the choice $\eta_i > 0$ which was meant to incentivize exploration. This significantly simplifies the analysis but the downside is that ϵ has to be known initially. Similarly, as in Section 6, we introduce a dummy arm which corresponds to skipping the round (i.e. pulling this arm yields a reward 0 and does not consume any resource) so that any basis can be mapped to one for which the time constraint is always binding, i.e. without loss of generality we assume that $C \in \mathcal{C}_x$ for any pseudo-basis x . Following the ideas developed in Section 6, we design load balancing algorithms for any basis x that pull arms in order to guarantee that, at any round t , the total amount of resource i consumed remains close to the target $t \cdot b(i)$ with high probability for any resource $i \in \mathcal{C}_x$. This is more involved than in Section 6 since we need to enforce this property for multiple resources but, as we show in this section, this can be done by perturbing the probability distribution solution to (8) taking into account whether we have over- or under-consumed in the past for each binding resource $i \in \mathcal{C}_x$.

Algorithm: Load balancing algorithm \mathcal{A}_x for any basis x

For any time period $t > t_{\text{ini}}$ and $i \in \mathcal{C}_x - \{C\}$, define $b_{x,t}(i)$ as the total amount of resource i consumed when selecting basis x in the past $t - 1$ rounds. Suppose that basis x is selected at time t and define the vector e_t^x by $e_{C,t}^x = 0$ and $e_{i,t}^x = -1$ (resp. $e_{i,t}^x = 1$) if $b_{x,t}(i) \geq n_{x,t} \cdot b(i)$ (resp. $b_{x,t}(i) < n_{x,t} \cdot b(i)$) for any $i \in \mathcal{C}_x - \{C\}$. Since x is selected at round t , $\bar{A}_{x,t}$ is invertible and we can define, for any $\delta \geq 0$, $p_{k,t}^x(\delta) = (\bar{A}_{x,t}^{-1}(b_{\mathcal{C}_x} + \delta \cdot e_t^x))_k$ for $k \in \mathcal{K}_x$ and $p_{k,t}^x(\delta) = 0$ otherwise, which together define the probability distribution $p_t^x(\delta) = (p_{k,t}^x(\delta))_{k \in \{1, \dots, K\}}$. Define

$$\delta_{x,t}^* = \max_{\substack{\delta \geq 0 \\ (\bar{A}_t p_t^x(\delta))_i \leq b(i), i \notin \mathcal{C}_x \\ p_t^x(\delta) \geq 0}} \delta$$

and $p_t^x = p_t^x(\delta_{x,t}^*)$. Note that $\delta_{x,t}^*$ is well defined as x must be feasible for (8) if it is selected at Step-Simplex. Pull an arm at random according to the distribution p_t^x .

Observe that the load balancing algorithms generalize the ones designed in Section 6 (up to the change $\eta_i = 0$). Indeed, when there is a single limited resource other than time, the probability distribution p_t^x is a Dirac supported at the arm with smallest (resp. largest) empirical cost when $b_{x,t} \geq n_{x,t} \cdot b$ (resp. $b_{x,t} < n_{x,t} \cdot b$). Similarly, as in Section 5, the load balancing algorithms \mathcal{A}_x may require a memory storage capacity exponential in C and polynomial in K , but, in practice, we expect that only a few bases will be selected at Step-Simplex, so that a hash table is an appropriate data structure to store the sequences $(b_{x,t}(i))_{i \in \mathcal{C}_x}$. Note, however, that the load balancing algorithms are computationally efficient because p_t^x can be computed in $O(C^2)$ running time if $\bar{A}_{x,t}^{-1}$ is available once we have computed an optimal basic feasible solution to (8), which is the case if we use the revised simplex algorithm.

Regret analysis. The regret analysis follows the same recipe as in Section 6 but the proofs are more technical and are thus deferred to Section F of the Appendix. First, we show that the initialization step guarantees that infeasible bases or pseudo-bases x with $\det(A_x) = 0$ cannot be selected more than $O(\ln(T))$ times on average at Step-Simplex.

LEMMA 13. For any basis $x \notin \mathcal{B}$, we have:

$$\mathbb{E}[n_{x,T}] \leq 2^9 \frac{(C + 3)!^4}{\epsilon^6}.$$

The same inequality holds if x is a pseudo-basis but not a basis for (3).

The next step is to show that the load balancing algorithms guarantee that, for any basis x , the amount of resource $i \in \mathcal{C}_x$ (resp. $i \notin \mathcal{C}_x$) consumed per round when selecting x remains close to (resp. below) the pacing target $b(i)$ with high probability. This enables us to show that the ratios $(\mathbb{E}[n_{k,T}^x]/\mathbb{E}[n_{l,T}^x])_{k,l \in \mathcal{K}_x}$ are close to the optimal ones $(\xi_k^x/\xi_l^x)_{k,l \in \mathcal{K}_x}$.

LEMMA 14. *For any feasible basis x and time period t , we have:*

$$\mathbb{P}[|b_{x,t}(i) - n_{x,t} \cdot b(i)| \geq u] \leq 2^5 \frac{(C+1)!^2}{\epsilon^4} \cdot \exp(-u \cdot (\frac{\epsilon^2}{4 \cdot (C+1)!})^2) + 2^9 \frac{(C+3)!^4}{\epsilon^6} \cdot \frac{1}{T} \quad \forall u \geq 1, \quad (12)$$

for any resource $i \in \mathcal{C}_x$ while

$$\mathbb{P}[b_{x,t}(i) - n_{x,t} \cdot b(i) \geq 2^8 \frac{(C+3)!^3}{\epsilon^6} \cdot \ln(T)] \leq 2^{10} \frac{(C+4)!^4}{\epsilon^6 \cdot T}, \quad (13)$$

for any resource $i \notin \mathcal{C}_x$. In particular, this implies that:

$$\mathbb{E}[n_{k,T}^x] \geq \xi_k^x \cdot \mathbb{E}[n_{x,T}] - 2^{10} \frac{(C+3)!^4}{\epsilon^9}, \quad (14)$$

for any arm $k \in \mathcal{K}_x$.

Next, we show that a suboptimal basis cannot be selected more than $O(\ln(T))$ times on average at Step-Simplex. Just like in Section 6, the exact definition of the load balancing algorithms has little impact on the result: we only need to know that, for any feasible basis x , each arm $k \in \mathcal{K}_x$ is pulled at least a fraction of the time x is selected with high probability.

LEMMA 15. *For any suboptimal basis $x \in \mathcal{B}$, we have:*

$$\mathbb{E}[n_{x,T}] \leq 2^{10} \frac{(C+3)!^3 \cdot \lambda^2}{\epsilon^6} \cdot \frac{\ln(T)}{(\Delta_x)^2} + 2^{11} \frac{(C+4)!^4}{\epsilon^6}.$$

We are now ready to derive both distribution-dependent and distribution-independent regret bounds.

THEOREM 8. *We have:*

$$R_{B(1), \dots, B(C-1), T} \leq 2^{10} \frac{(C+3)!^3 \cdot \lambda^2}{\epsilon^6} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln(T) + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right),$$

where the O notation hides universal constant factors.

THEOREM 9. *We have:*

$$R_{B(1), \dots, B(C-1), T} \leq 2^5 \frac{(C+3)!^2 \cdot \lambda}{\epsilon^3} \cdot \sqrt{|\mathcal{B}| \cdot T \cdot \ln(T)} + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right),$$

where the O notation hides universal constant factors.

Since the number of feasible bases is at most $2K^C$, we get the distribution-dependent regret bound $R_{B(1), \dots, B(C-1), T} = O(K^C \cdot (C+3)!^4/\epsilon^6 \cdot (\lambda^2/\Delta + \sigma) \cdot \ln(T))$ where $\Delta = \min_{x \in \mathcal{B} \mid \Delta_x > 0} \Delta_x$ and the distribution-independent bound $R_{B(1), \dots, B(C-1), T} = O((C+3)!^2 \cdot \lambda/\epsilon^3 \cdot \sqrt{K^C \cdot T \cdot \ln(T)})$. We stress that the dependence on K and C is not optimal since Agrawal and Devanur [1] obtain a $\tilde{O}(\sqrt{K \cdot T})$ distribution-independent bound on regret, where the \tilde{O} notation hides factors logarithmic in T . Just like in Section 6, we can also derive regret bounds that vanish with b under the assumption that pulling any arm incurs some positive amount of resource consumption in expectations for all resources, but this requires a minor tweak of the algorithm.

THEOREM 10. *Suppose that:*

$$\epsilon \leq \min_{\substack{i=1, \dots, C-1 \\ k=1, \dots, K}} \mu_k^c(i)$$

and that $b \leq \epsilon$. If the decision maker artificially constrains himself or herself to a time horizon $\tilde{T} = b \cdot T / \epsilon \leq T$, then the regret bounds derived in Theorems 8 and 9 hold with T substituted with \tilde{T} .

Similarly, if the decision maker is not constrained by a time horizon, artificially constraining himself or herself to a time horizon $\tilde{T} = \min_{i=1, \dots, C} B(i) / \epsilon$ yields the regret bounds derived in Theorems 8 and 9 with T substituted with \tilde{T} .

Applications. In dynamic pricing and online advertising applications, Assumption 7 is usually not satisfied as pulling an arm typically incurs the consumption of only a few resources. We can relax this assumption but this comes at the price of an additive term of order $O(\ln^2(T))$ in the final regret bounds.

THEOREM 11. *If Assumption 7 is not satisfied, the regret bounds derived in Theorems 8 and 9 hold with $\sigma = 0$ up to an additive term of order $O(\frac{(C+4)!^4 \cdot |B|^2}{b \cdot \epsilon^6} \cdot \ln^2(T))$.*

As for Assumption 8, the existence of degenerate optimal bases to (3) is determined by a complex interplay between the mean rewards and the mean amounts of resource consumption. However, we point out that the set of parameters $(\mu_k^r)_{k=1, \dots, K}$ and $(\mu_k^c(i))_{k=1, \dots, K, i=1, \dots, C}$ that satisfy these conditions has Lebesgue measure 0, hence such an event is unlikely to occur in practice. Additionally, while ϵ is typically not known in applications, taking ϵ as a vanishing function of T yields the same asymptotic regret bounds.

8. Concluding remark. In this paper, we develop an algorithm with a $O(K^C \cdot \ln(B) / \Delta)$ distribution-dependent bound on regret, where Δ is a parameter that generalizes the optimality gap for the standard MAB problem. It is however unclear whether the dependence on K is optimal. Extensions discussed in Section A of the Appendix suggest that it may be possible to achieve a linear dependence on K but this calls for the development of more efficient load balancing algorithms.

Appendix A: Extensions.

A.1. Improving the multiplicative factors in the regret bounds.

A.1.1. A single limited resource whose consumption is deterministic. If the amounts of resource consumed are deterministic, we can substitute the notation μ_k^c for c_k . Moreover, we can take $\lambda = 1$ and, going through the analysis of Lemma 5, we can slightly refine the regret bound. Specifically, we have $\mathbb{E}[n_{k, \tau^*}] \leq \frac{16}{(c_k)^2} \cdot \frac{\mathbb{E}[\ln(\tau^*)]}{(\Delta_k)^2} + \frac{\pi^2}{3}$, for any arm k such that $\Delta_k > 0$. Moreover, $\tau^* \leq B/\epsilon + 1$ in this setting since:

$$B \geq \sum_{t=1}^{\tau^*-1} c_{a_t, t} \geq (\tau^* - 1) \cdot \epsilon,$$

by definition of τ^* . As a result, the regret bound derived in Theorem 1 turns into:

$$R_B \leq 16 \left(\sum_{k | \Delta_k > 0} \frac{1}{c_k \cdot \Delta_k} \right) \cdot \ln\left(\frac{B}{\epsilon} + 1\right) + O(1),$$

which is identical (up to universal constant factors) to the bound obtained by Tran-Thanh et al. [24]. Note that this bound is scale-free.

A.1.2. Arbitrarily many limited resources whose consumptions are deterministic. We propose another load balancing algorithm that couples bases together. This is key to get a better dependence on K because, otherwise, we have to study each basis independently from the other ones.

Algorithm: Load balancing algorithm \mathcal{A}_x for a feasible basis $x \in \mathcal{B}$.

If x is selected at time t , pull any arm $a_t \in \arg \max_{k \in \mathcal{K}_x} \frac{\xi_k^x}{n_{k,t}}$.

Observe that this load balancing algorithm is computationally efficient with a $O(K)$ runtime (once we have computed an optimal basic feasible solution to (8)) and requires $O(K)$ memory space. The shortcoming of this approach is that, if there are multiple optimal bases to (3), the optimal load balance for each optimal basis will not be preserved since we take into account the number of times we have pulled each arm when selecting any basis (for which we strive to enforce different ratios). Hence, the following assumption will be required for the analysis.

ASSUMPTION 9. *There is a unique optimal basis to (3).*

Regret Analysis. All the proofs are deferred to Section G. We start by bounding, for each arm k , the number of times this arm can be pulled when selecting any of the suboptimal bases. This is in stark contrast with the analysis carried out in Section 5 where we bound the number of times each suboptimal basis has been selected.

LEMMA 16. *For any arm $k \in \{1, \dots, K\}$, we have:*

$$\mathbb{E} \left[\sum_{x \in \mathcal{B} \mid k \in \mathcal{K}_x, x \neq x^*} n_{k, \tau^*}^x \right] \leq 16 \frac{\rho \cdot (\sum_{i=1}^C b(i))^2}{\epsilon^2} \cdot \frac{\mathbb{E}[\ln(\tau^*)]}{(\Delta_k)^2} + K \cdot \frac{\pi^2}{3},$$

where $\Delta_k = \min_{x \in \mathcal{B} \mid k \in \mathcal{K}_x, x \neq x^*} \Delta_x$.

In contrast to Section 5, we can only guarantee that the ratios $(n_{k,t}^x/n_{l,t}^x)_{k,l \in \mathcal{K}_x}$ remain close to the optimal ones $(\xi_k^x/\xi_l^x)_{k,l \in \mathcal{K}_x}$ at all times for the optimal basis $x = x^*$. This will not allow us to derive distribution-free regret bounds for this particular class of load balancing algorithms.

LEMMA 17. *At any time t and for any arm $k \in \mathcal{K}_{x^*}$, we have:*

$$n_{k,t} \geq n_{x^*,t} \cdot \frac{\xi_k^{x^*}}{\sum_{l=1}^K \xi_l^{x^*}} - \rho \cdot \left(\sum_{x \in \mathcal{B}, x \neq x^*} n_{x,t} + 1 \right) \quad (15)$$

and

$$n_{k,t} \leq n_{x^*,t} \cdot \frac{\xi_k^{x^*}}{\sum_{l=1}^K \xi_l^{x^*}} + \sum_{x \in \mathcal{B}, x \neq x^*} n_{x,t} + 1. \quad (16)$$

Bringing everything together, we are now ready to establish regret bounds.

THEOREM 12. *We have:*

$$R_{B(1), \dots, B(C)} \leq 32 \frac{\rho^3 \cdot (\sum_{i=1}^C b(i))^3}{\epsilon^3 \cdot b} \cdot \left(\sum_{k=1}^K \frac{1}{(\Delta_k)^2} \right) \cdot \ln \left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1 \right) + O(1),$$

where the O notation hides universal constant factors.

We derive a distribution-dependent regret bound of order $O(\rho^3 \cdot K \cdot \frac{\ln(B)}{\Delta^2})$ where $\Delta = \min_{x \in \mathcal{B} \mid \Delta_x > 0} \Delta_x$ but no non-trivial distribution-free regret bound.

A.1.3. Arbitrarily many limited resources. A straightforward extension of the load balancing algorithm developed in the case of deterministic resource consumption in Section A.1.2 guarantees that the total number of times any suboptimal basis is pulled is of order $O(K \cdot \ln(T))$. However, in contrast to Section A.1.2, this is not enough to get logarithmic regret bounds as $\xi_{k,t}^x$ fluctuates around the optimal load balance $\xi_{k,t}^x$ with a magnitude of order at least $\sim 1/\sqrt{t}$, and, as a result, the ratios $(\mathbb{E}[n_{k,T}^x]/\mathbb{E}[n_{l,T}^x])_{k,l \in \mathcal{K}_x}$ might be very different from the optimal ones $(\xi_k^x/\xi_l^x)_{k,l \in \mathcal{K}_x}$.

Algorithm: Load balancing algorithm \mathcal{A}_x for a feasible basis $x \in \mathcal{B}$.

If x is selected at time t , pull any arm $a_t \in \arg \max_{k \in \mathcal{K}_x} \frac{\xi_{k,t}^x}{n_{k,t}}$.

LEMMA 18. *For any arm $k \in \{1, \dots, K\}$, we have:*

$$\mathbb{E} \left[\sum_{x \in \mathcal{B} \mid k \in \mathcal{K}_x, \Delta_x > 0} n_{k,T}^x \right] \leq 16C \cdot \lambda^2 \cdot \frac{\ln(T)}{(\Delta_k)^2} + 2^{12} \frac{K \cdot (C+3)!^2}{\epsilon^6},$$

where $\Delta_k = \min_{x \in \mathcal{B} \mid k \in \mathcal{K}_x, \Delta_x > 0} \Delta_x$.

A.2. Relaxing Assumption 1. The regret bounds obtained in Sections 5, 6, and 7 can be extended when the ratios converge as opposed to being fixed, as precisely stated below, but this requires slightly more work.

ASSUMPTION 10. *For any resource $i \in \{1, \dots, C\}$, the ratio $B(i)/B(C)$ converges to a finite value $b(i) \in (0, 1]$. Moreover, $b = \min_{i=1, \dots, C} b(i)$ is a positive quantity.*

To state the results, we need to redefine some notations and to work with the linear program:

$$\begin{aligned} & \sup_{(\xi_k)_{k=1, \dots, K}} \sum_{k=1}^K \mu_k^r \cdot \xi_k \\ & \text{subject to} \quad \sum_{k=1}^K \mu_k^c(i) \cdot \xi_k \leq \frac{B(i)}{B(C)}, \quad i = 1, \dots, C \\ & \quad \quad \quad \xi_k \geq 0, \quad k = 1, \dots, K. \end{aligned} \tag{17}$$

We redefine \mathcal{B} as the set of bases that are feasible to (17) and, for $x \in \mathcal{B}$, Δ_x is redefined as the optimality gap of x with respect to (17). We also redefine $\mathcal{O} = \{x \in \mathcal{B} \mid \Delta_x = 0\}$ as the set of optimal bases to (17). Moreover, we define \mathcal{B}_∞ (resp. \mathcal{O}_∞) as the set of feasible (resp. optimal) bases to (3) and, for $x \in \mathcal{B}_\infty$, Δ_x^∞ is the optimality gap of x with respect to (3). Our algorithm remains the same provided that we substitute $b(i)$ with $B(i)/B(C)$ for any resource $i \in \{1, \dots, C\}$. Specifically, Step-Simplex consists in solving:

$$\begin{aligned} & \sup_{(\xi_k)_{k=1, \dots, K}} \sum_{k=1}^K (\bar{r}_{k,t} + \lambda \cdot \epsilon_{k,t}) \cdot \xi_k \\ & \text{subject to} \quad \sum_{k=1}^K (\bar{c}_{k,t}(i) - \eta_i \cdot \epsilon_{k,t}) \cdot \xi_k \leq \frac{B(i)}{B(C)}, \quad i = 1, \dots, C \\ & \quad \quad \quad \xi_k \geq 0, \quad k = 1, \dots, K \end{aligned} \tag{18}$$

and Step-Load-Balance is identical to the substitution of $b(i)$ with $B(i)/B(C)$.

Regret Analysis. As it turns out, the logarithmic regret bounds established in Theorems 3, 5, and 8 do not always extend when Assumption 1 is relaxed even though these bounds appear to be very similar to the one derived in Theorem 1 when there is a single limited resource. The fundamental difference is that the set of optimal bases may not converge while it is always invariant in the case of a single limited resource. Typically, the ratios $(B(i)/B(C))_{i=1,\dots,C}$ may oscillate around $(b(i))_{i=1,\dots,C}$ in such a way that there exist two optimal bases for (3) while there is a unique optimal basis for this same optimization problem whenever the right-hand side of the inequality constraints is slightly perturbed around this limit. This alternately causes one of these two bases to be slightly suboptimal, a situation difficult to identify and to cope with for the decision maker. Nevertheless, this difficulty does not arise in several situations of interest which generalize Assumption 1, as precisely stated below. The proofs are deferred to Section G.

Arbitrarily many limited resources whose consumptions are deterministic.

THEOREM 13. *Suppose that Assumption 10 holds. If there exists a unique optimal basis to (3) or if $B(i)/B(C) - b(i) = O(\ln(B(C))/B(C))$ for all resources $i \in \{1, \dots, C-1\}$ then, we have:*

$$R_{B(1),\dots,B(C)} = O\left(\frac{\rho \cdot \sum_{i=1}^C b(i)}{\epsilon \cdot b} \cdot \left(\sum_{x \in \mathcal{B}_\infty \mid \Delta_x^\infty > 0} \frac{1}{\Delta_x^\infty}\right) \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B(C)}{\epsilon} + 1\right) + |\mathcal{O}_\infty| \cdot \frac{\ln(B(C))}{\epsilon \cdot b}\right),$$

where the O notation hides universal constant factors.

A time horizon and another limited resource.

THEOREM 14. *Suppose that Assumptions 4, 5, and 6 hold and that the ratio B/T converges to $b \in (0, 1]$. If there exists a unique optimal basis to (3) or if $B/T - b = O(\ln(T)/T)$, then, we have:*

$$R_{B,T} = O\left(\frac{\lambda^2}{\epsilon^3} \cdot \left(\sum_{x \in \mathcal{B}_\infty \mid \Delta_x^\infty > 0} \frac{1}{\Delta_x^\infty}\right) \cdot \ln(T) + \frac{K^2 \cdot \sigma}{\epsilon^3} \cdot \ln(T)\right),$$

where the O notation hides universal constant factors.

Arbitrarily many limited resources with a time horizon.

THEOREM 15. *Suppose that Assumptions 7, 8, and 10 hold. If there exists a unique optimal basis to (3) or if $B(i)/T - b(i) = O(\ln(T)/T)$ for all resources $i \in \{1, \dots, C-1\}$, then, we have:*

$$R_{B(1),\dots,B(C-1),T} = O\left(\frac{(C+3)^{13} \cdot \lambda^2}{\epsilon^6} \cdot \left(\sum_{x \in \mathcal{B}_\infty \mid \Delta_x^\infty > 0} \frac{1}{\Delta_x^\infty}\right) \cdot \ln(T) + \frac{\sigma \cdot |\mathcal{B}_\infty| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right),$$

where the O notation hides universal constant factors.

Appendix B: Proofs for Section 3.

B.1. Proof of Lemma 2. The proof can be found in Badanidiyuru et al. [8]. For the sake of completeness, we reproduce it here. The optimization problem (3) is a linear program whose dual reads:

$$\begin{aligned} & \inf_{(\zeta_i)_{i=1,\dots,C}} && \sum_{i=1}^C b(i) \cdot \zeta_i \\ & \text{subject to} && \sum_{i=1}^C \mu_k^c(i) \cdot \zeta_i \geq \mu_k^r, \quad k = 1, \dots, K \\ & && \zeta_i \geq 0, \quad i = 1, \dots, C. \end{aligned} \tag{19}$$

Observe that (3) is feasible therefore (3) and (19) have the same optimal value. Note that (3) is bounded under Assumption 2 as $\xi_k \in [0, b(i)/\mu_k^c(i)]$ for any feasible point and any resource $i \in \{1, \dots, C\}$ such that $\mu_k^c(i) > 0$. Hence, (19) has an optimal basic feasible solution $(\zeta_1^*, \dots, \zeta_C^*)$. Consider any non-anticipating algorithm. Let Z_t be the sum of the total payoff accumulated in rounds 1 to t plus the ‘‘cost’’ of the remaining resources, i.e. $Z_t = \sum_{\tau=1}^t r_{a_{\tau}, \tau} + \sum_{i=1}^C \zeta_i^* \cdot (B(i) - \sum_{\tau=1}^t c_{a_{\tau}, \tau}(i))$. Observe that $(Z_t)_t$ is a supermartingale with respect to the filtration $(\mathcal{F}_t)_t$ as $\mathbb{E}[Z_t | \mathcal{F}_{t-1}] = \mathbb{E}[\mu_{a_t}^r - \sum_{i=1}^C \zeta_i^* \cdot \mu_{a_t}^c(i) | \mathcal{F}_{t-1}] + Z_{t-1} \leq Z_{t-1}$ since $(\zeta_1^*, \dots, \zeta_C^*)$ is feasible for (19). Moreover, note that $(Z_t)_t$ has bounded increments since $|Z_t - Z_{t-1}| = |r_{a_t, t} - \sum_{i=1}^C \zeta_i^* \cdot c_{a_t, t}(i)| \leq 1 + \sum_{i=1}^C \zeta_i^* < \infty$. We also have $\mathbb{E}[\tau^*] < \infty$ as:

$$\begin{aligned} \mathbb{E}[\tau^*] &= \sum_{t=1}^{\infty} \mathbb{P}[\tau^* \geq t] \\ &\leq \sum_{t=1}^{\infty} \mathbb{P}\left[\sum_{\tau=1}^{t-1} c_{a_{\tau}, \tau}(i) \leq B(i), i=1, \dots, C\right] \\ &\leq 1 + \sum_{t=1}^{\infty} \mathbb{P}\left[\sum_{\tau=1}^t \sum_{i=1}^C c_{a_{\tau}, \tau}(i) \leq t \cdot \epsilon - (t \cdot \epsilon - \sum_{i=1}^C B(i))\right] \\ &\leq \left(\frac{\sum_{i=1}^C B(i)}{\epsilon} + 2\right) + \sum_{t \geq \frac{\sum_{i=1}^C B(i)}{\epsilon}}^{\infty} \exp\left(-\frac{2(t \cdot \epsilon - \sum_{i=1}^C B(i))^2}{t}\right) \\ &< \infty, \end{aligned}$$

where the third inequality results from an application of Lemma 1 and

$$\epsilon = \min_{\substack{k=1, \dots, K \\ i=1, \dots, C \\ \text{with } \mu_k^c(i) > 0}} \mu_k^c(i).$$

By Doob’s optional stopping theorem, $\mathbb{E}[Z_{\tau^*}] \leq \mathbb{E}[Z_0] = \sum_{i=1}^C \zeta_i^* \cdot B(i)$. Observe that:

$$\begin{aligned} \mathbb{E}[Z_{\tau^*}] &= \mathbb{E}\left[r_{a_{\tau^*}, \tau^*} - \sum_{i=1}^C \zeta_i^* \cdot c_{a_{\tau^*}, \tau^*}(i) + Z_{\tau^*-1}\right] \\ &\geq \mathbb{E}\left[-\sum_{i=1}^C \zeta_i^* + \sum_{t=1}^{\tau^*-1} r_{a_t, t}\right]. \end{aligned}$$

Using Assumption 2 and since $(\zeta_i^*)_{i=1, \dots, C}$ is a basic feasible solution, for every $i \in \{1, \dots, C\}$ such that $\zeta_i^* > 0$ there must exist $k \in \{1, \dots, K\}$ such that $\zeta_i^* \leq \mu_k^r / \mu_k^c(i)$ with $\mu_k^c(i) > 0$. We get:

$$\mathbb{E}[Z_{\tau^*}] \geq \mathbb{E}\left[\sum_{t=1}^{\tau^*-1} r_{a_t, t}\right] - \max_{\substack{k=1, \dots, K \\ i=1, \dots, C \\ \text{with } \mu_k^c(i) > 0}} \frac{\mu_k^r}{\mu_k^c(i)}$$

and finally:

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^{\tau^*-1} r_{a_t, t}\right] &\leq \sum_{i=1}^C \zeta_i^* \cdot B(i) + \max_{\substack{k=1, \dots, K \\ i=1, \dots, C \\ \text{with } \mu_k^c(i) > 0}} \frac{\mu_k^r}{\mu_k^c(i)} \\ &= B \cdot \sum_{i=1}^C \zeta_i^* \cdot b(i) + \max_{\substack{k=1, \dots, K \\ i=1, \dots, C \\ \text{with } \mu_k^c(i) > 0}} \frac{\mu_k^r}{\mu_k^c(i)}. \end{aligned}$$

By strong duality, $\sum_{i=1}^C \zeta_i^* \cdot b(i)$ is also the optimal value of (3).

Appendix C: Proofs for Section 4.

C.1. Proof of Lemma 4. By definition of τ^* , we have $\sum_{t=1}^{\tau^*-1} c_{a_t,t} \leq B$. Taking expectations on both sides yields:

$$\begin{aligned}
B &\geq \mathbb{E}\left[\sum_{t=1}^{\tau^*-1} c_{a_t,t}\right] \\
&\geq \mathbb{E}\left[\sum_{t=1}^{\tau^*} c_{a_t,t}\right] - 1 \\
&= \sum_{t=1}^{\infty} \mathbb{E}[I_{\tau^* \geq t} \cdot c_{a_t,t}] - 1 \\
&= \sum_{t=1}^{\infty} \mathbb{E}[I_{\tau^* \geq t} \cdot \mathbb{E}[c_{a_t,t} \mid \mathcal{F}_{t-1}]] - 1 \\
&= \sum_{t=1}^{\infty} \mathbb{E}[I_{\tau^* \geq t} \cdot \mu_{a_t}^c] - 1 \\
&\geq \sum_{t=1}^{\infty} \mathbb{E}[I_{\tau^* \geq t} \cdot \epsilon] - 1 \\
&= \mathbb{E}[\tau^*] \cdot \epsilon - 1,
\end{aligned}$$

where we use the fact that $c_{k,t} \leq 1$ for all arms k to derive the second inequality, the fact that τ^* is a stopping time for the second equality, the fact that a_t is deterministically determined by the past, i.e. $a_t \in \mathcal{F}_{t-1}$, for the third equality and Assumption 2 for the third inequality. We conclude that $\mathbb{E}[\tau^*] \leq \frac{B+1}{\epsilon}$.

C.2. Proof of Lemma 5. We break down the analysis in a series of facts. Consider any arm k such that $\Delta_k > 0$. We use the shorthand notation $\beta_k = 2^5 \left(\frac{\lambda}{\mu_k^c}\right)^2 \cdot \left(\frac{1}{\Delta_k}\right)^2$.

Fact 1

$$\mathbb{E}[n_{k,\tau^*}] \leq 2\beta_k \cdot \mathbb{E}[\ln(\tau^*)] + \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{a_t=k} \cdot I_{n_{k,t} \geq \beta_k \ln(t)}\right]. \quad (20)$$

Proof. Define the random variable $T_k = \beta_k \cdot \ln(\tau^*)$. We have:

$$\begin{aligned}
\mathbb{E}[n_{k,\tau^*}] &= \mathbb{E}[n_{k,\tau^*} \cdot I_{n_{k,\tau^*} < T_k}] + \mathbb{E}[n_{k,\tau^*} \cdot I_{n_{k,\tau^*} \geq T_k}] \\
&\leq \beta_k \cdot \mathbb{E}[\ln(\tau^*)] + \mathbb{E}[n_{k,\tau^*} \cdot I_{n_{k,\tau^*} \geq T_k}].
\end{aligned}$$

Define T_k^* as the first time t such that $n_{k,t} \geq T_k$ and $T_k^* = \infty$ if no such t exists. We have:

$$\begin{aligned}
\mathbb{E}[n_{k,\tau^*} \cdot I_{n_{k,\tau^*} \geq T_k}] &= \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{a_t=k} \cdot I_{n_{k,\tau^*} \geq T_k}\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T_k^*-1} I_{a_t=k} \cdot I_{n_{k,\tau^*} \geq T_k}\right] + \mathbb{E}\left[\sum_{t=T_k^*}^{\tau^*} I_{a_t=k} \cdot I_{n_{k,\tau^*} \geq T_k}\right] \\
&\leq \mathbb{E}[n_{k,T_k^*-1} \cdot I_{n_{k,\tau^*} \geq T_k}] + \mathbb{E}\left[\sum_{t=T_k^*}^{\tau^*} I_{a_t=k} \cdot I_{n_{k,t} \geq T_k}\right] \\
&\leq \beta_k \cdot \mathbb{E}[\ln(\tau^*)] + \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{a_t=k} \cdot I_{n_{k,t} \geq \beta_k \ln(t)}\right],
\end{aligned}$$

since, by definition of T_k^* , $n_{k, T_k^* - 1} \leq T_k$ if T_k^* is finite, which is always true if $n_{k, \tau^*} \geq T_k$ (the sequence $(n_{k, t})_t$ is non-decreasing and τ^* is finite almost surely as a byproduct of Lemma 4). Conversely, $n_{k, t} \geq T_k \geq \beta_k \ln(t)$ for $t \in \{T_k^*, \dots, \tau^*\}$. Wrapping up, we obtain:

$$\mathbb{E}[n_{k, \tau^*}] \leq 2\beta_k \cdot \mathbb{E}[\ln(\tau^*)] + \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{a_t=k} \cdot I_{n_{k, t} \geq \beta_k \ln(t)}\right].$$

Fact 1 enables us to assume that arm k has been pulled at least $\beta_k \ln(t)$ times out of the last t time periods. The remainder of this proof is dedicated to show that the second term of the right-hand side of (20) can be bounded by a constant. Let us first rewrite this term:

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{a_t=k} \cdot I_{n_{k, t} \geq \beta_k \ln(t)}\right] &\leq \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k, t} + E_{k, t} \geq \text{obj}_{k^*, t} + E_{k^*, t}} \cdot I_{n_{k, t} \geq \beta_k \ln(t)}\right] \\ &\leq \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k, t} \geq \text{obj}_k + E_{k, t}}\right] \end{aligned} \quad (21)$$

$$+ \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k^*, t} \leq \text{obj}_{k^*} - E_{k^*, t}}\right] \quad (22)$$

$$+ \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k^*} < \text{obj}_k + 2E_{k, t}} \cdot I_{n_{k, t} \geq \beta_k \ln(t)}\right]. \quad (23)$$

To derive this last inequality, simply observe that if $\text{obj}_{k, t} < \text{obj}_k + E_{k, t}$ and $\text{obj}_{k^*, t} > \text{obj}_{k^*} - E_{k^*, t}$ while $\text{obj}_{k, t} + E_{k, t} \geq \text{obj}_{k^*, t} + E_{k^*, t}$, it must be that $\text{obj}_{k^*} < \text{obj}_k + 2E_{k, t}$. Let us study (21), (22), and (23) separately.

Fact 2

$$\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k^*} < \text{obj}_k + 2E_{k, t}} \cdot I_{n_{k, t} \geq \beta_k \ln(t)}\right] \leq \frac{2\pi^2}{3\epsilon^2}.$$

Proof. Observe that when both $n_{k, t} \geq \beta_k \ln(t)$ and $\text{obj}_{k^*} < \text{obj}_k + 2E_{k, t}$, we have:

$$\begin{aligned} \frac{\Delta_k}{2} &< E_{k, t} \\ &\leq \frac{\lambda}{\bar{c}_{k, t}} \cdot \sqrt{\frac{2}{\beta_k}} \\ &\leq \frac{\mu_k^c}{2\bar{c}_{k, t}} \cdot \frac{\Delta_k}{2}, \end{aligned}$$

by definition of β_k . This implies that $\bar{c}_{k, t} \leq \mu_k^c/2$. Thus:

$$\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k^*} < \text{obj}_k + 2E_{k, t}} \cdot I_{n_{k, t} \geq \beta_k \ln(t)}\right] \leq \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\bar{c}_{k, t} < \mu_k^c/2} \cdot I_{n_{k, t} \geq \beta_k \ln(t)}\right]$$

We upper bound this last term using the concentration inequalities of Lemma 1 observing that:

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^{\infty} I_{\bar{c}_{k, t} < \mu_k^c/2} \cdot I_{n_{k, t} \geq \beta_k \ln(t)}\right] &= \sum_{t=1}^{\infty} \mathbb{P}\left[\bar{c}_{k, t} < \frac{\mu_k^c}{2}; n_{k, t} \geq \beta_k \ln(t)\right] \\ &\leq \sum_{t=1}^{\infty} \sum_{s=\beta_k \ln(t)}^t \mathbb{P}\left[\bar{c}_{k, t} < \mu_k^c - \frac{\mu_k^c}{2}; n_{k, t} = s\right]. \end{aligned}$$

Denote by t_1, \dots, t_s the first s random times at which arm k is pulled (these random variables are finite almost surely). We have:

$$\mathbb{P}[\bar{c}_{k,t} < \mu_k^c - \frac{\mu_k^c}{2}; n_{k,t} = s] \leq \mathbb{P}\left[\sum_{l=1}^s c_{k,t_l} < s \cdot \mu_k^c - s \cdot \frac{\mu_k^c}{2}\right].$$

Observe that, for any $l \leq s$:

$$\begin{aligned} \mathbb{E}[c_{k,t_l} \mid c_{k,t_1}, \dots, c_{k,t_{l-1}}] &= \mathbb{E}\left[\sum_{\tau=1}^{\infty} I_{t_l=\tau} \cdot \mathbb{E}[c_{k,\tau} \mid \mathcal{F}_{\tau-1}] \mid c_{k,t_1}, \dots, c_{k,t_{l-1}}\right] \\ &= \mathbb{E}\left[\sum_{\tau=1}^{\infty} I_{t_l=\tau} \cdot \mu_k^c \mid c_{k,t_1}, \dots, c_{k,t_{l-1}}\right] \\ &= \mu_k^c, \end{aligned}$$

since the algorithm is not randomized ($\{t_l = \tau\} \in \mathcal{F}_{\tau-1}$) and using the tower property. Hence, we can apply Lemma 1 to get:

$$\begin{aligned} \sum_{t=1}^{\infty} \mathbb{P}[\bar{c}_{k,t} < \frac{\mu_k^c}{2}; n_{k,t} \geq \beta_k \ln(t)] &\leq \sum_{t=1}^{\infty} \sum_{s=\beta_k \ln(t)}^{\infty} \exp(-s \cdot \frac{(\mu_k^c)^2}{2}) \\ &\leq \sum_{t=1}^{\infty} \frac{\exp(-\frac{(\mu_k^c)^2}{2} \cdot \beta_k \ln(t))}{1 - \exp(-\frac{(\mu_k^c)^2}{2})} \\ &\leq \frac{1}{1 - \exp(-\frac{(\mu_k^c)^2}{2})} \sum_{t=1}^{\infty} \frac{1}{t^2} \\ &\leq \frac{2\pi^2}{3(\mu_k^c)^2} \\ &\leq \frac{2\pi^2}{3\epsilon^2}, \end{aligned}$$

where we use the fact that $\beta_k \geq 2^5 \left(\frac{1+\kappa}{\mu_k^c}\right)^2 \cdot \left(\frac{\mu_k^c}{\mu_k^*}\right)^2 \geq 2^5 \left(\frac{1+\frac{1}{\kappa}}{\mu_k^c}\right)^2 \geq \frac{4}{(\mu_k^c)^2}$ for the third inequality (using Assumption 3), the fact that $\exp(-x) \leq 1 - \frac{x}{2}$ for $x \in [0, 1]$ for the fourth inequality, and Assumption 2 for the last inequality.

Let us now elaborate on (21).

Fact 3

$$\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k,t} \geq \text{obj}_k + E_{k,t}}\right] \leq \frac{\pi^2}{3}.$$

Proof. Note that if $\bar{r}_{k,t}/\bar{c}_{k,t} = \text{obj}_{k,t} \geq \text{obj}_k + E_{k,t} = \mu_k^r/\mu_k^c + E_{k,t}$, then either $\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}$ or $\bar{c}_{k,t} \leq \mu_k^c - \epsilon_{k,t}$, otherwise we would have:

$$\begin{aligned} \frac{\bar{r}_{k,t}}{\bar{c}_{k,t}} - \frac{\mu_k^r}{\mu_k^c} &= \frac{(\bar{r}_{k,t} - \mu_k^r)\mu_k^c + (\mu_k^c - \bar{c}_{k,t})\mu_k^r}{\bar{c}_{k,t} \cdot \mu_k^c} \\ &< \frac{\epsilon_{k,t}}{\bar{c}_{k,t}} + \frac{\epsilon_{k,t}}{\bar{c}_{k,t}} \cdot \frac{\mu_k^r}{\mu_k^c} \\ &\leq (1 + \kappa) \cdot \frac{\epsilon_{k,t}}{\bar{c}_{k,t}} \\ &= E_{k,t}, \end{aligned}$$

a contradiction. Therefore:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k,t} \geq \text{obj}_k + E_{k,t}}\right] &\leq \sum_{t=1}^{\infty} \mathbb{P}[\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}] + \mathbb{P}[\bar{c}_{k,t} \leq \mu_k^c - \epsilon_{k,t}] \\
&\leq \sum_{t=1}^{\infty} \sum_{s=1}^t \mathbb{P}[\bar{r}_{k,t} \geq \mu_k^r + \sqrt{\frac{2 \ln(t)}{s}}; n_{k,t} = s] \\
&\quad + \sum_{t=1}^{\infty} \sum_{s=1}^t \mathbb{P}[\bar{c}_{k,t} \leq \mu_k^c - \sqrt{\frac{2 \ln(t)}{s}}; n_{k,t} = s] \\
&= \sum_{t=1}^{\infty} \sum_{s=1}^t \mathbb{P}\left[\sum_{l=1}^s r_{k,t_l} \geq s \cdot \mu_k^r + \sqrt{s \cdot 2 \ln(t)}; n_{k,t} = s\right] \\
&\quad + \sum_{t=1}^{\infty} \sum_{s=1}^t \mathbb{P}\left[\sum_{l=1}^s c_{k,t_l} \leq s \cdot \mu_k^c - \sqrt{s \cdot 2 \ln(t)}; n_{k,t} = s\right] \\
&\leq \sum_{t=1}^{\infty} \sum_{s=1}^t 2 \exp(-4 \ln(t)) \\
&= \frac{\pi^2}{3},
\end{aligned}$$

where the random variables $(t_l)_l$ are defined similarly as in the proof of Fact 2 and the last inequality results from an application of Lemma 1.

What remains to be done is to bound (22).

Fact 4

$$\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k^*,t} \leq \text{obj}_{k^*} - E_{k^*,t}}\right] \leq \frac{\pi^2}{3}.$$

Proof. We proceed along the same lines as in the proof of Fact 3. As a matter of fact, the situation is perfectly symmetric because, in the course of proving Fact 3, we do not rely on the fact that we have pulled arm k more than $\beta_k \ln(t)$ times at any time t . If $\bar{r}_{k^*,t}/\bar{c}_{k^*,t} = \text{obj}_{k^*,t} \leq \text{obj}_{k^*} - E_{k^*,t} = \mu_{k^*}^r/\mu_{k^*}^c - E_{k^*,t}$, then we have either $\bar{r}_{k^*,t} \leq \mu_{k^*}^r - \epsilon_{k^*,t}$ or $\bar{c}_{k^*,t} \geq \mu_{k^*}^c + \epsilon_{k^*,t}$, otherwise we would have:

$$\begin{aligned}
\frac{\bar{r}_{k^*,t}}{\bar{c}_{k^*,t}} - \frac{\mu_{k^*}^r}{\mu_{k^*}^c} &= \frac{(\bar{r}_{k^*,t} - \mu_{k^*}^r)\mu_{k^*}^c + (\mu_{k^*}^c - \bar{c}_{k^*,t})\mu_{k^*}^r}{\bar{c}_{k^*,t} \cdot \mu_{k^*}^c} \\
&> -\frac{\epsilon_{k^*,t}}{\bar{c}_{k^*,t}} - \frac{\epsilon_{k^*,t}}{\bar{c}_{k^*,t}} \cdot \frac{\mu_{k^*}^r}{\mu_{k^*}^c} \\
&\geq -(1 + \kappa) \cdot \frac{\epsilon_{k^*,t}}{\bar{c}_{k^*,t}} \\
&= -E_{k^*,t},
\end{aligned}$$

a contradiction. Therefore:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{k^*,t} \leq \text{obj}_{k^*} - E_{k^*,t}}\right] &\leq \mathbb{E}\left[\sum_{t=1}^{\infty} I_{\bar{r}_{k^*,t} \leq \mu_{k^*}^r - \epsilon_{k^*,t}} + I_{\bar{c}_{k^*,t} \geq \mu_{k^*}^c + \epsilon_{k^*,t}}\right] \\
&\leq \sum_{t=1}^{\infty} \sum_{s=1}^t \mathbb{P}[\bar{r}_{k^*,t} \leq \mu_{k^*}^r - \sqrt{\frac{2 \ln(t)}{s}}; n_{k^*,t} = s] \\
&\quad + \sum_{t=1}^{\infty} \sum_{s=1}^t \mathbb{P}[\bar{c}_{k^*,t} \geq \mu_{k^*}^c + \sqrt{\frac{2 \ln(t)}{s}}; n_{k^*,t} = s]
\end{aligned}$$

$$\begin{aligned} &\leq \sum_{t=1}^{\infty} \sum_{s=1}^t \frac{2}{t^4} \\ &= \frac{\pi^2}{3}, \end{aligned}$$

where the third inequality is obtained using Lemma 1 in the same fashion as in Fact 3.

We conclude:

$$\mathbb{E}[n_{k,\tau^*}] \leq 2\beta_k \cdot \mathbb{E}[\ln(\tau^*)] + \frac{4\pi^2}{3\epsilon^2}.$$

C.3. Proof of Theorem 1. First observe that:

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t,t}\right] &= \sum_{t=1}^{\infty} \mathbb{E}[I_{\tau^* \geq t} \cdot \mathbb{E}[r_{a_t,t} \mid \mathcal{F}_{t-1}]] \\ &= \sum_{t=1}^{\infty} \mathbb{E}[I_{\tau^* \geq t} \cdot \mu_{a_t}^r] \\ &= \sum_{t=1}^{\infty} \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[I_{\tau^* \geq t} \cdot I_{a_t=k}] \\ &= \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}\left[\sum_{t=1}^{\infty} I_{\tau^* \geq t} \cdot I_{a_t=k}\right] \\ &= \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[n_{k,\tau^*}], \end{aligned}$$

since τ^* is a stopping time. Plugging this equality into (9) yields:

$$\begin{aligned} R_B &\leq B \cdot \frac{\mu_{k^*}^r}{\mu_{k^*}^c} - \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[n_{k,\tau^*}] + O(1) \\ &= \frac{\mu_{k^*}^r}{\mu_{k^*}^c} \cdot (B - \sum_{k \mid \Delta_k=0} \mu_k^c \cdot \mathbb{E}[n_{k,\tau^*}]) - \sum_{k \mid \Delta_k>0} \mu_k^r \cdot \mathbb{E}[n_{k,\tau^*}] + O(1). \end{aligned}$$

Along the same lines as for the rewards, we can show that $\mathbb{E}[\sum_{t=1}^{\tau^*} c_{a_t,t}] = \sum_{k=1}^K \mu_k^c \cdot \mathbb{E}[n_{k,\tau^*}]$. By definition of τ^* , we have $B \leq \sum_{t=1}^{\tau^*} c_{a_t,t}$ almost surely. Taking expectations on both sides yields:

$$\begin{aligned} B &\leq \sum_{k=1}^K \mu_k^c \cdot \mathbb{E}[n_{k,\tau^*}] \\ &= \sum_{k \mid \Delta_k=0} \mu_k^c \cdot \mathbb{E}[n_{k,\tau^*}] + \sum_{k \mid \Delta_k>0} \mu_k^c \cdot \mathbb{E}[n_{k,\tau^*}]. \end{aligned}$$

Plugging this inequality back into the regret bound, we get:

$$\begin{aligned} R_B &\leq \sum_{k \mid \Delta_k>0} \left(\frac{\mu_{k^*}^r}{\mu_{k^*}^c} \cdot \mu_k^c - \mu_k^r \right) \cdot \mathbb{E}[n_{k,\tau^*}] + O(1) \\ &= \sum_{k \mid \Delta_k>0} \mu_k^c \cdot \Delta_k \cdot \mathbb{E}[n_{k,\tau^*}] + O(1). \end{aligned} \tag{24}$$

Using the upper bound of Lemma 4, the concavity of the logarithmic function, and Lemma 5, we derive:

$$\begin{aligned} R_B &\leq 2^6 \lambda^2 \cdot \left(\sum_{k \mid \Delta_k > 0} \frac{1}{\mu_k^c \cdot \Delta_k} \right) \cdot \ln\left(\frac{B+1}{\epsilon}\right) + \frac{4\pi^2}{3\epsilon^2} \cdot \left(\sum_{k \mid \Delta_k > 0} \mu_k^c \cdot \Delta_k \right) + O(1) \\ &\leq 2^6 \lambda^2 \cdot \left(\sum_{k \mid \Delta_k > 0} \frac{1}{\mu_k^c \cdot \Delta_k} \right) \cdot \ln\left(\frac{B+1}{\epsilon}\right) + K \cdot \frac{4\pi^2 \kappa}{3\epsilon^2} + O(1), \end{aligned}$$

since $\Delta_k \leq \mu_{k^*}^r / \mu_{k^*}^c \leq \kappa$ and $\mu_k^c \leq 1$ for any arm k .

Appendix D: Proofs for Section 5.

D.1. Proof of Lemma 8. Consider any suboptimal basis $x \in \mathcal{B}$. The proof is along the same lines as for Lemma 5 and follows the exact same steps. We use the shorthand notation $\beta_x = 8\rho \cdot \left(\frac{\sum_{k=1}^K \xi_k^x}{\Delta_x}\right)^2$.

Fact 5

$$\mathbb{E}[n_{x,\tau^*}] \leq 2\beta_x \cdot \mathbb{E}[\ln(\tau^*)] + \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right]. \quad (25)$$

We omit the proof as it is analogous to the proof of Fact 1. As in Lemma 5, we break down the second term in three terms and bound each of them by a constant:

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] &= \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x,t} + E_{x,t} \geq \text{obj}_{x^*,t} + E_{x^*,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\ &\leq \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}}\right] \end{aligned} \quad (26)$$

$$+ \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x^*,t} \leq \text{obj}_{x^*} - E_{x^*,t}}\right] \quad (27)$$

$$+ \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right]. \quad (28)$$

Fact 6

$$\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] = 0.$$

Proof. If $\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}$, we get:

$$\begin{aligned} \frac{\Delta_x}{2} &< \sum_{k \in \mathcal{K}_x} \xi_k^x \cdot \sqrt{\frac{2 \ln(t)}{n_{k,t}}} \\ &\leq \sum_{k \in \mathcal{K}_x} \xi_k^x \cdot \sqrt{\frac{2 \ln(t)}{\rho + n_{k,t}}} \\ &\leq \sqrt{\sum_{k \in \mathcal{K}_x} \xi_k^x} \cdot \sum_{k \in \mathcal{K}_x} \sqrt{\xi_k^x} \cdot \sqrt{\frac{2 \ln(t)}{n_{x,t}}}, \end{aligned}$$

where we use (10) and Lemma 6 for each $k \in \mathcal{K}_x$ such that $\xi_k^x \neq 0$ (otherwise, if $\xi_k^x = 0$, the inequality is trivial). This implies:

$$\begin{aligned} n_{x,t} &< 8\rho \cdot \left(\frac{\sum_{k=1}^K \xi_k^x}{\Delta_x} \right)^2 \cdot \ln(t) \\ &= \beta_x \cdot \ln(t), \end{aligned}$$

using the Cauchy–Schwarz inequality and the fact that a basis involves at most ρ arms.

Fact 7

$$\mathbb{E} \left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \right] \leq \rho \cdot \frac{\pi^2}{6}.$$

Proof. If $\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}$, there must exist $k \in \mathcal{K}_x$ such that $\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}$, otherwise:

$$\begin{aligned} \text{obj}_{x,t} - \text{obj}_x &= \sum_{k \in \mathcal{K}_x} (\bar{r}_{k,t} - \mu_k^r) \cdot \xi_k^x \\ &< \sum_{k \in \mathcal{K}_x} \epsilon_{k,t} \cdot \xi_k^x \\ &= E_{x,t}, \end{aligned}$$

where the inequality is strict because there must exist $l \in \mathcal{K}_x$ such that $\xi_l^x > 0$ as a result of Assumption 2 (at least one resource constraint is binding for a feasible basis to (3) aside from the basis \tilde{x} associated with $\mathcal{K}_{\tilde{x}} = \emptyset$). We obtain:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \right] &\leq \sum_{k \in \mathcal{K}_x} \sum_{t=1}^{\infty} \mathbb{P}[\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}] \\ &\leq \rho \cdot \frac{\pi^2}{6}, \end{aligned}$$

where the last inequality is derived along the same lines as in the proof of Fact 3.

Fact 8

$$\mathbb{E} \left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x^*,t} \leq \text{obj}_{x^*} - E_{x^*,t}} \right] \leq \rho \cdot \frac{\pi^2}{6}.$$

Proof. Similar to Fact 7.

D.2. Proof of Theorem 3. The proof proceeds along the same lines as for Theorem 1. We build upon (4):

$$\begin{aligned} R_{B(1), \dots, B(C)} &\leq B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E} \left[\sum_{t=1}^{\tau^*} r_{a_t,t} \right] + O(1) \\ &= B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{t=1}^{\infty} \mathbb{E} [I_{\tau^* \geq t} \cdot \sum_{k=1}^K \sum_{x \in \mathcal{B}} r_{k,t} \cdot I_{x_t=x, a_t=k}] + O(1) \\ &= B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{t=1}^{\infty} \mathbb{E} [I_{\tau^* \geq t} \cdot \sum_{k=1}^K \sum_{x \in \mathcal{B}} I_{x_t=x, a_t=k} \cdot \mathbb{E}[r_{k,t} \mid \mathcal{F}_{t-1}]] + O(1) \end{aligned}$$

$$\begin{aligned}
&= B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \sum_{k=1}^K \mu_k^r \cdot \mathbb{E} \left[\sum_{t=1}^{\tau^*} I_{x_t=x, a_t=k} \right] + O(1) \\
&= B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \sum_{k=1}^K \mu_k^r \cdot \mathbb{E} [n_{x, \tau^*}^x] + O(1),
\end{aligned}$$

where we use the fact that x_t and a_t are determined by the events of the first $t - 1$ rounds and that τ^* is a stopping time. Using the properties of the load balancing algorithm established in Lemma 6, we derive:

$$\begin{aligned}
R_{B(1), \dots, B(C)} &\leq B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \sum_{k \in \mathcal{K}_x} \left\{ \mu_k^r \cdot \frac{\xi_k^x}{\sum_{l=1}^K \xi_l^x} \cdot \mathbb{E} [n_{x, \tau^*}] - \rho \right\} + O(1) \\
&= B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \left\{ \frac{\mathbb{E} [n_{x, \tau^*}]}{\sum_{k=1}^K \xi_k^x} \cdot \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x \right) - (\rho)^2 \right\} + O(1) \\
&= \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \right) \cdot \left(B - \sum_{x \in \mathcal{B} \mid \Delta_x=0} \frac{\mathbb{E} [n_{x, \tau^*}]}{\sum_{k=1}^K \xi_k^x} \right) \\
&\quad - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \left\{ \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x \right) \cdot \frac{\mathbb{E} [n_{x, \tau^*}]}{\sum_{k=1}^K \xi_k^x} \right\} + O(1).
\end{aligned}$$

Now observe that, by definition, at least one resource is exhausted at time τ^* . Hence, there exists $i \in \{1, \dots, C\}$ such that the following holds almost surely:

$$\begin{aligned}
B(i) &\leq \sum_{x \in \mathcal{B}} \sum_{k \in \mathcal{K}_x} c_k(i) \cdot n_{k, \tau^*}^x \\
&\leq \sum_{x \in \mathcal{B}} \sum_{k \in \mathcal{K}_x} \left[c_k(i) \cdot \left(\frac{\xi_k^x}{\sum_{l=1}^K \xi_l^x} \cdot n_{x, \tau^*} + 1 \right) \right] \\
&= |\mathcal{B}| \cdot \rho + \sum_{x \in \mathcal{B}} \frac{n_{x, \tau^*}}{\sum_{k=1}^K \xi_k^x} \cdot \sum_{k \in \mathcal{K}_x} c_k(i) \cdot \xi_k^x \\
&\leq |\mathcal{B}| \cdot \rho + b(i) \cdot \sum_{x \in \mathcal{B}} \frac{n_{x, \tau^*}}{\sum_{k=1}^K \xi_k^x},
\end{aligned}$$

where we use Lemma 6 again and the fact that any basis $x \in \mathcal{B}$ satisfies all the constraints of (3). We conclude that the inequality:

$$\sum_{x \in \mathcal{B} \mid \Delta_x=0} \frac{n_{x, \tau^*}}{\sum_{k=1}^K \xi_k^x} \geq B - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{n_{x, \tau^*}}{\sum_{k=1}^K \xi_k^x} - \frac{|\mathcal{B}| \cdot \rho}{b}$$

holds almost surely. Taking expectations on both sides and plugging the result back into the regret bound yields:

$$R_{B(1), \dots, B(C)} \leq \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{\left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{k=1}^K \mu_k^r \cdot \xi_k^x \right)}{\sum_{k=1}^K \xi_k^x} \cdot \mathbb{E} [n_{x, \tau^*}] \quad (29)$$

$$\begin{aligned}
&+ \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \right) \cdot \frac{|\mathcal{B}| \cdot \rho}{b} + O(1) \\
&\leq \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{\Delta_x}{\sum_{k=1}^K \xi_k^x} \cdot \mathbb{E} [n_{x, \tau^*}] + O(1), \quad (30)
\end{aligned}$$

where we use the fact that:

$$\begin{aligned} \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} &\leq \sum_{k=1}^K \frac{\sum_{i=1}^C c_k(i)}{\epsilon} \cdot \xi_k^{x^*} \\ &= \frac{1}{\epsilon} \cdot \sum_{i=1}^C \sum_{k=1}^K c_k(i) \cdot \xi_k^{x^*} \\ &\leq \frac{\sum_{i=1}^C b(i)}{\epsilon}, \end{aligned} \tag{31}$$

using Assumption 2 and the fact that x^* is a feasible basis to (3). Using Lemma 7, Lemma 8, and the concavity of the logarithmic function, we obtain:

$$\begin{aligned} R_{B(1), \dots, B(C)} &\leq 16\rho \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{\sum_{k=1}^K \xi_k^x}{\Delta_x} \right) \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1 \right) \\ &\quad + \frac{\pi^2}{3} \rho \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{\Delta_x}{\sum_{k=1}^K \xi_k^x} \right) + O(1) \\ &\leq 16 \frac{\rho \cdot \sum_{i=1}^C b(i)}{\epsilon} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1 \right) + O(1). \end{aligned}$$

To derive this last inequality, we use: (i) $\Delta_x \leq \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \leq \sum_{i=1}^C b(i)/\epsilon$ (see (31)), (ii) the fact that, for any basis $x \in \mathcal{B}$, at least one of the first C inequalities is binding in (3), which implies that there exists $i \in \{1, \dots, C\}$ such that:

$$\begin{aligned} \sum_{k=1}^K \xi_k^x &\geq \sum_{k=1}^K c_k(i) \cdot \xi_k^x \\ &= b(i) \\ &\geq b, \end{aligned}$$

and (iii) the inequality:

$$\begin{aligned} \sum_{k=1}^K \xi_k^x &\leq \sum_{k=1}^K \frac{\sum_{i=1}^C c_k(i)}{\epsilon} \cdot \xi_k^x \\ &= \frac{1}{\epsilon} \cdot \sum_{i=1}^C \sum_{k=1}^K c_k(i) \cdot \xi_k^x \\ &\leq \frac{\sum_{i=1}^C b(i)}{\epsilon}, \end{aligned}$$

for any basis $x \in \mathcal{B}$.

As a side note, observe a possibly better regret bound is given by:

$$R_{B(1), \dots, B(C)} \leq 16\rho \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln(T) + O(1),$$

if time is a limited resource since, in this case, $\tau^* \leq T$ and the constraint $\sum_{k=1}^K \xi_k^x \leq 1$ is part of (3).

D.3. Proof of Theorem 4. Along the same lines as for the case of a single limited resource, we start from inequality (30) derived in the proof of Theorem 3 and apply Lemma 8 only if Δ_x is big enough, taking into account the fact that:

$$\sum_{x \in \mathcal{B}} \mathbb{E}[n_{x, \tau^*}] \leq \mathbb{E}[\tau^*] \leq \frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1.$$

Specifically, we have:

$$\begin{aligned}
& R_{B(1), \dots, B(C)} \\
& \leq \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq \frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1}} \left\{ \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \min\left(\frac{\Delta_x}{\sum_{k=1}^K \xi_k^x} \cdot n_x, \right. \right. \\
& \quad \left. \left. 16\rho \cdot \frac{\sum_{k=1}^K \xi_k^x}{\Delta_x} \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1\right) + \frac{\pi^2}{3} \rho \cdot \frac{\Delta_x}{\sum_{k=1}^K \xi_k^x}\right)\right\} + O(1) \\
& \leq \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq \frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1}} \left\{ \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \min\left(\frac{\Delta_x}{\sum_{k=1}^K \xi_k^x} \cdot n_x, \right. \right. \\
& \quad \left. \left. 16\rho \cdot \frac{\sum_{k=1}^K \xi_k^x}{\Delta_x} \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1\right)\right)\right\} + \frac{\pi^2}{3} \frac{|\mathcal{B}| \cdot \rho}{\epsilon} + O(1) \\
& \leq \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq \frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1}} \left\{ \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \sqrt{16\rho \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1\right) \cdot n_x}\right\} + O(1) \\
& \leq 4\sqrt{\rho \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1\right) \cdot \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq \frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1}} \left\{ \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \sqrt{n_x}\right\}} + O(1) \\
& \leq 4\sqrt{\rho \cdot |\mathcal{B}| \cdot \left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1\right) \cdot \ln\left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1\right) + O(1)},
\end{aligned}$$

where we use $\Delta_x \leq \sum_{i=1}^C b(i)/\epsilon$ and $\sum_{k=1}^K \xi_k^x \geq b$ for the second inequality (see the end of the proof of Theorem 3), we maximize over $\Delta_x / \sum_{k=1}^K \xi_k^x \geq 0$ for each $x \in \mathcal{B}$ to derive the third inequality, and we use Cauchy-Schwartz for the last inequality.

Appendix E: Proofs for Section 6.

E.1. Proof of Lemma 9. For any $T \in \mathbb{N}$ and any arm $k \in \{1, \dots, K\}$, we denote by $n_{k,T}^{\text{opt}}$ the expected number of times that arm k is pulled by the optimal non-anticipating algorithm (which is characterized by a high-dimensional dynamic program) when the time horizon is T and the budget is $B = b \cdot T$. We prove the claim in two steps. First, we show that if $T - n_{k^*,T}^{\text{opt}} = \Omega(\sqrt{T})$ (Case A) or $T - n_{k^*,T}^{\text{opt}} = o(\sqrt{T})$ (Case B) then $\text{ER}_{\text{OPT}}(B, T) = T \cdot \text{obj}_{x^*} - \Omega(\sqrt{T})$. This is enough to establish the result because if $T - n_{k^*,T}^{\text{opt}} \neq \Omega(\sqrt{T})$ then we can extract a subsequence of $(T - n_{k^*,T}^{\text{opt}})/\sqrt{T}$ that converges to 0 and we can conclude with Case B.

Case A: $T - n_{k^*,T}^{\text{opt}} = \Omega(\sqrt{T})$. Consider the linear program:

$$\begin{aligned}
& \sup_{(\xi_k)_{k=1, \dots, K}} \sum_{k=1}^K \mu_k^T \cdot \xi_k \\
& \text{subject to} \quad \sum_{k=1}^K \mu_k^c \cdot \xi_k \leq b \\
& \quad \sum_{k=1}^K \xi_k \leq 1 \\
& \quad \xi_{k^*} \leq \Gamma \\
& \quad \xi_k \geq 0, \quad k = 1, \dots, K
\end{aligned} \tag{32}$$

parametrized by Γ and its dual:

$$\begin{aligned} & \inf_{(\zeta_1, \zeta_2, \zeta_3)} && b \cdot \zeta_1 + \zeta_2 + \Gamma \cdot \zeta_3 \\ \text{subject to} &&& \mu_k^c \cdot \zeta_1 + \zeta_2 \geq \mu_k^r, \quad k \neq k^* \\ &&& \mu_{k^*}^c \cdot \zeta_1 + \zeta_2 + \zeta_3 \geq \mu_{k^*}^r \\ &&& \zeta_1, \zeta_2, \zeta_3 \geq 0. \end{aligned} \quad (33)$$

Since the vector $(\xi_k)_{k=1, \dots, K}$ determined by $\xi_{k^*} = 1$ and $\xi_k = 0$ for $k \neq k^*$ is the only optimal solution to (32) when $\Gamma = 1$ (by assumption), we can find a strictly complementary optimal solution to the dual (33) $\zeta_1^*, \zeta_2^*, \zeta_3^* > 0$. Moreover, by definition of $n_{k^*, T}^{\text{opt}}$, we can show, along the same lines as in the proof of Lemma 2, that $\text{ER}_{\text{OPT}}(B, T)$ is no larger than T times the optimal value of (32) when $\Gamma = n_{k^*, T}^{\text{opt}}/T$ (up to a constant additive term of order $O(1)$). By weak duality, and since $(\zeta_1^*, \zeta_2^*, \zeta_3^*)$ is feasible for (33) for any Γ , this implies:

$$\begin{aligned} \text{ER}_{\text{OPT}}(B, T) &\leq T \cdot (b \cdot \zeta_1^* + \zeta_2^* + \frac{n_{k^*, T}^{\text{opt}}}{T} \cdot \zeta_3^*) + O(1) \\ &\leq T \cdot (b \cdot \zeta_1^* + \zeta_2^* + \zeta_3^* - \frac{T - n_{k^*, T}^{\text{opt}}}{T} \cdot \zeta_3^*) + O(1) \\ &\leq T \cdot \text{obj}_{x^*} - \Omega(\sqrt{T}), \end{aligned}$$

where we use the fact that $b \cdot \zeta_1^* + \zeta_2^* + \zeta_3^*$ is the optimal value of (32) when $\Gamma = 1$ by strong duality (both (32) and (33) are feasible) and $\zeta_3^* > 0$.

Case B: $T - n_{k^*, T}^{\text{opt}} = o(\sqrt{T})$. First observe that since the vector $(\xi_k)_{k=1, \dots, K}$ determined by $\xi_{k^*} = 1$ and $\xi_k = 0$ for $k \neq k^*$ is the only optimal solution to (3), it must be that $\mu_{k^*}^r > 0$ (since 0 is a feasible solution to (3) with objective value 0). For any $t \in \mathbb{N}$, denote by a_t the arm pulled by the optimal non-anticipating algorithm at time t and define τ_T^* as the corresponding stopping time when the time horizon is T . We have:

$$\begin{aligned} \text{ER}_{\text{OPT}}(B, T) &= \mathbb{E}\left[\sum_{t=1}^{\tau_T^*-1} r_{a_t, t}\right] \\ &= \sum_{k=1}^K \mu_k^r \cdot n_{k, T}^{\text{opt}} \\ &\leq \sum_{k \neq k^*} n_{k, T}^{\text{opt}} + \mu_{k^*}^r \cdot n_{k^*, T}^{\text{opt}} \\ &\leq (T - n_{k^*, T}^{\text{opt}}) + \mu_{k^*}^r \cdot (\mathbb{E}[\tau_T^*] - 1) \\ &= T \cdot \mu_{k^*}^r - \mu_{k^*}^r \cdot (T - \mathbb{E}[\tau_T^*] + 1) + o(\sqrt{T}) \\ &= T \cdot \text{obj}_{x^*} - \mu_{k^*}^r \cdot \mathbb{E}\left[\sum_{t=\tau_T^*}^T 1\right] + o(\sqrt{T}) \\ &\leq T \cdot \text{obj}_{x^*} - \mu_{k^*}^r \cdot \mathbb{E}\left[\left(\sum_{t=\tau_T^*}^T c_{k^*, t} + \sum_{t=1}^{\tau_T^*-1} c_{a_t, t} - B\right)_+\right] \\ &\leq T \cdot \text{obj}_{x^*} - \mu_{k^*}^r \cdot \left(\mathbb{E}\left[\left(\sum_{t=1}^T \{c_{k^*, t} - b\}\right)_+\right] - \sum_{k \neq k^*} n_{k, T}^{\text{opt}}\right) + o(\sqrt{T}) \\ &= T \cdot \text{obj}_{x^*} - \mu_{k^*}^r \cdot \mathbb{E}\left[\left(\sum_{t=1}^T \{c_{k^*, t} - b\}\right)_+\right] + o(\sqrt{T}). \end{aligned}$$

The first inequality is obtained using the fact that the rewards are bounded by 1. The second inequality is obtained using the fact that $\sum_{k=1}^K n_{k, T}^{\text{opt}} = \mathbb{E}[\tau_T^*] - 1 \leq T$. The third inequality is obtained along the

same lines as in the proof of Lemma 3, using $\sum_{t=1}^{\tau_T^*-1} c_{a_t,t} \leq B$ by definition of τ_T^* . We use the inequality $(y+z)_+ \geq y_+ - |z|$ (true for any $(y,z) \in \mathbb{R}^2$) and the fact the amount of resource consumed at any step is no larger than 1 for the fourth inequality. Since $(c_{k^*,t} - b)_{t \in \mathbb{N}}$ is an i.i.d. zero-mean bounded stochastic process with positive variance, $\frac{1}{\sqrt{T}} \cdot \mathbb{E}[(\sum_{t=1}^T \{c_{k^*,t} - b\})_+]$ converges to a positive value and we conclude:

$$\text{ER}_{\text{OPT}}(B, T) \leq T \cdot \text{obj}_{x^*} - \Omega(\sqrt{T}),$$

since $\mu_{k^*}^r > 0$.

E.2. Proof of Lemma 10. Consider x either an infeasible basis to (3) or a pseudo-basis for (3). Without loss of generality, we can assume that x involves two arms (one of which may be a dummy arm introduced in the specification of the algorithm given in Section 6) and that $\mathcal{K}_x = \{k, l\}$ with $\mu_k^c, \mu_l^c > b$ (the situation is symmetric if the reverse inequality holds). Defining $\beta_x = 32/\epsilon^3$, we have:

$$\begin{aligned} \mathbb{E}[n_{x,T}] &\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\ &\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{b_{x,t} \geq n_{x,t} \cdot (b+\epsilon/2)}\right] \\ &\quad + \sum_{t=1}^T \sum_{s=\beta_x \ln(t)}^t \mathbb{P}[b_{x,t} < s \cdot (b+\epsilon) - s \cdot \epsilon/2, n_{x,t} = s] \\ &\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{b_{x,t} \geq n_{x,t} \cdot (b+\epsilon/2)}\right] + \sum_{t=1}^T \sum_{s=\beta_x \ln(t)}^{\infty} \exp(-s \frac{\epsilon^2}{2}) \\ &\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{b_{x,t} \geq n_{x,t} \cdot b + \epsilon \beta_x / 4 \cdot \ln(t)}\right] + \frac{2\pi^2}{3\epsilon^2}. \end{aligned}$$

The first inequality is derived along the same lines as in Fact 1. The third inequality is obtained by observing that, as a result of Assumption 6, the average amount of resource consumed any time basis x is selected at Step-Simplex is at least $b + \epsilon$ no matter which of arm k or l is pulled. Finally, we use the same bounds as in Fact 2 for the last two inequalities. Observe that if x is selected at time t , either $\bar{c}_{l,t} - \epsilon_{l,t} \leq b$ or $\bar{c}_{l,t} - \epsilon_{l,t} \leq b$, otherwise x would have been infeasible for (8). Moreover, if $n_{x,t} \geq \beta_x \ln(T)$, then we have either $n_{k,t}^x \geq \beta_x/2 \ln(T)$ or $n_{l,t}^x \geq \beta_x/2 \ln(T)$ since there are only two arms in \mathcal{K}_x . By symmetry, we study the first situation and look at:

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{k,t}^x \geq \beta_x/2 \cdot \ln(t)} \cdot I_{b_{x,t} \geq n_{x,t} \cdot b + \epsilon \beta_x / 4 \cdot \ln(t)}\right] \\ &\leq \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{k,t}^x \geq \beta_x/2 \cdot \ln(t)} \cdot I_{b_{x,t} \geq n_{x,t} \cdot b + \epsilon \beta_x / 4 \cdot \ln(t)} \cdot I_{\bar{c}_{k,\tau_q} - \epsilon_{k,\tau_q} \geq b, q=\tau_t - \epsilon \beta_x / 4 \ln(t), \dots, \tau_t}\right] \\ &\quad + \sum_{t=1}^T \sum_{\tau=1}^t \sum_{s=\beta_x/4 \cdot \ln(t)}^t \mathbb{P}[\bar{c}_{k,\tau} < b + \frac{\epsilon}{2}, n_{k,\tau} = s] \\ &\leq \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{k,t}^x \geq \beta_x/2 \cdot \ln(t)} \cdot I_{b_{x,t} \geq n_{x,t} \cdot b + \epsilon \beta_x / 4 \cdot \ln(t)} \cdot I_{\bar{c}_{k,\tau_q} - \epsilon_{k,\tau_q} \geq b, q=\tau_t - \epsilon \beta_x / 4 \ln(t), \dots, \tau_t}\right] \\ &\quad + \sum_{t=1}^T \sum_{\tau=1}^t \sum_{s=\beta_x/4 \cdot \ln(t)}^{\infty} \exp(-s \cdot \frac{\epsilon^2}{2}) \end{aligned}$$

$$\begin{aligned} &\leq \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{k,t}^x \geq \beta_x/2 \cdot \ln(t)} \cdot I_{b_{x,t} \geq n_{x,t} \cdot b + \epsilon \beta_x/4 \cdot \ln(t)} \cdot I_{\bar{c}_{k,\tau_q} - \epsilon_{k,\tau_q} \geq b, q=q_t - \epsilon \beta_x/4 \ln(t), \dots, q_t}\right] \\ &\quad + \frac{2\pi^2}{3\epsilon^2}, \end{aligned}$$

where $(\tau_q)_{q \in \mathbb{N}}$ denote the random times at which basis x is selected and, for a time t at which basis x is selected, q_t denotes the index $q \in \mathbb{N}$ such that $\tau_q = t$. The first inequality is a consequence of $n_{k,\tau_q}^x = n_{k,t}^x - (q_t - q) \geq n_{k,t}^x - \epsilon \beta_x/4 \ln(t) \geq \beta_x/4 \ln(t)$ for $q = q_t - \epsilon \beta_x/4 \ln(t), \dots, q_t$ and $n_{k,\tau_q} \geq n_{k,\tau_q}^x$, which implies $\epsilon_{k,\tau_q} \leq \epsilon/2$. We use the same bounds as in Fact 2 for the last two inequalities. Now observe that, for any $q \in \{q_t - \epsilon \beta_x/4 \ln(t), \dots, q_t\}$, we have $\bar{c}_{l,\tau_q} - \epsilon_{l,\tau_q} \leq b$ since $\bar{c}_{k,\tau_q} - \epsilon_{k,\tau_q} \geq b$ and since x is feasible basis to (8) at time τ_q (by definition). This implies that, for any $q \in \{q_t - \epsilon \beta_x/4 \ln(t), \dots, q_t\}$, arm l was pulled at time τ_q by definition of the load balancing algorithm since the amount of resource consumed at any round cannot be larger than 1 and $b_{x,\tau_q} \geq b_{x,t} - (q_t - q) \geq n_{x,t} \cdot b + \epsilon \beta_x/4 \cdot \ln(t) - (q_t - q) \geq n_{x,t} \cdot b \geq n_{x,\tau_q} b$. Hence, we get:

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{k,t}^x \geq \beta_x/2 \cdot \ln(t)} \cdot I_{b_{x,t} \geq n_{x,t} \cdot b + \epsilon \beta_x/4 \cdot \ln(t)}\right] \\ &\leq \mathbb{E}\left[\sum_{t=1}^T I_{n_{l,t}^x \geq \epsilon \beta_x/4 \cdot \ln(t)} \cdot I_{\bar{c}_{l,t} - \epsilon_{l,t} \leq b}\right] + \frac{2\pi^2}{3\epsilon^2} \\ &\leq \sum_{t=1}^T \mathbb{P}[\bar{c}_{l,t} \leq b + \frac{\epsilon}{2}, n_{l,t}^x \geq \epsilon \beta_x/4 \cdot \ln(t)] + \frac{2\pi^2}{3\epsilon^2} \\ &\leq \sum_{t=1}^T \sum_{s=\epsilon \beta_x/4 \cdot \ln(t)}^t \mathbb{P}[\bar{c}_{l,t} \leq b + \frac{\epsilon}{2}, n_{l,t} = s] + \frac{2\pi^2}{3\epsilon^2} \\ &\leq \sum_{t=1}^T \sum_{s=\epsilon \beta_x/4 \cdot \ln(t)}^{\infty} \exp(-s \cdot \frac{\epsilon^2}{2}) + \frac{2\pi^2}{3\epsilon^2} \\ &\leq \frac{4\pi^2}{3\epsilon^2}. \end{aligned}$$

Bringing everything together, we derive:

$$\mathbb{E}[n_{x,T}] \leq \frac{2^6}{\epsilon^3} \cdot \ln(T) + \frac{10\pi^2}{3\epsilon^2}.$$

E.3. Proof of Lemma 11. Without loss of generality, we can assume that x involves two arms (one of which may be a dummy arm introduced in the specification of the algorithm given in Section 6) and that $\mathcal{K}_x = \{k, l\}$ with $\mu_k^c > b > \mu_l^c$. We say that a “swap” occurred at time τ if basis x was selected at time τ and $\bar{c}_{k,\tau} - \epsilon_{k,\tau} \leq b \leq \bar{c}_{l,\tau} - \epsilon_{l,\tau}$. We define $n_{x,t}^{\text{swap}}$ as the total number of swaps that have occurred before time t , i.e. $n_{x,t}^{\text{swap}} = \sum_{\tau=1}^{t-1} I_{x_\tau=x} \cdot I_{\bar{c}_{k,\tau} - \epsilon_{k,\tau} \leq b \leq \bar{c}_{l,\tau} - \epsilon_{l,\tau}}$. Consider $u \geq 1$ and define $\gamma_x = (4/\epsilon)^2$. First note that:

$$\begin{aligned} \mathbb{P}[n_{x,t}^{\text{swap}} \geq \gamma_x \ln(t)] &\leq \sum_{q=\gamma_x \ln(t)}^t \mathbb{P}[\bar{c}_{k,\tau_q} - \epsilon_{k,\tau_q} \leq b \leq \bar{c}_{l,\tau_q} - \epsilon_{l,\tau_q}] \\ &\leq \sum_{q=\gamma_x \ln(t)}^t \mathbb{P}[\bar{c}_{k,\tau_q} \leq b + \frac{\epsilon}{2}, n_{k,\tau_q} \geq \frac{\gamma_x}{2} \ln(t)] \\ &\quad + \sum_{q=\gamma_x \ln(t)}^t \mathbb{P}[b - \frac{\epsilon}{2} \leq \bar{c}_{l,\tau_q}, n_{l,\tau_q} \geq \frac{\gamma_x}{2} \ln(t)] \end{aligned}$$

$$\begin{aligned} &\leq 2 \sum_{q=1}^t \sum_{s=\gamma_x/2 \cdot \ln(t)}^{\infty} \exp(-s \cdot \frac{\epsilon^2}{2}) \\ &\leq \frac{8}{\epsilon^2 \cdot t^2}, \end{aligned}$$

where $(\tau_q)_{q \in \mathbb{N}}$ are defined as the times at which basis x is selected. The first inequality is derived observing that if $n_{x,t}^{\text{swap}} \geq \gamma_x \ln(t)$ then it must be that basis x was selected for the q th time, for some $q \geq \gamma_x \ln(t)$, and that we had $\bar{c}_{k,\tau_q} - \epsilon_{k,\tau_q} \leq b \leq \bar{c}_{l,\tau_q} - \epsilon_{l,\tau_q}$. To obtain the second inequality, we observe that, at any time τ , at least one of arm k and l must have been pulled $n_{x,\tau}/2$ times and that $\epsilon_{k,\tau} \leq \epsilon/2$ when $n_{k,\tau} \geq \gamma_x/2 \ln(t)$ (a similar inequality holds for arm l). The last two inequalities are derived in the same fashion as in Lemma 10. This yields:

$$\begin{aligned} &\mathbb{P}[|b_{x,t} - n_{x,t} \cdot b| \geq u + \gamma_x \ln(t)] \\ &\leq \mathbb{P}[|b_{x,t} - n_{x,t} \cdot b| \geq u + \gamma_x \ln(t); n_{x,t}^{\text{swap}} \leq \gamma_x \ln(t)] + \mathbb{P}[n_{x,t}^{\text{swap}} \geq \gamma_x \ln(t)] \\ &\leq \mathbb{P}[|b_{x,t} - n_{x,t} \cdot b| \geq u + \gamma_x \ln(t); n_{x,t}^{\text{swap}} \leq \gamma_x \ln(t)] + \frac{8}{\epsilon^2 \cdot t^2}. \end{aligned}$$

Note that, by definition of the load balancing algorithm, we are led to pull arm k (resp. arm l) at time τ_q if the budget spent so far when selecting basis x , denoted by b_{x,τ_q} , is below (resp. above) the ‘‘target’’ of $n_{x,\tau_q} \cdot b$ assuming there is no ‘‘swap’’ at time τ_q (i.e. $\bar{c}_{k,\tau_q} - \epsilon_{k,\tau_q} \geq \bar{c}_{l,\tau_q} - \epsilon_{l,\tau_q}$). Hence, if $b_{x,t} - n_{x,t} \cdot b \geq u + \gamma_x \ln(t)$ and $n_{x,t}^{\text{swap}} \leq \gamma_x \ln(t)$, we must have been pulling arm l for at least $s \geq \lfloor u \rfloor$ rounds $t_1 \leq \dots \leq t_s \leq t-1$ where basis x was selected since the last time, denoted by t_0 , where basis x was selected and the budget was below the target, i.e. $b_{x,t_0} \leq n_{x,t_0} \cdot b$ (because the amounts of resource consumed at each round are almost surely bounded by 1). Moreover, we have:

$$\begin{aligned} &\sum_{\tau=t_0+1}^{t-1} I_{x_\tau=x} \cdot (c_{k,\tau} \cdot I_{\bar{c}_{k,\tau} - \epsilon_{k,\tau} \geq \bar{c}_{l,\tau} - \epsilon_{l,\tau}} + c_{l,\tau} \cdot I_{\bar{c}_{k,\tau} - \epsilon_{k,\tau} < \bar{c}_{l,\tau} - \epsilon_{l,\tau}}) \\ &= b_{x,t} - b_{x,t_0+1} \\ &\geq (n_{x,t} - n_{x,t_0}) \cdot b + u - 1 + \gamma_x \ln(t) \\ &\geq s \cdot b + u - 1 + \gamma_x \ln(t). \end{aligned}$$

This implies:

$$\sum_{q=1}^s c_{l,t_q} \geq s \cdot b + u - 1$$

since $\sum_{\tau=t_0+1}^{t-1} I_{x_\tau=x} \cdot I_{\bar{c}_{k,\tau} - \epsilon_{k,\tau} < b \cdot c_{k,\tau}} \leq n_{x,t}^{\text{swap}} \leq \gamma_x \ln(t)$. Hence, if $u \geq 1$:

$$\begin{aligned} &\mathbb{P}[b_{x,t} - n_{x,t} \cdot b \geq u + \gamma_x \ln(t); n_{x,t}^{\text{swap}} \leq \gamma_x \ln(t)] \\ &\leq \sum_{s=\lfloor u \rfloor}^t \mathbb{P}[\sum_{q=1}^s c_{l,t_q} \geq s \cdot b + u - 1] \\ &= \sum_{s=\lfloor u \rfloor}^t \mathbb{P}[\sum_{q=1}^s c_{l,t_q} \geq s \cdot \mu_l^c + s \cdot (b - \mu_l^c)] \\ &\leq \sum_{s=\lfloor u \rfloor}^t \mathbb{P}[\sum_{q=1}^s c_{l,t_q} \geq s \cdot \mu_l^c + s \cdot \epsilon] \\ &\leq \sum_{s=\lfloor u \rfloor}^t \exp(-2\epsilon^2 \cdot s) \\ &\leq \frac{\exp(-2\epsilon^2 \cdot \lfloor u \rfloor)}{1 - \exp(-2\epsilon^2)} \\ &\leq \frac{2}{\epsilon^2} \cdot \exp(-\epsilon^2 \cdot u), \end{aligned}$$

where we use Lemma 1 for the third inequality and the fact that $\exp(-2v) \leq 1 - v/2$ for $v \in [0, 1]$ for the last inequality. With a similar argument, we conclude:

$$\mathbb{P}[|b_{x,t} - n_{x,t}| \cdot b \geq u + \gamma_x \ln(t); n_{x,t}^{\text{swap}} \leq \gamma_x \ln(t)] \leq \frac{4}{\epsilon^2} \cdot \exp(-\epsilon^2 \cdot u).$$

This last result enables us to show that:

$$\begin{aligned} & |\mathbb{E}[b_{x,T}] - \mathbb{E}[n_{x,T}] \cdot b| \\ & \leq \mathbb{E}[|b_{x,T} - n_{x,T} \cdot b|] \\ & = \int_0^T \mathbb{P}[|b_{x,T} - n_{x,T} \cdot b| \geq u] du \\ & \leq \int_0^T \mathbb{P}[|b_{x,T} - n_{x,T} \cdot b| \geq u; n_{x,T}^{\text{swap}} \leq \gamma_x \ln(T)] du + T \cdot \mathbb{P}[n_{x,T}^{\text{swap}} \geq \gamma_x \ln(T)] \\ & \leq \int_0^T \mathbb{P}[|b_{x,T} - n_{x,T} \cdot b| \geq u + 1 + \gamma_x \ln(T); n_{x,T}^{\text{swap}} \leq \gamma_x \ln(T)] du + 1 + \gamma_x \ln(T) + \frac{8}{\epsilon^2} \\ & \leq \frac{4}{\epsilon^2} \cdot \int_0^T \exp(-\epsilon^2 \cdot u) du + 1 + \gamma_x \ln(T) + \frac{8}{\epsilon^2} \\ & = \frac{13}{\epsilon^4} + \left(\frac{4}{\epsilon}\right)^2 \ln(T). \end{aligned}$$

We get:

$$\mathbb{E}[n_{k,T}^x] \cdot \mu_k^c + \mathbb{E}[n_{l,T}^x] \cdot \mu_l^c = \mathbb{E}[b_{x,T}] \geq \mathbb{E}[n_{x,T}] \cdot b - \frac{13}{\epsilon^4} - \left(\frac{4}{\epsilon}\right)^2 \ln(T),$$

which, in combination with $\mathbb{E}[n_{k,T}^x] + \mathbb{E}[n_{l,T}^x] = \mathbb{E}[n_{x,T}]$, shows that:

$$\begin{aligned} \mathbb{E}[n_{k,T}^x] & \geq \left(\frac{b - \mu_l^c}{\mu_k^c - \mu_l^c}\right) \cdot \mathbb{E}[n_{x,T}] - \frac{13}{\epsilon^4 \cdot (\mu_k^c - \mu_l^c)} - \frac{4^2}{\epsilon^2 \cdot (\mu_k^c - \mu_l^c)} \ln(T) \\ & \geq \xi_k^x \cdot \mathbb{E}[n_{x,T}] - \frac{13}{\epsilon^5} - \frac{16}{\epsilon^3} \ln(T). \end{aligned}$$

Symmetrically, we get:

$$\mathbb{E}[n_{l,T}^x] \geq \xi_l^x \cdot \mathbb{E}[n_{x,T}] - \frac{13}{\epsilon^5} - \frac{16}{\epsilon^3} \ln(T).$$

E.4. Proof of Lemma 12. Consider any suboptimal basis $x \in \mathcal{B}$. We use the shorthand notation $\beta_x = \frac{2^8}{\epsilon^3} \cdot \left(\frac{\lambda}{\Delta_x}\right)^2$. Without loss of generality, we can assume that both x and x^* involve two arms (one of which may be a dummy arm introduced in the specification of the algorithm given in Section 6) and that $\mathcal{K}_{x^*} = \{k^*, l^*\}$ with $\mu_{k^*}^c > b > \mu_{l^*}^c$ and $\mathcal{K}_x = \{k, l\}$ with $\mu_k^c > b > \mu_l^c$. The proof is along the same lines as for Lemmas 5 and 8. We break down the analysis in a series of facts where we emphasize the main differences. We start off with an inequality analogous to Fact 1.

$$\begin{aligned} \mathbb{E}[n_{x,T}] & \leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\ & \leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{\bar{c}_{k^*,t} - \epsilon_{k^*,t} \leq \mu_{k^*}^c} \cdot I_{\bar{c}_{l^*,t} - \epsilon_{l^*,t} \leq \mu_{l^*}^c}\right] \\ & \quad + \sum_{t=1}^T \mathbb{P}[\bar{c}_{l^*,t} > \mu_{l^*}^c + \epsilon_{l^*,t}] + \mathbb{P}[\bar{c}_{k^*,t} > \mu_{k^*}^c + \epsilon_{k^*,t}] \\ & \leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{\bar{c}_{k^*,t} - \epsilon_{k^*,t} \leq \mu_{k^*}^c} \cdot I_{\bar{c}_{l^*,t} - \epsilon_{l^*,t} \leq \mu_{l^*}^c}\right] + \frac{\pi^2}{3} \end{aligned}$$

$$\begin{aligned}
&\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{\text{obj}_{x,t}+E_{x,t} \geq \sum_{k \in \{k^*, l^*\}} (\bar{r}_{k,t} + \lambda \epsilon_{k,t}) \cdot \xi_k^{x^*}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] + \frac{\pi^2}{3} \\
&\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{\text{obj}_{x,t}+E_{x,t} \geq \text{obj}_{x^*}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\
&\quad + \sum_{t=1}^T \mathbb{P}[\bar{r}_{l^*,t} < \mu_{l^*}^r - \epsilon_{l^*,t}] + \mathbb{P}[\bar{r}_{k^*,t} < \mu_{k^*}^r - \epsilon_{k^*,t}] + \frac{\pi^2}{3} \\
&\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{\text{obj}_{x,t}+E_{x,t} \geq \text{obj}_{x^*}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] + \frac{2\pi^2}{3}.
\end{aligned}$$

The first inequality is derived in the same fashion as in Fact 1 substituting k with x . The third and last inequalities are obtained using Lemma 1 in the same fashion as in Fact 3. The fourth inequality is obtained by observing that (i) if $x_t = x$ then x_t is optimal for (8) and (ii) $(\xi_k^*)_{k=1, \dots, K}$ is feasible for (8) if $\bar{c}_{l^*,t} - \epsilon_{l^*,t} \leq \mu_{l^*}^c$ and $\bar{c}_{k^*,t} - \epsilon_{k^*,t} \leq \mu_{k^*}^c$. The fifth inequality results from $\lambda \geq 1$ and $\text{obj}_{x^*} = \sum_{k \in \{k^*, l^*\}} \mu_k^r \cdot \xi_k^{x^*}$. The second term in the last upper bound can be broken down in two terms similarly as in Lemmas 5 and 8:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^T I_{x_t=x} \cdot I_{\text{obj}_{x,t}+E_{x,t} \geq \text{obj}_{x^*}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^T I_{x_t \in \mathcal{B}_t} \cdot I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \tag{34}
\end{aligned}$$

$$+ \mathbb{E}\left[\sum_{t=1}^T I_{x_t \in \mathcal{B}_t} \cdot I_{\text{obj}_{x^*} \leq \text{obj}_x + 2E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right]. \tag{35}$$

We carefully study each term separately.

Fact 9

$$\mathbb{E}\left[\sum_{t=1}^T I_{x \in \mathcal{B}_t} \cdot I_{\text{obj}_{x^*} \leq \text{obj}_x + 2E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \leq \frac{6\pi^2}{\epsilon^2}.$$

Proof. Using the shorthand notations $\alpha_x = 8\left(\frac{\lambda}{\Delta_x}\right)^2$ and $\gamma_x = \left(\frac{4}{\epsilon}\right)^2$, we have:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^T I_{x \in \mathcal{B}_t} \cdot I_{\text{obj}_{x^*} \leq \text{obj}_x + 2E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^T I_{\Delta_x \leq 2\lambda \cdot \max(\epsilon_{k,t}, \epsilon_{l,t})} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\
&= \mathbb{E}\left[\sum_{t=1}^T I_{\min(n_{k,t}, n_{l,t}) \leq \alpha_x \ln(t)} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\
&\leq \sum_{t=1}^T \mathbb{P}[n_{l,t} \leq \alpha_x \ln(t); n_{x,t} \geq \beta_x \ln(t)] + \sum_{t=1}^T \mathbb{P}[n_{k,t} \leq \alpha_x \ln(t); n_{x,t} \geq \beta_x \ln(t)].
\end{aligned}$$

The first inequality is derived using $E_{x,t} = \lambda \cdot (\xi_{k,t}^x \cdot \epsilon_{k,t} + \xi_{l,t}^x \cdot \epsilon_{l,t})$ and $\xi_{k,t}^x + \xi_{l,t}^x \leq 1$ (this is imposed as a constraint in (8)). Observe now that α_x / β_x is a constant factor independent of Δ_x . Thus, we just have to show that if x has been selected at least $\beta_x \ln(t)$ times, then both k and l have been pulled at least a

constant fraction of the time with high probability. This is the only time the load balancing algorithm comes into play in the proof of Lemma 12. We study the first term and we conclude the study by symmetry. We have:

$$\begin{aligned}
& \mathbb{P}[n_{l,t} \leq \alpha_x \ln(t) ; n_{x,t} \geq \beta_x \ln(t)] \\
& \leq \mathbb{P}[n_{l,t} \leq \alpha_x \ln(t) ; n_{x,t} \geq \beta_x \ln(t) ; \sum_{\tau=1}^t I_{x_\tau=x} \cdot I_{a_\tau=k} \cdot c_{k,\tau} \geq (b + \epsilon/2) \cdot n_{k,t}^x] \\
& + \sum_{s=4/\epsilon^2 \ln(t)}^t \mathbb{P}[\sum_{\tau=1}^t I_{x_\tau=x} \cdot I_{a_\tau=k} \cdot c_{k,\tau} \leq (b + \epsilon/2) \cdot s ; n_{k,t}^x = s] \\
& \leq \mathbb{P}[n_{l,t} \leq \alpha_x \ln(t) ; n_{x,t} \geq \beta_x \ln(t) ; b_{x,t} - n_{x,t} \cdot b \geq \epsilon/2 \cdot n_{k,t}^x - n_{l,t}^x] \\
& + \sum_{s=4/\epsilon^2 \ln(t)}^t \exp(-s \cdot \frac{\epsilon^2}{2}) \\
& \leq \mathbb{P}[b_{x,t} - n_{x,t} \cdot b \geq 2\gamma_x \ln(t)] + (\frac{2}{\epsilon \cdot t})^2 \\
& \leq \frac{16}{\epsilon^2 \cdot t^2}.
\end{aligned}$$

The first inequality is obtained observing that if $n_{l,t} \leq \alpha_x \ln(t)$ and $n_{x,t} \geq \beta_x \ln(t)$, we have:

$$n_{k,t}^x = n_{x,t} - n_{l,t}^x \geq (\frac{2^8}{\epsilon^3} - 8) \cdot (\frac{\lambda}{\Delta_x})^2 \cdot \ln(t) \geq \frac{4}{\epsilon^2} \cdot \ln(t)$$

because $\lambda \geq 1$ and $\Delta_x \leq \text{obj}_{x^*} = \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \leq \sum_{k=1}^K \xi_k^{x^*} \leq 1$ since x^* is a feasible basis to (3). To derive the second inequality, we use Lemma 1 for the second term and remark that:

$$b_{x,t} - n_{x,t} \cdot b \geq (\sum_{\tau=1}^t I_{x_\tau=x} \cdot I_{a_\tau=k} \cdot c_{k,\tau} - n_{k,t}^x \cdot b) - n_{l,t}^x$$

since $b \leq 1$. The third inequality is derived using:

$$\epsilon/2 \cdot n_{k,t}^x - n_{l,t}^x \geq \epsilon/2 \cdot n_{x,t} - 2 \cdot n_{l,t}^x \geq (\frac{2^7}{\epsilon^2} - 16) \cdot \ln(t) \geq 2\gamma_x \ln(t)$$

and the last inequality is obtained with Lemma 11.

Fact 10

$$\mathbb{E}[\sum_{t=1}^T I_{x_t \in \mathcal{B}_t} \cdot I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}] \leq \frac{3\pi^2}{\epsilon^2}.$$

Proof. First observe that:

$$\begin{aligned}
& \mathbb{E}[\sum_{t=1}^T I_{x_t \in \mathcal{B}_t} \cdot I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}] \\
& \leq \mathbb{E}[\sum_{t=1}^T I_{x_t \in \mathcal{B}_t} \cdot I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{n_{k,t} \geq \beta_x/2 \cdot \ln(t)}] \\
& + \mathbb{E}[\sum_{t=1}^T I_{x_t \in \mathcal{B}_t} \cdot I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{n_{l,t} \geq \beta_x/2 \cdot \ln(t)}] \\
& \leq \mathbb{E}[\sum_{t=1}^T I_{x_t \in \mathcal{B}_t} \cdot I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{\bar{c}_{k,t} - \epsilon_{k,t} > b}]
\end{aligned}$$

$$\begin{aligned}
& + \mathbb{E}\left[\sum_{t=1}^T I_{x_t \in \mathcal{B}_t} \cdot I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{\bar{c}_{l,t} - \epsilon_{l,t} < b}\right] \\
& + \sum_{t=1}^T \sum_{s=\beta_x/2 \cdot \ln(t)}^t \mathbb{P}[\bar{c}_{k,t} - \epsilon_{k,t} \leq b, n_{k,t} = s] + \mathbb{P}[\bar{c}_{l,t} - \epsilon_{l,t} \geq b, n_{l,t} = s] \\
& \leq 2 \cdot \mathbb{E}\left[\sum_{t=1}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{\bar{c}_{k,t} - \epsilon_{k,t} \geq b \geq \bar{c}_{l,t} - \epsilon_{l,t}} \cdot I_{\bar{c}_{k,t} - \epsilon_{k,t} > \bar{c}_{l,t} - \epsilon_{l,t}}\right] \\
& + \sum_{t=1}^T \sum_{s=\beta_x/2 \cdot \ln(t)}^t \mathbb{P}[\bar{c}_{k,t} \leq b + \epsilon/2, n_{k,t} = s] + \mathbb{P}[\bar{c}_{l,t} \geq b + \epsilon/2, n_{l,t} = s] \\
& \leq 2 \cdot \mathbb{E}\left[\sum_{t=1}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{\bar{c}_{k,t} - \epsilon_{k,t} \geq b \geq \bar{c}_{l,t} - \epsilon_{l,t}} \cdot I_{\bar{c}_{k,t} - \epsilon_{k,t} > \bar{c}_{l,t} - \epsilon_{l,t}}\right] \\
& + 2 \cdot \sum_{t=1}^T \sum_{s=\beta_x/2 \cdot \ln(t)}^t \exp(-s \frac{\epsilon^2}{2}) \\
& \leq 2 \cdot \mathbb{E}\left[\sum_{t=1}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{\bar{c}_{k,t} - \epsilon_{k,t} \geq b \geq \bar{c}_{l,t} - \epsilon_{l,t}} \cdot I_{\bar{c}_{k,t} - \epsilon_{k,t} > \bar{c}_{l,t} - \epsilon_{l,t}}\right] + \frac{4\pi^2}{3\epsilon^2}.
\end{aligned}$$

The third inequality is obtained by observing that $\epsilon_{k,t}, \epsilon_{l,t} \leq \frac{\epsilon}{2}$ for $n_{k,t}, n_{l,t} \geq \frac{\beta_x}{2} \ln(t)$ (because $\lambda \geq 1$ and $\Delta_x \leq \text{obj}_{x^*} = \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \leq \sum_{k=1}^K \xi_k^{x^*} \leq 1$) and that, if $x_t \in \mathcal{B}_t$ and (for example) $\bar{c}_{l,t} - \epsilon_{l,t} < b$, it must be that $\bar{c}_{k,t} - \epsilon_{k,t} \geq b$. The last two inequalities are obtained in the same fashion as in Lemma 10 observing that $\beta_x/2 \geq 4/\epsilon^2$. At this point, the key observation is that if $\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}$, $\bar{c}_{k,t} - \epsilon_{k,t} \geq b \geq \bar{c}_{l,t} - \epsilon_{l,t}$, and $\bar{c}_{k,t} - \epsilon_{k,t} > \bar{c}_{l,t} - \epsilon_{l,t}$, at least one of the following six events occurs: $\{\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}\}$, $\{\bar{r}_{l,t} \geq \mu_l^r + \epsilon_{l,t}\}$, $\{\bar{c}_{k,t} \leq \mu_k^c - \epsilon_{k,t}\}$, $\{\bar{c}_{k,t} \geq \mu_k^c + \epsilon_{k,t}\}$, $\{\bar{c}_{l,t} \leq \mu_l^c - \epsilon_{l,t}\}$ or $\{\bar{c}_{l,t} \geq \mu_l^c + \epsilon_{l,t}\}$. Otherwise, using the shorthand notations $\tilde{c}_k = \bar{c}_{k,t} - \epsilon_{k,t}$ and $\tilde{c}_l = \bar{c}_{l,t} - \epsilon_{l,t}$, we have:

$$\begin{aligned}
& \text{obj}_{x,t} - \text{obj}_x \\
& = \left[\frac{\tilde{c}_k - b}{\tilde{c}_k - \tilde{c}_l} \cdot \bar{r}_{l,t} + \frac{b - \tilde{c}_l}{\tilde{c}_k - \tilde{c}_l} \cdot \bar{r}_{k,t} \right] - \left[\frac{\mu_k^c - b}{\mu_k^c - \mu_l^c} \cdot \mu_l^r + \frac{b - \mu_l^c}{\mu_k^c - \mu_l^c} \cdot \mu_k^r \right] \\
& < \left[\frac{\tilde{c}_k - b}{\tilde{c}_k - \tilde{c}_l} \cdot (\mu_l^r + \epsilon_{l,t}) + \frac{b - \tilde{c}_l}{\tilde{c}_k - \tilde{c}_l} \cdot (\mu_k^r + \epsilon_{k,t}) \right] - \left[\frac{\mu_k^c - b}{\mu_k^c - \mu_l^c} \cdot \mu_l^r + \frac{b - \mu_l^c}{\mu_k^c - \mu_l^c} \cdot \mu_k^r \right] \\
& = \frac{1}{\lambda} \cdot E_{x,t} + (\mu_k^r - \mu_l^r) \cdot \left[\frac{b - \tilde{c}_l}{\tilde{c}_k - \tilde{c}_l} - \frac{b - \mu_l^c}{\mu_k^c - \mu_l^c} \right] \\
& = \frac{1}{\lambda} \cdot E_{x,t} + \frac{(\mu_k^r - \mu_l^r)}{(\tilde{c}_k - \tilde{c}_l) \cdot (\mu_k^c - \mu_l^c)} \cdot [(\mu_k^c - b)(\mu_l^c - \tilde{c}_l) + (b - \mu_l^c)(\mu_k^c - \tilde{c}_k)] \\
& \leq \frac{1}{\lambda} \cdot E_{x,t} + \frac{\kappa}{\tilde{c}_k - \tilde{c}_l} \cdot |(\mu_k^c - b)(\mu_l^c - \tilde{c}_l) + (b - \mu_l^c)(\mu_k^c - \tilde{c}_k)| \\
& \leq \frac{1}{\lambda} \cdot E_{x,t} + 2 \frac{\kappa}{\tilde{c}_k - \tilde{c}_l} \cdot [(\tilde{c}_k - b) \cdot \epsilon_{l,t} + (b - \tilde{c}_l) \cdot \epsilon_{k,t}] \\
& = \frac{1}{\lambda} \cdot E_{x,t} + \frac{2\kappa}{\lambda} \cdot E_{x,t} \\
& = E_{x,t},
\end{aligned}$$

a contradiction. The first inequality is strict because either $\tilde{c}_k > b$ or $\tilde{c}_l < b$. The second inequality is derived using Assumption 5. The third inequality is derived from the observation that the expression $(\mu_k^c - b)(\mu_l^c - \tilde{c}_l) + (b - \mu_l^c)(\mu_k^c - \tilde{c}_k)$ is a linear function of (μ_k^c, μ_l^c) (since the cross term $\mu_k^c \cdot \mu_l^c$ cancels out) so that $|(\mu_k^c - b)(\mu_l^c - \tilde{c}_l) + (b - \mu_l^c)(\mu_k^c - \tilde{c}_k)|$ is convex in (μ_k^c, μ_l^c) and the maximum of this expression over the polyhedron $[\bar{c}_{k,t} - \epsilon_{k,t}, \bar{c}_{k,t} + \epsilon_{k,t}] \times [\bar{c}_{l,t} - \epsilon_{l,t}, \bar{c}_{l,t} + \epsilon_{l,t}]$ is attained at an extreme point. We obtain:

$$\mathbb{E}\left[\sum_{t=1}^T I_{x \in \mathcal{B}_t} I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right]$$

$$\begin{aligned}
&\leq \sum_{t=1}^{\infty} \mathbb{P}[\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}] + \mathbb{P}[\bar{r}_{l,t} \geq \mu_l^r + \epsilon_{l,t}] \\
&+ \sum_{t=1}^{\infty} \mathbb{P}[\bar{c}_{l,t} \geq \mu_l^c + \epsilon_{l,t}] + \mathbb{P}[\bar{c}_{k,t} \geq \mu_k^c + \epsilon_{k,t}] \\
&+ \sum_{t=1}^{\infty} \mathbb{P}[\bar{c}_{k,t} \leq \mu_k^c - \epsilon_{k,t}] + \mathbb{P}[\bar{c}_{l,t} \leq \mu_l^c - \epsilon_{l,t}] + \frac{4\pi^2}{3\epsilon^2} \\
&\leq \pi^2 + \frac{4\pi^2}{3\epsilon^2} \\
&\leq \frac{3\pi^2}{\epsilon^2},
\end{aligned}$$

using the same argument as in Fact 3.

E.5. Proof of Theorem 5. We build upon (4):

$$\begin{aligned}
R_{B,T} &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t,t}\right] + O(1) \\
&= T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^T r_{a_t,t}\right] + \mathbb{E}\left[\sum_{t=\tau^*+1}^T r_{a_t,t}\right] + O(1) \\
&\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^T r_{a_t,t}\right] + \sigma \cdot \mathbb{E}\left[\sum_{t=\tau^*+1}^T c_{a_t,t}\right] + O(1) \\
&\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^T r_{a_t,t}\right] + \sigma \cdot \mathbb{E}\left[\left(\sum_{t=1}^T c_{a_t,t} - B\right)_+\right] + O(1).
\end{aligned}$$

The second inequality is a consequence of Assumption 4:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=\tau^*+1}^T c_{a_t,t}\right] &= \mathbb{E}\left[\sum_{t=1}^T c_{a_t,t}\right] - \mathbb{E}\left[\sum_{t=1}^{\tau^*} c_{a_t,t}\right] \\
&= \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}[c_{a_t,t} \mid \mathcal{F}_{t-1}]\right] - \mathbb{E}\left[\sum_{t=1}^{\infty} I_{\tau^* \geq t} \cdot \mathbb{E}[c_{a_t,t} \mid \mathcal{F}_{t-1}]\right] \\
&= \mathbb{E}\left[\sum_{t=1}^T \mu_{a_t}^c\right] - \mathbb{E}\left[\sum_{t=1}^{\infty} I_{\tau^* \geq t} \cdot \mu_{a_t}^c\right] \\
&= \mathbb{E}\left[\sum_{t=\tau^*+1}^T \mu_{a_t}^c\right] \\
&\geq \frac{1}{\sigma} \cdot \mathbb{E}\left[\sum_{t=\tau^*+1}^T \mu_{a_t}^r\right] = \frac{1}{\sigma} \cdot \mathbb{E}\left[\sum_{t=\tau^*+1}^T r_{a_t,t}\right],
\end{aligned}$$

since τ^* is a stopping time. To derive the third inequality, observe that if $\tau^* = T + 1$, we have:

$$\sum_{t=\tau^*+1}^T c_{a_t,t} = 0 \leq \left(\sum_{t=1}^T c_{a_t,t} - B\right)_+,$$

while if $\tau^* < T + 1$ we have run out of resources before round T , i.e. $\sum_{t=1}^{\tau^*} c_{a_t,t} \geq B$, which implies:

$$\sum_{t=\tau^*+1}^T c_{a_t,t} \leq \sum_{t=\tau^*+1}^T c_{a_t,t} + \sum_{t=1}^{\tau^*} c_{a_t,t} - B$$

$$\leq \left(\sum_{t=1}^T c_{a_t, t} - B \right)_+.$$

Now observe that:

$$\begin{aligned} \mathbb{E}\left[\left(\sum_{t=1}^T c_{a_t, t} - B\right)_+\right] &= \mathbb{E}\left[\left(\sum_{x \text{ pseudo-basis for (3)}} \{b_{x, T} - n_{x, T} \cdot b\}\right)_+\right] \\ &\leq \sum_{x \in \mathcal{B}} \mathbb{E}[|b_{x, T} - n_{x, T} \cdot b|] + \sum_{x \notin \mathcal{B}} \mathbb{E}[n_{x, T}] + \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x)=0}} \mathbb{E}[n_{x, T}] \\ &= O\left(\frac{K^2}{\epsilon^3} \ln(T)\right), \end{aligned}$$

where we use the fact that $c_{k, t} \leq 1$ at any time t and for all arms k for the first inequality and Lemma 10 along with the proof of Lemma 11 for the last equality. Plugging this last inequality back into the regret bound yields:

$$\begin{aligned} R_{B, T} &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^T r_{a_t, t}\right] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[n_{k, T}^x] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x, T}] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &= \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \cdot \left(T - \sum_{x \in \mathcal{B} \mid \Delta_x = 0} \mathbb{E}[n_{x, T}]\right) - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x, T}] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &= \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \mathbb{E}[n_{x, T}] + \sum_{x \notin \mathcal{B}} \mathbb{E}[n_{x, T}] + \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x)=0}} \mathbb{E}[n_{x, T}]\right) \\ &\quad - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x, T}] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &\leq \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \Delta_x \cdot \mathbb{E}[n_{x, T}] + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\ &\leq 2^9 \frac{\lambda^2}{\epsilon^3} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x}\right) \cdot \ln(T) + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right), \end{aligned} \tag{36}$$

where we use Lemma 11 for the third inequality, Lemma 10 along with $\sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \leq \sum_{k=1}^K \xi_k^{x^*} \leq 1$ for the fourth inequality, and Lemma 12 for the last inequality.

E.6. Proof of Theorem 6. Along the same lines as for the case of a single limited resource, we start from inequality (36) derived in the proof of Theorem 5 and apply Lemma 12 only if Δ_x is big enough, taking into account the fact that:

$$\sum_{x \in \mathcal{B}} \mathbb{E}[n_{x, T}] \leq T.$$

Specifically, we have:

$$R_{B, T} \leq \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq T}} \left\{ \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \min\left(\Delta_x \cdot n_x, 2^9 \frac{\lambda^2}{\epsilon^3} \cdot \frac{\ln(T)}{\Delta_x} + \frac{10\pi^2}{3\epsilon^2} \cdot \Delta_x\right) \right\} + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right)$$

$$\begin{aligned}
&\leq \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq T}} \left\{ \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \min(\Delta_x \cdot n_x, 2^9 \frac{\lambda^2}{\epsilon^3} \cdot \frac{\ln(T)}{\Delta_x}) \right\} + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\
&\leq \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq T}} \left\{ \sum_{x \in \mathcal{B}} \sqrt{2^9 \frac{\lambda^2}{\epsilon^3} \cdot \ln(T) \cdot n_x} \right\} + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\
&\leq 2^5 \frac{\lambda}{\epsilon^{3/2}} \cdot \sqrt{\ln(T)} \cdot \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq T}} \left\{ \sum_{x \in \mathcal{B}} \sqrt{n_x} \right\} + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right) \\
&\leq 2^5 \frac{\lambda}{\epsilon^{3/2}} \cdot \sqrt{|\mathcal{B}| \cdot T \cdot \ln(T)} + O\left(\frac{K^2 \cdot \sigma}{\epsilon^3} \ln(T)\right),
\end{aligned}$$

where we use the fact that $\Delta_x \leq \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x*} \leq \sum_{k=1}^K \xi_k^{x*} \leq 1$ for the second inequality, we maximize over each $\Delta_x \geq 0$ to derive the third inequality, and we use Cauchy-Schwartz for the last inequality.

E.7. Proof of Theorem 7. When $b \leq \epsilon/2$, the analysis almost falls back to the case of a single limited resource. Indeed, we have $\tau^* = \tau(B)$ with high probability given that:

$$\begin{aligned}
\mathbb{P}[\tau(B) > T + t] &= \mathbb{P}\left[\sum_{\tau=1}^{T+t} c_{a_\tau, \tau} \leq B\right] \\
&\leq \mathbb{P}\left[\frac{1}{T+t} \cdot \sum_{t=1}^{T+t} c_{a_t, t} \leq \epsilon - (\epsilon - b)\right] \\
&\leq \exp\left(- (T+t) \cdot \frac{\epsilon^2}{2}\right),
\end{aligned}$$

for any $t \in \mathbb{N}$ using Lemma 1. Now observe that, since $b \leq \epsilon/2$, the feasible bases for (3) are exactly the bases x such that $\mathcal{K}_x = \{k\}$ and $\mathcal{C}_x = \{1\}$ for some $k \in \{1, \dots, K\}$, which we denote by $(x_k)_{k=1, \dots, K}$. This shows that $\text{ER}_{\text{OPT}}(B, T) = B \cdot \max_{k=1, \dots, K} \mu_k^r / \mu_k^c$. Moreover note that Assumption 6 is automatically satisfied when $b \leq \epsilon/2$. Hence, with a minor modification of the proof of Lemma 10, we get:

$$\mathbb{E}[n_{x, \tau(B)}] \leq \frac{2^{12}}{\epsilon^3} \cdot \mathbb{E}[\ln(\tau(B))] + \frac{40\pi^2}{3\epsilon^2},$$

for any pseudo-basis x involving two arms or any basis x such that $\mathcal{K}_x = \{k\}$ and $\mathcal{C}_x = \{2\}$ for some arm $k \in \{1, \dots, K\}$. Similarly, a minor modification of Lemma 12 yields:

$$\mathbb{E}[n_{x_k, \tau(B)}] \leq 2^{12} \frac{\lambda^2}{\epsilon^3} \cdot \frac{\mathbb{E}[\ln(\tau(B))]}{\Delta_{x_k}^2} + \frac{40\pi^2}{\epsilon^2},$$

for any $k \in \{1, \dots, K\}$. What is left is to refine the analysis of Theorem 5 as follows:

$$\begin{aligned}
R_{B, T} &\leq B \cdot \max_{k=1, \dots, K} \frac{\mu_k^r}{\mu_k^c} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t, t}\right] + O(1) \\
&\leq B \cdot \max_{k=1, \dots, K} \frac{\mu_k^r}{\mu_k^c} - \mathbb{E}\left[\sum_{t=1}^{\tau(B)} r_{a_t, t}\right] + T \cdot \mathbb{P}[\tau(B) > T] + \mathbb{E}[(\tau(B) - T)_+] + O(1) \\
&\leq B \cdot \max_{k=1, \dots, K} \frac{\mu_k^r}{\mu_k^c} - \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[n_{x_k, \tau(B)}] + O(1) \\
&\leq \max_{k=1, \dots, K} \frac{\mu_k^r}{\mu_k^c} \cdot \left(B - \sum_{k \mid \Delta_{x_k} = 0} \mu_k^c \cdot \mathbb{E}[n_{x_k, \tau(B)}]\right) - \sum_{k \mid \Delta_{x_k} > 0} \mu_k^r \cdot \mathbb{E}[n_{x_k, \tau(B)}] + O(1).
\end{aligned}$$

By definition of $\tau(B)$, we have $B \leq \sum_{t=1}^{\tau(B)} c_{a_t, t}$. Taking expectations on both sides yields:

$$\begin{aligned} B &\leq \sum_{k=1}^K \mu_k^c \cdot \mathbb{E}[n_{k, \tau(B)}] \\ &= \sum_{k \mid \Delta_{x_k} = 0} \mu_k^c \cdot \mathbb{E}[n_{k, \tau(B)}^{x_k}] + \sum_{k \mid \Delta_{x_k} > 0} \mu_k^c \cdot \mathbb{E}[n_{k, \tau(B)}^{x_k}] + \sum_{x \notin \mathcal{B}} \mathbb{E}[n_{x, \tau(B)}]. \end{aligned}$$

Plugging this inequality back into the regret bound, we get:

$$\begin{aligned} R_{B, T} &= \sum_{k \mid \Delta_{x_k} > 0} \mu_k^c \cdot \Delta_{x_k} \cdot \mathbb{E}[n_{k, \tau(B)}^{x_k}] + \max_{k=1, \dots, K} \frac{\mu_k^r}{\mu_k^c} \cdot \sum_{x \notin \mathcal{B}} \mathbb{E}[n_{x, \tau(B)}] + O(1) \\ &\leq \sum_{k \mid \Delta_{x_k} > 0} \mu_k^c \cdot \Delta_{x_k} \cdot \mathbb{E}[n_{k, \tau(B)}^{x_k}] + \frac{2^{12} K^2 \cdot \kappa}{\epsilon^3} \cdot \mathbb{E}[\ln(\tau(B))] + O(1) \\ &\leq 2^{12} \frac{\lambda^2}{\epsilon^3} \cdot \left(\sum_{k \mid \Delta_{x_k} > 0} \frac{1}{\Delta_{x_k}} \right) \cdot \ln\left(\frac{B+1}{\epsilon}\right) + \frac{2^{12} K^2 \cdot \kappa}{\epsilon^3} \cdot \ln\left(\frac{B+1}{\epsilon}\right) + O(1), \end{aligned}$$

where we use Assumption 5 for the second inequality and Lemma 4 for the third inequality. A distribution-independent regret bound of order $O(\sqrt{K \cdot \frac{B+1}{\epsilon}} \cdot \ln(\frac{B+1}{\epsilon}) + \frac{K^2 \cdot \kappa}{\epsilon^3} \cdot \ln(\frac{B+1}{\epsilon}))$ can be derived from the penultimate inequality along the same lines as in Theorem 6.

Appendix F: Proofs for Section 7.

F.1. Preliminary work for the proofs of Section 7.

Concentration inequality. We will use the following inequality repeatedly. For a given round $\tau \geq t_{\text{ini}}$ and a basis x :

$$\begin{aligned} &\mathbb{P}[\exists(k, i) \in \mathcal{K}_x \times \{1, \dots, C\}, |\bar{c}_{k, \tau}(i) - \mu_k^c(i)| > \frac{\epsilon^3}{16 \cdot (C+2)!^2}] \\ &\leq \sum_{s=t_{\text{ini}}/K}^T \sum_{\substack{k \in \mathcal{K}_x \\ i \in \{1, \dots, C\}}} \mathbb{P}[|\bar{c}_{k, \tau}(i) - \mu_k^c(i)| > \frac{\epsilon^3}{16 \cdot (C+2)!^2}, n_{k, \tau} = s] \\ &\leq 2 \cdot C^2 \cdot \sum_{s=t_{\text{ini}}/K}^{\infty} \exp\left(-s \cdot \frac{\epsilon^6}{2^7 \cdot (C+2)!^4}\right) \\ &\leq 2^9 \frac{(C+3)!^4}{\epsilon^6 \cdot T^2}, \end{aligned} \tag{37}$$

using Lemma 1, the inequality $\exp(-x) \leq 1 - x/2$ for $x \in [0, 1]$, and the fact that we pull each arm $t_{\text{ini}}/K \geq 2^8 \frac{(C+2)!^4}{\epsilon^6} \cdot \ln(T)$ times during the initialization phase.

Useful matrix inequalities. For any basis x , assume that $\{|\bar{c}_{k, t}(i) - \mu_k^c(i)| \leq \frac{\epsilon^3}{16 \cdot (C+2)!^2}\}$, for any arm $k \in \mathcal{K}_x$ and resource $i \in \mathcal{C}_x$. We have:

$$\begin{aligned} |\det(\bar{A}_{x, t}) - \det(A_x)| &= \left| \sum_{\sigma \in S(\mathcal{K}_x, \mathcal{C}_x)} \left[\prod_{k \in \mathcal{K}_x} \bar{c}_{k, t}(\sigma(k)) - \prod_{k \in \mathcal{K}_x} \mu_k^c(\sigma(k)) \right] \right| \\ &= \left| \sum_{\sigma \in S(\mathcal{K}_x, \mathcal{C}_x)} \sum_{l \in \mathcal{K}_x} \prod_{k < l} \bar{c}_{k, t}(\sigma(k)) \cdot [\bar{c}_{l, t}(\sigma(l)) - \mu_l^c(\sigma(l))] \cdot \prod_{k > l} \mu_k^c(\sigma(k)) \right| \\ &\leq \sum_{\sigma \in S(\mathcal{K}_x, \mathcal{C}_x)} \sum_{l \in \mathcal{K}_x} |\bar{c}_{l, t}(\sigma(l)) - \mu_l^c(\sigma(l))| \\ &\leq \frac{\epsilon^3}{16 \cdot (C+2)!} \\ &\leq \frac{\epsilon}{2}, \end{aligned}$$

since the amounts of resources consumed at any round are no larger than 1. This yields:

$$|\det(\bar{A}_{x,t})| \geq |\det(A_x)| - |\det(\bar{A}_{x,t}) - \det(A_x)| \geq \frac{\epsilon}{2}, \quad (38)$$

using Assumption 8. Now consider any vector c such that $\|c\|_\infty \leq 1$, we have:

$$\begin{aligned} & |c^\top \bar{A}_{x,t}^{-1} b_{\mathcal{K}_x} - c^\top A_x^{-1} b_{\mathcal{K}_x}| \\ &= \frac{1}{|\det(A_x)| \cdot |\det(\bar{A}_{x,t})|} \cdot |\det(A_x) \cdot c^\top \text{adj}(\bar{A}_{x,t}) b_{\mathcal{K}_x} - \det(\bar{A}_{x,t}) \cdot c^\top \text{adj}(A_x) b_{\mathcal{K}_x}| \\ &\leq \frac{1}{|\det(\bar{A}_{x,t})|} \cdot |c^\top (\text{adj}(\bar{A}_{x,t}) - \text{adj}(A_x)) b_{\mathcal{K}_x}| \\ &\quad + \frac{1}{|\det(A_x)| \cdot |\det(\bar{A}_{x,t})|} \cdot |\det(\bar{A}_{x,t}) - \det(A_x)| \cdot |c^\top \text{adj}(A_x) b_{\mathcal{K}_x}| \\ &\leq \frac{\epsilon^2}{8} + \frac{\epsilon}{8 \cdot (C+2)!} \cdot \|c\|_2 \cdot \|\text{adj}(A_x) b_{\mathcal{K}_x}\|_2 \\ &\leq \frac{\epsilon}{4}. \end{aligned}$$

The second inequality is obtained using Assumption 8 and (38) by proceeding along the same lines as above to bound the difference between two determinants for each component of $\text{adj}(\bar{A}_{x,t}) - \text{adj}(A_x)$. The last inequality is obtained using $\|c\|_2 \leq \sqrt{C}$ and the fact that each component of A_x is smaller than 1. If we take $c = e_k$ for $k \in \mathcal{K}_x$, this yields:

$$|\xi_{k,t}^x - \xi_k^x| \leq \frac{\epsilon}{4}. \quad (39)$$

If we take $c = (\mu_k^c(i))_{k \in \mathcal{K}_x}$, for any $i \in \{1, \dots, C\}$, we get:

$$\left| \sum_{k \in \mathcal{K}_x} \mu_k^c(i) \cdot \xi_{k,t}^x - \sum_{k \in \mathcal{K}_x} \mu_k^c(i) \cdot \xi_k^x \right| \leq \frac{\epsilon}{4} \quad (40)$$

and

$$\begin{aligned} \left| \sum_{k \in \mathcal{K}_x} \bar{c}_{k,t}(i) \cdot \xi_{k,t}^x - \sum_{k \in \mathcal{K}_x} \mu_k^c(i) \cdot \xi_k^x \right| &= |\bar{c}^\top \bar{A}_{x,t}^{-1} b_{\mathcal{K}_x} - c^\top A_x^{-1} b_{\mathcal{K}_x}| \\ &\leq |(\bar{c} - c)^\top \bar{A}_{x,t}^{-1} b_{\mathcal{K}_x}| + |c^\top \bar{A}_{x,t}^{-1} b_{\mathcal{K}_x} - c^\top A_x^{-1} b_{\mathcal{K}_x}| \\ &\leq \|\bar{c} - c\|_2 \cdot \frac{1}{|\det(\bar{A}_{x,t})|} \cdot \|\text{adj}(\bar{A}_{x,t}) b_{\mathcal{K}_x}\|_2 + \frac{\epsilon}{4} \\ &\leq \sqrt{C} \cdot \frac{\epsilon^3}{16 \cdot (C+2)!^2} \cdot \frac{2}{\epsilon} \cdot C! + \frac{\epsilon}{4} \\ &\leq \frac{\epsilon}{2}, \end{aligned} \quad (41)$$

where $\bar{c} = (\bar{c}_{k,t}(i))_{k \in \mathcal{K}_x}$.

F.2. Proof of Lemma 13. First, consider a basis $x \notin \mathcal{B}$. Since x is ϵ -non-degenerate by Assumption 8, there must exist $k \in \mathcal{K}_x$ such that $\xi_k^x \leq -\epsilon$ or $i \in \{1, \dots, C\}$ such that $\sum_{k=1}^K \mu_k^c(i) \cdot \xi_k^x \geq b(i) + \epsilon$. Let us assume that we are in the first situation (the proof is symmetric in the other scenario). Using (39) in the preliminary work of Section F.1, we have:

$$\xi_{k,t}^x = \xi_k^x + (\xi_{k,t}^x - \xi_k^x) \leq -\frac{\epsilon}{2}, \quad (42)$$

if $\{|\bar{c}_{k,t}(i) - \mu_k^c(i)| \leq \frac{\epsilon^3}{16 \cdot (C+2)!^2}\}$ for all arms $k \in \mathcal{K}_x$ and resources $i \in \{1, \dots, C\}$. Hence:

$$\mathbb{E}[n_{x,T}] = \mathbb{E}\left[\sum_{t=v_{\text{ini}}}^T I_{x_t=x}\right]$$

$$\begin{aligned}
&\leq \sum_{t=t_{\text{ini}}}^T \mathbb{P}[\xi_{k,t}^x \geq 0] \\
&\leq \sum_{t=t_{\text{ini}}}^T \sum_{\substack{k \in \mathcal{K}_x \\ i \in \{1, \dots, C\}}} \mathbb{P}[|\bar{c}_{k,t}(i) - \mu_k^c(i)| > \frac{\epsilon^3}{16 \cdot (C+2)!^2}] \\
&\leq 2^9 \frac{(C+3)!^4}{\epsilon^6},
\end{aligned}$$

where the third inequality is derived with (37).

Second, consider a pseudo-basis x for (3) that is not a basis. Since $\det(A_x) = 0$, either every component of A_x is 0, in which case $\det(\bar{A}_{x,t}) = 0$ at every round t and x can never be selected, or there exists a basis \tilde{x} for (3) with $\mathcal{K}_{\tilde{x}} \subset \mathcal{K}_x$ and $\mathcal{C}_{\tilde{x}} \subset \mathcal{C}_x$ along with coefficients $(\alpha_{kl})_{k \in \mathcal{K}_x - \mathcal{K}_{\tilde{x}}, l \in \mathcal{K}_{\tilde{x}}}$ such that $\mu_k^c(i) = \sum_{l \in \mathcal{K}_{\tilde{x}}} \alpha_{kl} \cdot \mu_l^c(i)$ for any resource $i \in \mathcal{C}_x$. Assuming we are in the second scenario and since \tilde{x} is ϵ -non-degenerate by Assumption 8, we have $\sum_{k=1}^K \mu_k^c(j) \cdot \xi_k^{\tilde{x}} \leq b(j) - \epsilon$ for any $j \notin \mathcal{C}_{\tilde{x}}$. Take $i \in \mathcal{C}_x - \mathcal{C}_{\tilde{x}}$. Suppose that x is feasible for (8) at round t and assume by contradiction that $\{|\bar{c}_{k,t}(j) - \mu_k^c(j)| \leq \frac{\epsilon^3}{16 \cdot (C+2)!^2}\}$ for all arms $k \in \mathcal{K}_x$ and resources $j \in \{1, \dots, C\}$. Using the notations $\tilde{\xi}_{k,t}^x = \xi_{k,t}^x + \sum_{l \in \mathcal{K}_x - \mathcal{K}_{\tilde{x}}} \alpha_{kl} \cdot \xi_{l,t}^x$, we have, for any resource $j \in \mathcal{C}_{\tilde{x}}$:

$$\begin{aligned}
b(j) &= \sum_{l \in \mathcal{K}_x} \bar{c}_{l,t}(j) \cdot \xi_{l,t}^x \\
&= \alpha(j) + \sum_{l \in \mathcal{K}_x} \mu_l^c(j) \cdot \xi_{l,t}^x \\
&= \alpha(j) + \sum_{l \in \mathcal{K}_{\tilde{x}}} \mu_l^c(j) \cdot \tilde{\xi}_{l,t}^x,
\end{aligned}$$

where $|\alpha(j)| \leq \frac{\epsilon^3}{16 \cdot (C+2)!^2}$ since $\xi_k \in [0, 1] \forall k \in \{1, \dots, K\}$ for any feasible solution to (8). We get $A_{\tilde{x}} \tilde{\xi}_{\mathcal{K}_{\tilde{x}},t}^x = b_{\mathcal{C}_{\tilde{x}}} - \alpha_{\mathcal{C}_{\tilde{x}}}$ while $A_{\tilde{x}} \xi_{\mathcal{K}_{\tilde{x}}}^{\tilde{x}} = b_{\mathcal{C}_{\tilde{x}}}$. We derive:

$$\begin{aligned}
& \left| \sum_{k=1}^K \bar{c}_{k,t}(i) \cdot \xi_{k,t}^x - \sum_{k=1}^K \mu_k^c(i) \cdot \xi_k^{\tilde{x}} \right| \\
& \leq \left| \sum_{k=1}^K \mu_k^c(i) \cdot \xi_{k,t}^x - \sum_{k=1}^K \mu_k^c(i) \cdot \xi_k^{\tilde{x}} \right| + \frac{\epsilon^3}{16 \cdot (C+2)!^2} \\
& \leq \left| \sum_{k \in \mathcal{K}_{\tilde{x}}} \mu_k^c(i) \cdot (\tilde{\xi}_{k,t}^x - \xi_k^{\tilde{x}}) \right| + \frac{\epsilon}{4} \\
& = \frac{\epsilon}{4} + |\mu_{\mathcal{K}_{\tilde{x}}}^c(i)^\top A_{\tilde{x}}^{-1} \alpha_{\mathcal{C}_{\tilde{x}}}| \\
& \leq \frac{\epsilon}{4} + \sqrt{C} \cdot \frac{1}{|\det(A_{\tilde{x}})|} \cdot \|\text{adj}(A_{\tilde{x}}) \alpha_{\mathcal{C}_{\tilde{x}}}\|_2 \\
& \leq \frac{\epsilon}{4} + \frac{(C+1)!}{\epsilon} \cdot \frac{\epsilon^3}{16 \cdot (C+2)!^2} \\
& \leq \frac{\epsilon}{2}.
\end{aligned}$$

Thus we obtain:

$$\sum_{k=1}^K \bar{c}_{k,t}(i) \cdot \xi_{k,t}^x \leq \sum_{k=1}^K \mu_k^c(i) \cdot \xi_k^{\tilde{x}} + \frac{\epsilon}{2} \leq b(i) - \frac{\epsilon}{2} < b(i),$$

a contradiction since this inequality must be binding by definition if x is selected at round t . We finally conclude:

$$\begin{aligned} \mathbb{E}[n_{x,T}] &= \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{x_t=x}\right] \\ &\leq \sum_{t=t_{\text{ini}}}^T \sum_{\substack{k \in \mathcal{K}_x \\ j \in \{1, \dots, C\}}} \mathbb{P}[|\bar{c}_{k,t}(j) - \mu_k^c(j)| > \frac{\epsilon^3}{16 \cdot (C+2)!^2}] \\ &\leq 2^9 \frac{(C+3)!^4}{\epsilon^6}. \end{aligned}$$

F.3. Proof of Lemma 14.

Proof of (12). Consider a resource $i \in \mathcal{C}_x$ and $u \geq 1$. We study $\mathbb{P}[b_{x,t}(i) - n_{x,t} \cdot b(i) \geq u]$ but the same technique can be used to bound $\mathbb{P}[b_{x,t}(i) - n_{x,t} \cdot b(i) \leq -u]$. If $b_{x,t}(i) - n_{x,t} \cdot b(i) \geq u$, it must be that $e_{i,\tau}^x = -1$ for at least $s \geq \lfloor u \rfloor$ rounds $\tau = t_1 \leq \dots \leq t_s \leq t-1$ where x was selected at Step-Simplex since the last time, denoted by $t_0 < t_1$, where x was selected at Step-Simplex and the budget was below the target, i.e. $b_{x,t_0}(i) \leq n_{x,t_0} \cdot b(i)$ (because the amounts of resources consumed at each round are bounded by 1). Moreover, we have:

$$\begin{aligned} \sum_{q=1}^s c_{a_{t_q}, t_q}(i) &= \sum_{\tau=t_0+1}^{t-1} I_{x_\tau=x} \cdot c_{a_\tau, \tau}(i) \\ &= b_{x,t}(i) - b_{x,t_0+1}(i) \\ &\geq (n_{x,t} - n_{x,t_0}) \cdot b(i) + u - 1 \\ &\geq s \cdot b(i) + u - 1. \end{aligned}$$

Hence:

$$\begin{aligned} &\mathbb{P}[b_{x,t}(i) - n_{x,t} \cdot b(i) \geq u] \\ &\leq \sum_{s=\lfloor u \rfloor}^t \mathbb{P}\left[\sum_{q=1}^s c_{a_{t_q}, t_q}(i) \geq s \cdot b(i) + u - 1; e_{i,t_q}^x = -1 \forall q \in \{1, \dots, s\}\right] \\ &\leq \sum_{s=\lfloor u \rfloor}^t \mathbb{P}\left[\sum_{q=1}^s c_{a_{t_q}, t_q}(i) \geq s \cdot b(i); \sum_{k=1}^K \mu_k^c(i) \cdot p_{k,t_q}^x \leq b(i) - \frac{\epsilon^2}{4 \cdot (C+1)!} \forall q \in \{1, \dots, s\}\right] \\ &\quad + \sum_{\tau=t_{\text{ini}}}^T \sum_{\substack{k \in \mathcal{K}_x \\ j \in \{1, \dots, C\}}} \mathbb{P}[|\bar{c}_{k,\tau}(j) - \mu_k^c(j)| > \frac{\epsilon^3}{16 \cdot (C+2)!^2}] \\ &\leq \sum_{s=\lfloor u \rfloor}^{\infty} \exp(-2s \cdot (\frac{\epsilon^2}{4 \cdot (C+1)!})^2) + 2^9 \frac{(C+3)!^4}{\epsilon^6} \cdot \frac{1}{T} \\ &\leq 16 \frac{(C+1)!^2}{\epsilon^4} \cdot \exp(-u \cdot (\frac{\epsilon^2}{4 \cdot (C+1)!})^2) + 2^9 \frac{(C+3)!^4}{\epsilon^6} \cdot \frac{1}{T}. \end{aligned}$$

The last inequality is obtained using $\exp(-x) \leq 1 - x/2$ for $x \in [0, 1]$. The third inequality is derived using Lemma 1 for the first term and (37) for the second term. The second inequality is obtained by observing that if x was selected at time τ and $|\bar{c}_{k,\tau}(j) - \mu_k^c(j)| \leq \frac{\epsilon^3}{16 \cdot (C+2)!^2}$ for any arm $k \in \mathcal{K}_x$ and resource $j \in \{1, \dots, C\}$, then we must have $\delta_{x,\tau}^* \geq \frac{\epsilon^2}{4 \cdot (C+1)!}$. Indeed, using (38) and (39), we have, for $\delta \leq \frac{\epsilon^2}{4 \cdot (C+1)!}$ and any arm $k \in \mathcal{K}_x$:

$$c^\top \bar{A}_{x,\tau}^{-1} (b_{\mathcal{C}_x} + \delta \cdot e_\tau^x) = \xi_{k,\tau}^x + \delta \cdot c^\top \bar{A}_{x,\tau}^{-1} e_\tau^x$$

$$\begin{aligned}
&\geq \xi_k^x - |\xi_k^x - \xi_{k,\tau}^x| - \delta \cdot |c^\top \bar{A}_{x,\tau}^{-1} e_\tau^x| \\
&\geq \frac{\epsilon}{2} - \delta \cdot \|\bar{A}_{x,\tau}^{-1} e_\tau^x\|_2 \\
&\geq \frac{\epsilon}{2} - \delta \cdot \frac{1}{|\det(\bar{A}_{x,\tau})|} \cdot \|\text{adj}(\bar{A}_{x,\tau}) e_\tau^x\|_2 \\
&\geq \frac{\epsilon}{2} - \frac{2\delta}{\epsilon} \cdot \sqrt{C} \cdot C! \\
&\geq 0,
\end{aligned}$$

where $c = e_k$ and since x is ϵ -non-degenerate by Assumption 8. Similarly, using (38) and (41), we have, for $\delta \leq \frac{\epsilon^2}{4 \cdot (C+1)!}$ and any resource $j \notin \mathcal{C}_x$:

$$\begin{aligned}
c^\top \bar{A}_{x,\tau}^{-1} (b_{\mathcal{C}_x} + \delta \cdot e_\tau^x) &= \sum_{k \in \mathcal{K}_x} \bar{c}_{k,\tau}(j) \cdot \xi_{k,\tau}^x + \delta \cdot c^\top \bar{A}_{x,\tau}^{-1} e_\tau^x \\
&\leq \sum_{k \in \mathcal{K}_x} \mu_k^c(j) \cdot \xi_{k,\tau}^x + \left| \sum_{k \in \mathcal{K}_x} \bar{c}_{k,\tau}(j) \cdot \xi_{k,\tau}^x - \sum_{k \in \mathcal{K}_x} \mu_k^c(j) \cdot \xi_{k,\tau}^x \right| + \delta \cdot |c^\top \bar{A}_{x,\tau}^{-1} e_\tau^x| \\
&\leq b(j) - \frac{\epsilon}{2} + \delta \cdot \sqrt{C} \cdot \|\bar{A}_{x,\tau}^{-1} e_\tau^x\|_2 \\
&\leq b(j) - \frac{\epsilon}{2} + \frac{2\delta}{\epsilon} \cdot (C+1)! \\
&\leq b(j),
\end{aligned}$$

where $c = (\bar{c}_{k,\tau}(j))_{k \in \mathcal{K}_x}$ and since x is ϵ -non-degenerate by Assumption 8.

Proof of (14). First observe that, using (12), we have:

$$\begin{aligned}
&\max_{i \in \mathcal{C}_x} |\mathbb{E}[b_{x,T}(i)] - \mathbb{E}[n_{x,T}] \cdot b(i)| \\
&\leq \mathbb{E}[\max_{i \in \mathcal{C}_x} |b_{x,T}(i) - n_{x,T} \cdot b(i)|] \\
&= \int_0^T \mathbb{P}[\max_{i \in \mathcal{C}_x} |b_{x,T}(i) - n_{x,T} \cdot b(i)| \geq u] du \\
&= \sum_{i \in \mathcal{C}_x} \int_0^T \mathbb{P}[|b_{x,T}(i) - n_{x,T} \cdot b(i)| \geq u] du \\
&\leq 32C \cdot \frac{(C+1)!^2}{\epsilon^4} \cdot \int_0^T \exp(-u \cdot (\frac{\epsilon^2}{4 \cdot (C+1)!})^2) du + C + C \cdot 2^9 \frac{(C+3)!^4}{\epsilon^6} \\
&= 2^9 C \cdot \frac{(C+1)!^4}{\epsilon^8} + C + C \cdot 2^9 \frac{(C+3)!^4}{\epsilon^6} \\
&\leq 2^{10} C \cdot \frac{(C+3)!^4}{\epsilon^8}.
\end{aligned}$$

Now observe that, for any resource $i \in \mathcal{C}_x$, we have $\mathbb{E}[b_{x,T}(i)] = \sum_{k \in \mathcal{K}_x} \mu_k^c(i) \cdot \mathbb{E}[n_{k,T}^x]$. Hence, defining the vector $p = (\frac{\mathbb{E}[n_{k,T}^x]}{\mathbb{E}[n_{x,T}]})_{k \in \mathcal{K}_x}$, we get:

$$\begin{aligned}
\mathbb{E}[n_{x,T}] \cdot \|p - \xi^x\|_2 &= \mathbb{E}[n_{x,T}] \cdot \|A_x^{-1} A_x (p - \xi^x)\|_2 \\
&\leq \frac{1}{\|\det(A_x)\|_2} \cdot \|\mathbb{E}[n_{x,T}] \cdot A_x (p - \xi^x)\|_2 \\
&= \frac{1}{|\det(A_x)|} \cdot \|\text{adj}(A_x)\|_2 \cdot \|(\mathbb{E}[b_{x,T}(i)])_{i \in \mathcal{C}_x} - (\mathbb{E}[n_{x,T}] \cdot b(i))_{i \in \mathcal{C}_x}\|_2 \\
&\leq \frac{1}{\epsilon} \cdot (C+1)! \cdot \sqrt{C} \cdot 2^{10} C \cdot \frac{(C+3)!^4}{\epsilon^8} \\
&\leq 2^{10} \frac{(C+3)!^5}{\epsilon^9},
\end{aligned}$$

using Assumption 8. Finally we obtain:

$$\begin{aligned} \mathbb{E}[n_{x,T}] \cdot \xi_k^x - \mathbb{E}[n_{k,T}^x] &\leq \mathbb{E}[n_{x,T}] \cdot \|p - \xi^x\|_2 \\ &\leq 2^{10} \frac{(C+3)!^5}{\epsilon^9}, \end{aligned}$$

for any arm $k \in \mathcal{K}_x$.

Proof of (13). Consider a resource $i \notin \mathcal{C}_x$ and assume that $b_{x,t}(i) - n_{x,t} \cdot b(i) \geq 2^8 \frac{(C+3)!^3}{\epsilon^6} \cdot \ln(T)$. By contradiction, suppose that:

- $|b_{x,t}(j) - n_{x,t} \cdot b(j)| \leq 16 \frac{(C+1)!^2}{\epsilon^4} \cdot \ln(T)$ for all resources $j \in \mathcal{C}_x$,
- $|\mu_k^c(j) - \tilde{c}_k(j)| \leq \frac{\epsilon^2}{8 \cdot (C+2)!}$ for all resources $j \in \{1, \dots, C\}$ and for all arms $k \in \mathcal{K}_x$ such that $n_{k,t}^x \geq 2^6 \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(T)$, where $\tilde{c}_k(j)$ denotes the empirical average amount of resource j consumed when selecting basis x and pulling arm k , i.e. $\tilde{c}_k(j) = \frac{1}{n_{k,t}^x} \cdot \sum_{\tau=t_{\text{ini}}}^{t-1} I_{x_\tau=x} \cdot I_{a_\tau=k} \cdot c_{k,\tau}(j)$.

Observe that if $b_{x,t}(i) - n_{x,t} \cdot b(i) \geq 2^8 \frac{(C+3)!^3}{\epsilon^6} \cdot \ln(T)$, it must be that x has been selected at least $2^8 \frac{(C+3)!^3}{\epsilon^6} \cdot \ln(T)$ times at Step-Simplex since t_{ini} , i.e. $n_{x,t} \geq 2^8 \frac{(C+3)!^3}{\epsilon^6} \cdot \ln(T)$. We can partition \mathcal{K}_x into two sets \mathcal{K}_x^1 and \mathcal{K}_x^2 such that $n_{k,t}^x \geq 2^6 \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(T)$ for all $k \in \mathcal{K}_x^1$ and $n_{k,t}^x < 2^6 \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(T)$ for all $k \in \mathcal{K}_x^2$. We get, for any $j \in \mathcal{C}_x$:

$$\begin{aligned} 16 \frac{(C+1)!^2}{\epsilon^4} \cdot \ln(T) &\geq |b_{x,t}(j) - n_{x,t} \cdot b(j)| \\ &\geq n_{x,t} \cdot \left| \sum_{k \in \mathcal{K}_x^1} \tilde{c}_k(j) \cdot p_k - b(j) \right| - \sum_{k \in \mathcal{K}_x^2} n_{k,t}^x \\ &\geq n_{x,t} \cdot \left| \sum_{k \in \mathcal{K}_x^1} \tilde{c}_k(j) \cdot p_k - b(j) \right| - 2^6 C \cdot \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(T), \end{aligned}$$

where $p_k = \frac{n_{k,t}^x}{n_{x,t}}$ for $k \in \mathcal{K}_x^1$ and $p_k = 0$ otherwise. Hence:

$$\begin{aligned} \left| \sum_{k \in \mathcal{K}_x} \mu_k^c(j) \cdot p_k - b(j) \right| &\leq \max_{k \in \mathcal{K}_x^1} |\mu_k^c(j) - \tilde{c}_k(j)| + \left| \sum_{k \in \mathcal{K}_x^1} \tilde{c}_k(j) \cdot p_k - b(j) \right| \\ &\leq \frac{\epsilon^2}{8 \cdot (C+2)!} + \frac{(16 \frac{(C+1)!^2}{\epsilon^4} + 2^6 C \cdot \frac{(C+2)!^2}{\epsilon^4}) \cdot \ln(T)}{n_{x,t}} \\ &\leq \frac{\epsilon^2}{4 \cdot (C+2)!}, \end{aligned}$$

where we use the fact that $\sum_{k=1}^K p_k \leq 1$ and $p_k \geq 0$ for any arm k for the first inequality and $n_{x,t} \geq 2^8 \frac{(C+3)!^3}{\epsilon^6} \cdot \ln(T)$ for the last one. We get:

$$\begin{aligned} \|p - \xi^x\|_2 &= \|A_x^{-1} A_x (p - \xi^x)\|_2 \\ &\leq \|A_x^{-1}\|_2 \cdot \|A_x (p - \xi^x)\|_2 \\ &\leq \frac{1}{|\det(A_x)|} \cdot \|\text{adj}(A_x)\|_2 \cdot \sqrt{C} \cdot \frac{\epsilon^2}{4 \cdot (C+2)!} \\ &\leq \frac{1}{\epsilon} \cdot (C+1)! \cdot \sqrt{C} \cdot \frac{\epsilon^2}{4 \cdot (C+2)!} \\ &\leq \frac{\epsilon}{4 \cdot \sqrt{C}}, \end{aligned}$$

using Assumption 8. Hence:

$$2^8 \frac{(C+3)!^3}{\epsilon^6} \cdot \ln(T) \leq b_{x,t}(i) - n_{x,t} \cdot b(i)$$

$$\begin{aligned}
&\leq n_{x,t} \cdot \left(\sum_{k \in \mathcal{K}_x^1} \tilde{c}_k(i) \cdot p_k - b(i) \right) + \sum_{k \in \mathcal{K}_x^2} n_{k,t}^x \\
&\leq n_{x,t} \cdot \left(\sum_{k \in \mathcal{K}_x^1} \tilde{c}_k(i) \cdot p_k - b(i) \right) + 2^6 C \cdot \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(T).
\end{aligned}$$

Using the shorthand notation $c = (\mu_k^c(i))_{k \in \mathcal{K}_x}$, this implies:

$$\begin{aligned}
0 &\leq \sum_{k \in \mathcal{K}_x^1} \tilde{c}_k(i) \cdot p_k - b(i) \\
&\leq \sum_{k \in \mathcal{K}_x^1} \mu_k^c(i) \cdot p_k + \frac{\epsilon^2}{8 \cdot (C+2)!} - \left(\sum_{k \in \mathcal{K}_x} \mu_k^c(i) \cdot \xi_k^x + \epsilon \right) \\
&\leq c^\top (p - \xi^x) - \frac{\epsilon}{2} \\
&\leq \sqrt{C} \cdot \|p - \xi^x\|_2 - \frac{\epsilon}{2} \\
&< 0,
\end{aligned}$$

a contradiction. Note that we use the fact that $\sum_{k=1}^K p_k \leq 1$ and $p_k \geq 0$ for any arm k and Assumption 8 for the second inequality. We conclude that:

$$\begin{aligned}
&\mathbb{P}[b_{x,t}(i) - n_{x,t} \cdot b(i) \geq 2^8 \frac{(C+3)!^3}{\epsilon^6} \cdot \ln(T)] \\
&\leq \sum_{j \in \mathcal{C}_x} \mathbb{P}[|b_{x,t}(j) - n_{x,t} \cdot b(j)| \geq 16 \frac{(C+1)!^2}{\epsilon^4} \cdot \ln(T)] \\
&+ \sum_{\substack{k \in \mathcal{K}_x \\ j \in \{1, \dots, C\}}} \mathbb{P}[|\mu_k^c(j) - \tilde{c}_k(j)| \geq \frac{\epsilon^2}{8 \cdot (C+2)!}; n_{k,t}^x \geq 2^6 \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(T)] \\
&\leq 2^5 C \cdot \frac{(C+1)!^2}{\epsilon^4 \cdot T} + C \cdot T \cdot 2^9 \frac{(C+3)!^4}{\epsilon^6 \cdot T^2} \\
&+ \sum_{\substack{k \in \mathcal{K}_x \\ j \in \{1, \dots, C\}}} \sum_{s=2^6 \cdot (C+2)!^2 / \epsilon^4 \cdot \ln(T)} \mathbb{P}[|\mu_k^c(j) - \tilde{c}_k(j)| \geq \frac{\epsilon^2}{8 \cdot (C+2)!}; n_{k,t}^x = s] \\
&\leq 2C \cdot T \cdot 2^9 \frac{(C+3)!^4}{\epsilon^6 \cdot T^2} + 2^8 C^2 \cdot \frac{(C+2)!^2}{\epsilon^4 \cdot T^2} \\
&\leq 2^{10} \frac{(C+4)!^4}{\epsilon^6 \cdot T},
\end{aligned}$$

where we use (12) for the second inequality and Lemma 1 for the third inequality.

F.4. Proof of Lemma 15. Consider any suboptimal basis $x \in \mathcal{B}$. The proof is along the same lines as for Lemmas 5, 8, and 12. We break down the analysis in a series of facts where we emphasize the main differences. We start off with an inequality similar to Fact 1. We use the shorthand notation $\beta_x = 2^{10} \frac{(C+3)!^3}{\epsilon^6} \cdot \left(\frac{\lambda}{\Delta_x}\right)^2$.

Fact 11

$$\begin{aligned}
\mathbb{E}[n_{x,T}] &\leq 2\beta_x \cdot \ln(T) + 2^9 \frac{(C+3)!^4}{\epsilon^6} \\
&+ \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{x^* \in \mathcal{B}_t} \right].
\end{aligned} \tag{43}$$

Proof. Similarly as in Fact 1, we have:

$$\mathbb{E}[n_{x,T}] \leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right].$$

This yields:

$$\mathbb{E}[n_{x,T}] \leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{x^* \in \mathcal{B}_t}\right] + \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{x^* \notin \mathcal{B}_t}\right].$$

Using (39), (41), and Assumption 8, we have:

$$\xi_{k,t}^{x^*} = \xi_k^{x^*} - (\xi_k^{x^*} - \xi_{k,t}^{x^*}) \geq \frac{\epsilon}{2} \geq 0,$$

for any $k \in \mathcal{K}_{x^*}$ and

$$\begin{aligned} \sum_{k \in \mathcal{K}_{x^*}} \bar{c}_{k,t}(i) \cdot \xi_{k,t}^{x^*} &= \sum_{k \in \mathcal{K}_{x^*}} \mu_k^c(i) \cdot \xi_k^{x^*} + \left(\sum_{k \in \mathcal{K}_{x^*}} \bar{c}_{k,t}(i) \cdot \xi_{k,t}^{x^*} - \sum_{k \in \mathcal{K}_{x^*}} \mu_k^c(i) \cdot \xi_k^{x^*} \right) \\ &\leq b(i) - \epsilon + \frac{\epsilon}{2} \\ &\leq b(i) - \frac{\epsilon}{2} \\ &\leq b(i), \end{aligned}$$

for any resource $i \notin \mathcal{C}_{x^*}$ if $\{|\bar{c}_{l,t}(j) - \mu_l^c(j)| \leq \frac{\epsilon^3}{16 \cdot (C+2)!^2}\}$ for any arm $l \in \mathcal{K}_{x^*}$ and resource $j \in \{1, \dots, C\}$. Hence:

$$\begin{aligned} \mathbb{E}[n_{x,T}] &\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{x^* \in \mathcal{B}_t}\right] \\ &\quad + \sum_{t=t_{\text{ini}}}^T \sum_{\substack{l \in \mathcal{K}_{x^*} \\ j \in \{1, \dots, C\}}} \mathbb{P}[|\bar{c}_{l,t}(j) - \mu_l^c(j)| > \frac{\epsilon^3}{16 \cdot (C+2)!^2}] \\ &\leq 2\beta_x \cdot \ln(T) + \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{x^* \in \mathcal{B}_t}\right] + 2^9 \frac{(C+3)!^4}{\epsilon^6}, \end{aligned}$$

where we bound the third term appearing in the right-hand side using (37).

The remainder of this proof is dedicated to show that the last term in (43) can be bounded by a constant. This term can be broken down in three terms similarly as in Lemmas 5 and 8.

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{x_t=x} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{x^* \in \mathcal{B}_t}\right] \\ &\leq \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x,t} + E_{x,t} \geq \text{obj}_{x^*,t} + E_{x^*,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)} \cdot I_{x \in \mathcal{B}_t, x^* \in \mathcal{B}_t}\right] \\ &\leq \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{x \in \mathcal{B}_t}\right] \tag{44} \end{aligned}$$

$$+ \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*,t} \leq \text{obj}_{x^*} - E_{x^*,t}} \cdot I_{x^* \in \mathcal{B}_t}\right] \tag{45}$$

$$+ \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}} \cdot I_{x \in \mathcal{B}_t} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right]. \tag{46}$$

Fact 12

$$\mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}} \cdot I_{x \in \mathcal{B}_t} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \leq 2^{11} \frac{(C+4)!^4}{\epsilon^6}.$$

Proof. Using the shorthand notation $\alpha_x = 8\left(\frac{\lambda}{\Delta_x}\right)^2$, we have:

$$\begin{aligned} & \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}} \cdot I_{x \in \mathcal{B}_t} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\ & \leq \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\Delta_x < 2\lambda \cdot \max_{k \in \mathcal{K}_x} \epsilon_{k,t}} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\ & \leq \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\min_{k \in \mathcal{K}_x} n_{k,t} \leq \alpha_x \ln(t)} \cdot I_{n_{x,t} \geq \beta_x \ln(t)}\right] \\ & \leq \sum_{t=t_{\text{ini}}}^T \sum_{k \in \mathcal{K}_x} \mathbb{P}[n_{k,t} \leq \alpha_x \ln(t); n_{x,t} \geq \beta_x \ln(t)], \end{aligned}$$

since $\sum_{l=1}^K \xi_{l,t}^x \leq 1$ and $\xi_{l,t}^x \geq 0$ for any arm l when x is feasible for (8) at time t . Consider $k \in \mathcal{K}_x$ and assume that $n_{k,t} \leq \alpha_x \cdot \ln(t)$ and $n_{x,t} \geq \beta_x \cdot \ln(t)$. Suppose, by contradiction, that $|b_{x,t}(i) - n_{x,t} \cdot b(i)| \leq 32 \frac{(C+1)!^2}{\epsilon^4} \cdot \ln(t)$ for any resource $i \in \mathcal{C}_x$ and that $|\mu_l^c(i) - \tilde{c}_l(i)| \leq \frac{\epsilon^2}{8 \cdot (C+2)!}$ for any resource $i \in \{1, \dots, C\}$ for all arms $l \in \mathcal{K}_x$ such that $n_{l,t}^x \geq 2^7 \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(t)$, where $\tilde{c}_l(i)$ is the empirical average amount of resource i consumed when selecting basis x and pulling arm l , i.e. $\tilde{c}_l(i) = \frac{1}{n_{l,t}^x} \cdot \sum_{\tau=t_{\text{ini}}}^{t-1} I_{x_\tau=x} \cdot I_{a_\tau=l} \cdot c_{l,\tau}(i)$. We can partition $\mathcal{K}_x - \{k\}$ into two sets \mathcal{K}_x^1 and \mathcal{K}_x^2 such that $n_{l,t}^x \geq 2^7 \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(t)$ for all $l \in \mathcal{K}_x^1$ and $n_{l,t}^x < 2^7 \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(t)$ for all $l \in \mathcal{K}_x^2$. Similarly as in the proof of Lemma 14, we have, for any resource $i \in \mathcal{C}_x$:

$$\begin{aligned} 32 \frac{(C+1)!^2}{\epsilon^4} \cdot \ln(t) & \geq |b_{x,t}(i) - n_{x,t} \cdot b(i)| \\ & \geq n_{x,t} \cdot \left| \sum_{l \in \mathcal{K}_x^1} \tilde{c}_l(i) \cdot p_l - b(i) \right| - n_{k,t} - \sum_{l \in \mathcal{K}_x^2} n_{l,t}^x \\ & \geq n_{x,t} \cdot \left| \sum_{l \in \mathcal{K}_x^1} \tilde{c}_l(i) \cdot p_l - b(i) \right| - \alpha_x \cdot \ln(t) - C \cdot \frac{2^7 \cdot (C+2)!^2}{\epsilon^4} \cdot \ln(t), \end{aligned}$$

where $p_l = \frac{n_{l,t}^x}{n_{x,t}}$ for $l \in \mathcal{K}_x^1$ and $p_l = 0$ otherwise. Hence:

$$\begin{aligned} \left| \sum_{l \in \mathcal{K}_x} \mu_l^c(i) \cdot p_l - b(i) \right| & \leq \max_{l \in \mathcal{K}_x^1} |\mu_l^c(i) - \tilde{c}_l(i)| + \left| \sum_{l \in \mathcal{K}_x^1} \tilde{c}_l(i) \cdot p_l - b(i) \right| \\ & \leq \frac{\epsilon^2}{8 \cdot (C+2)!} + \frac{(\alpha_x + C \cdot \frac{2^7 \cdot (C+2)!^2}{\epsilon^4}) \cdot \ln(t)}{n_{x,t}} \\ & \leq \frac{\epsilon^2}{4 \cdot (C+2)!}. \end{aligned}$$

To derive the first inequality, we use the fact that $\sum_{l=1}^K p_l \leq 1$ and $p_l \geq 0$ for any arm l . For the last inequality, we use $n_{x,t} \geq 2^{10} \frac{(C+3)!^3}{\epsilon^8} \cdot \left(\frac{\lambda}{\Delta_x}\right)^2 \cdot \ln(t)$ along with $\lambda \geq 1$ and $\Delta_x \leq \text{obj}_{x^*} = \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \leq \sum_{k=1}^K \xi_k^{x^*} \leq 1$ because of the time constraint imposed in (3). We get:

$$\xi_k^x \leq \|p - \xi^x\|_2$$

$$\begin{aligned}
&= \|A_x^{-1} A_x (p - \xi^x)\|_2 \\
&\leq \|A_x^{-1}\|_2 \cdot \|A_x (p - \xi^x)\|_2 \\
&\leq \frac{1}{|\det(A_x)|} \cdot \|\text{adj}(A_x)\|_2 \cdot \sqrt{C} \cdot \frac{\epsilon^2}{4 \cdot (C+2)!} \\
&\leq \frac{1}{\epsilon} \cdot (C+1)! \cdot \sqrt{C} \cdot \frac{\epsilon^2}{4 \cdot (C+2)!} \\
&\leq \frac{\epsilon}{2},
\end{aligned}$$

a contradiction since x is ϵ -non-degenerate by Assumption 8. We conclude that:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*} < \text{obj}_x + 2E_{x,t}} \cdot I_{x \in \mathcal{B}_t} \cdot I_{n_{x,t} \geq \beta_x \cdot \ln(t)}\right] \\
&\leq C \cdot \sum_{t=t_{\text{ini}}}^T \sum_{i \in \mathcal{C}_x} \mathbb{P}[|b_{x,t}(i) - n_{x,t} \cdot b(i)| \geq 32 \frac{(C+1)!^2}{\epsilon^4} \cdot \ln(t)] \\
&+ C \cdot \sum_{t=t_{\text{ini}}}^T \sum_{\substack{l \in \mathcal{K}_x \\ i \in \mathcal{C}_x}} \mathbb{P}[|\mu_l^c(i) - \tilde{c}_l(i)| \geq \frac{\epsilon^2}{8 \cdot (C+2)!}; n_{l,t}^x \geq 2^7 \frac{(C+2)!^2}{\epsilon^4} \cdot \ln(t)] \\
&\leq 32C^2 \cdot \frac{(C+1)!^2}{\epsilon^4} \cdot \frac{\pi^2}{6} + C \cdot 2^9 \frac{(C+3)!^4}{\epsilon^6} \\
&+ \sum_{\substack{l \in \mathcal{K}_x \\ j \in \mathcal{C}_x}} \sum_{s=2^7 \cdot (C+2)!^2 / \epsilon^4 \cdot \ln(T)} \mathbb{P}[|\mu_l^c(j) - \tilde{c}_l(j)| \geq \frac{\epsilon^2}{8 \cdot (C+2)!}; n_{l,T}^x = s] \\
&\leq 4C \cdot 2^9 \frac{(C+3)!^4}{\epsilon^6} + 2^7 C^2 \cdot \frac{(C+2)!^2}{\epsilon^4} \cdot \frac{\pi^2}{6} \\
&\leq 2^{11} \frac{(C+4)!^4}{\epsilon^6},
\end{aligned}$$

where we use (12) for the second inequality and Lemma 1 for the third inequality.

Fact 13

$$\mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{x \in \mathcal{B}_t}\right] \leq 2^{10} \frac{(C+3)!^2}{\epsilon^6}.$$

Proof. First observe that:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{x \in \mathcal{B}_t}\right] &\leq \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{x \in \mathcal{B}_t} \cdot I_{|\det(\bar{A}_{x,t})| \geq \epsilon/2}\right] \\
&+ \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{|\det(\bar{A}_{x,t})| < \epsilon/2}\right] \\
&\leq \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{x \in \mathcal{B}_t} \cdot I_{|\det(\bar{A}_{x,t})| \geq \epsilon/2}\right] \\
&+ \sum_{t=t_{\text{ini}}}^T \sum_{\substack{k \in \mathcal{K}_x \\ i \in \mathcal{C}_x}} \mathbb{P}[|\bar{c}_{k,t}(i) - \mu_k^c(i)| > \frac{\epsilon^3}{16 \cdot (C+2)!^2}] \\
&\leq \mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{x \in \mathcal{B}_t} \cdot I_{|\det(\bar{A}_{x,t})| \geq \epsilon/2}\right] \\
&+ 2^9 \frac{(C+3)!^4}{\epsilon^6},
\end{aligned}$$

where we use the preliminary work of Section F.1 and in particular (38). The key observation now is that if $\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}$, $x \in \mathcal{B}_t$, and $|\det(\bar{A}_{x,t})| \geq \epsilon/2$, at least one of the following events $\{\bar{r}_{k,t} \geq \mu_k^r + \epsilon_{k,t}\}$, for $k \in \mathcal{K}_x$, or $\{|\bar{c}_{k,t}(i) - \mu_k^c(i)| \geq \epsilon_{k,t}\}$, for $k \in \mathcal{K}_x$ and $i \in \{1, \dots, C\}$, occurs. Otherwise we have:

$$\begin{aligned} \text{obj}_{x,t} - \text{obj}_x &= (\bar{A}_{x,t}^{-1} b_{\mathcal{C}_x})^\top \bar{r}_{\mathcal{K}_x,t} - (A_x^{-1} b_{\mathcal{C}_x})^\top \mu_{\mathcal{K}_x}^r \\ &< (\bar{A}_{x,t}^{-1} b_{\mathcal{C}_x})^\top (\mu_{\mathcal{K}_x}^r + \epsilon_{\mathcal{K}_x,t}) - (A_x^{-1} b_{\mathcal{C}_x})^\top \mu_{\mathcal{K}_x}^r \\ &= \frac{1}{\lambda} \cdot E_{x,t} + ((\bar{A}_{x,t}^{-1} - A_x^{-1}) b_{\mathcal{C}_x})^\top \mu_{\mathcal{K}_x}^r, \end{aligned}$$

where the first inequality is a consequence of the fact that x is feasible for (8), i.e. $\bar{A}_{x,t}^{-1} b_{\mathcal{C}_x} \geq 0$ with at least one non-zero coordinate since $\bar{c}_{k,t}(i) \geq \epsilon$ for all arms k and resource i by Assumption 8. Writing $\mu_k^c(i) = \bar{c}_{k,t}(i) + u_{i,k} \cdot \epsilon_{k,t}$, with $u_{i,k} \in [-1, 1]$ for all $(i, k) \in \mathcal{C}_x \times \mathcal{K}_x$, and defining the matrix $U = (u_{i,k})_{(i,k) \in \mathcal{C}_x \times \mathcal{K}_x}$, we get:

$$\begin{aligned} &\text{obj}_{x,t} - \text{obj}_x \\ &< \frac{E_{x,t}}{\lambda} + |(\bar{A}_{x,t}^{-1} - (\bar{A}_{x,t} + U \text{diag}(\epsilon_{\mathcal{K}_x,t}))^{-1}) b_{\mathcal{C}_x}|^\top \mu_{\mathcal{K}_x}^r \\ &= \frac{E_{x,t}}{\lambda} + |(\bar{A}_{x,t}^{-1} U (I + \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} U)^{-1} \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} b_{\mathcal{C}_x})^\top \mu_{\mathcal{K}_x}^r| \\ &\leq \frac{E_{x,t}}{\lambda} + \frac{1}{\epsilon} \cdot |\det(\bar{A}_{x,t} + U \text{diag}(\epsilon_{\mathcal{K}_x,t}))| \cdot |(\bar{A}_{x,t}^{-1} U (I + \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} U)^{-1} \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} b_{\mathcal{C}_x})^\top \mu_{\mathcal{K}_x}^r| \\ &= \frac{E_{x,t}}{\lambda} + \frac{1}{\epsilon} \cdot |(\text{adj}(\bar{A}_{x,t}) U \text{adj}(I + \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} U) \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} b_{\mathcal{C}_x})^\top \mu_{\mathcal{K}_x}^r| \\ &\leq \frac{E_{x,t}}{\lambda} + \frac{1}{\epsilon} \cdot \|\text{adj}(\bar{A}_{x,t}) U \text{adj}(I + \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} U)\|_2 \cdot \|\text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} b_{\mathcal{C}_x}\|_2 \cdot \|\mu_{\mathcal{K}_x}^r\|_2 \\ &\leq \frac{E_{x,t}}{\lambda} + \frac{1}{\epsilon} \cdot \|\text{adj}(\bar{A}_{x,t}) U\|_2 \cdot \|\text{adj}(I + \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} U)\|_2 \cdot \frac{E_{x,t}}{\lambda} \cdot \sqrt{C} \\ &\leq \frac{E_{x,t}}{\lambda} + \frac{1}{\epsilon} \cdot (C+1)! \cdot \|\text{adj}(I + \text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} U)\|_2 \cdot \frac{E_{x,t}}{\lambda} \\ &\leq \frac{E_{x,t}}{\lambda} + \frac{2}{\epsilon} \cdot (C+1)!^2 \cdot \frac{E_{x,t}}{\lambda} \\ &= E_{x,t}. \end{aligned}$$

We use the Woodbury matrix identity to derive the first equality and the matrix determinant lemma for the second equality. The second inequality is derived from Assumption 8 since $\epsilon \leq \det(A_x) = \det(\bar{A}_{x,t} + U \text{diag}(\epsilon_{\mathcal{K}_x,t}))$ by definition of U . The fourth inequality is derived from the observation that $\text{diag}(\epsilon_{\mathcal{K}_x,t}) \bar{A}_{x,t}^{-1} b_{\mathcal{C}_x}$ is the vector $(\epsilon_{k,t} \cdot \xi_{k,t}^x)_{k \in \mathcal{K}_x}$. The fifth inequality is obtained by observing that the components of $\bar{A}_{x,t}$ and U are all smaller than 1 in absolute value. The sixth inequality is obtained by observing that the elements of $\bar{A}_{x,t}^{-1}$ are smaller than $\frac{2}{\epsilon} \cdot (C-1)!$ since $\det(\bar{A}_{x,t}) \geq \frac{\epsilon}{2}$ and that $\epsilon_{k,t} \leq \frac{\epsilon}{2 \cdot C!}$ for all arms k as a result of the initialization phase. We get:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x,t} \geq \text{obj}_x + E_{x,t}} \cdot I_{x \in \mathcal{B}_t} \right] &\leq \sum_{t=t_{\text{ini}}}^T \sum_{\substack{k \in \mathcal{K}_x \\ i \in \mathcal{C}_x}} \mathbb{P}[|c_{k,t}(i) - \mu_k^c(i)| \geq \epsilon_{k,t}] \\ &\quad + \sum_{t=t_{\text{ini}}}^T \sum_{k \in \mathcal{K}_x} \mathbb{P}[r_{k,t} \geq \mu_k^r + \epsilon_{k,t}] \\ &\quad + 2^9 \frac{(C+3)!^4}{\epsilon^6} \\ &\leq 2 \cdot 2^9 \frac{(C+3)!^4}{\epsilon^6}. \end{aligned}$$

Fact 14

$$\mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*,t} \leq \text{obj}_{x^*} - E_{x,t}} \cdot I_{x^* \in \mathcal{B}_t}\right] \leq 2^{10} \frac{(C+3)!^2}{\epsilon^6}.$$

We omit the proof since it is almost identical to the proof of Fact 13.

F.5. Proof of Theorem 8. Along the same lines as for Theorem 5, we build upon (4):

$$\begin{aligned} R_{B(1), \dots, B(C-1), T} &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t, t}\right] + O(1) \\ &= T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t, t}\right] + \mathbb{E}\left[\sum_{t=\tau^*+1}^T r_{a_t, t}\right] + O(1) \\ &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t, t}\right] + \sigma \cdot \mathbb{E}\left[\min_{i=1, \dots, C} \sum_{t=\tau^*+1}^T c_{a_t, t}(i)\right] + O(1) \\ &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t, t}\right] + \sigma \cdot \sum_{i=1}^C \mathbb{E}\left[\left(\sum_{t=1}^{\tau^*} c_{a_t, t}(i) - B(i)\right)_+\right] + O(1). \end{aligned}$$

The second inequality is a direct consequence of Assumption 7. To derive the last inequality, observe that if $\tau^* = T + 1$, we have:

$$\sum_{t=\tau^*+1}^T c_{a_t, t}(j) = 0 \leq \sum_{i=1}^C \mathbb{E}\left[\left(\sum_{t=1}^T c_{a_t, t}(i) - B(i)\right)_+\right],$$

for any $j \in \{1, \dots, C\}$ while if $\tau^* < T + 1$ we have run out of resources before the end of the game, i.e. there exists $j \in \{1, \dots, C\}$ such that $\sum_{t=1}^{\tau^*} c_{a_t, t}(j) \geq B(j)$, which implies that:

$$\begin{aligned} \min_{i=1, \dots, C} \sum_{t=\tau^*+1}^T c_{a_t, t}(i) &\leq \sum_{t=\tau^*+1}^T c_{a_t, t}(j) \\ &\leq \sum_{t=\tau^*+1}^T c_{a_t, t}(j) + \sum_{t=1}^{\tau^*} c_{a_t, t}(j) - B(j) \\ &= \left(\sum_{t=1}^{\tau^*} c_{a_t, t}(j) - B(j)\right)_+ \\ &\leq \sum_{i=1}^C \left(\sum_{t=1}^{\tau^*} c_{a_t, t}(i) - B(i)\right)_+. \end{aligned}$$

Now observe that, for any resource $i \in \{1, \dots, C\}$:

$$\begin{aligned} \mathbb{E}\left[\left(\sum_{t=1}^T c_{a_t, t}(i) - B\right)_+\right] &\leq \mathbb{E}\left[\left(\sum_{t=t_{\text{ini}}}^T c_{a_t, t}(i) - b(i)\right)_+\right] + K \cdot 2^8 \frac{(C+2)!^4}{\epsilon^6} \cdot \ln(T) \\ &= \mathbb{E}\left[\left(\sum_{\substack{x \text{ basis for (3)}}} \{b_{x, T}(i) - n_{x, T} \cdot b(i)\}\right)_+\right] + O\left(K \cdot \frac{(C+2)!^4}{\epsilon^6} \cdot \ln(T)\right) \\ &\leq \sum_{x \in \mathcal{B}} \mathbb{E}\left[(b_{x, T}(i) - n_{x, T} \cdot b(i))_+\right] + \sum_{\substack{x \notin \mathcal{B}}} \mathbb{E}[n_{x, T}] \\ &+ \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x)=0}} \mathbb{E}[n_{x, T}] + O\left(K \cdot \frac{(C+2)!^4}{\epsilon^6} \cdot \ln(T)\right) \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{x \in \mathcal{B}} \int_0^T \mathbb{P}[b_{x,T}(i) - n_{x,T} \cdot b(i) \geq u] du + O\left(K \cdot \frac{(C+3)!^4}{\epsilon^6} \cdot \ln(T)\right) \\
&\leq \sum_{x \in \mathcal{B}} T \cdot \mathbb{P}[b_{x,T}(i) - n_{x,T} \cdot b(i) \geq 2^8 \frac{(C+3)!^4}{\epsilon^6}] \cdot \ln(T) \\
&\quad + 2^8 |\mathcal{B}| \cdot \frac{(C+3)!^4}{\epsilon^6} \cdot \ln(T) + O\left(K \cdot \frac{(C+3)!^4}{\epsilon^6} \cdot \ln(T)\right) \\
&= O\left(\frac{|\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right),
\end{aligned}$$

where we use the fact that the amounts of resources consumed at any time period are no larger than 1 for the first and second inequalities, Lemma 13 for the third inequality and inequalities (12) and (13) from Lemma 14 along with the fact that there are at least K feasible basis for (3) (corresponding to single-armed strategies) for the last equality. Plugging this back into the regret bound yields:

$$\begin{aligned}
R_{B(1), \dots, B(C-1), T} &\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \mathbb{E}\left[\sum_{t=\text{t}_{\text{ini}}}^T r_{a_t, t}\right] + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right) \\
&= T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[n_{k,T}^x] + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right) \\
&\leq T \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{x \in \mathcal{B}} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x,T}] + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right) \\
&= \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \cdot \left(T - \sum_{x \in \mathcal{B} \mid \Delta_x = 0} \mathbb{E}[n_{x,T}]\right) \tag{47} \\
&\quad - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x,T}] + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right) \\
&= \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \cdot \left(\text{t}_{\text{ini}} + \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \mathbb{E}[n_{x,T}] + \sum_{x \notin \mathcal{B}} \mathbb{E}[n_{x,T}] + \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x) = 0}} \mathbb{E}[n_{x,T}]\right) \\
&\quad - \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \left(\sum_{k=1}^K \mu_k^r \cdot \xi_k^x\right) \cdot \mathbb{E}[n_{x,T}] + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right) \\
&\leq \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \Delta_x \cdot \mathbb{E}[n_{x,T}] + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right) \tag{48} \\
&\leq 2^{10} \frac{(C+3)!^3 \cdot \lambda^2}{\epsilon^6} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x}\right) \cdot \ln(T) + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)!^4}{\epsilon^6} \cdot \ln(T)\right),
\end{aligned}$$

where we use (14) from Lemma 14 for the second inequality, Lemma 13 for the third inequality. To derive the last inequality, we use Lemma 15 and the fact that $\Delta_x \leq \text{obj}_{x^*} = \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} \leq \sum_{k=1}^K \xi_k^{x^*} \leq 1$.

F.6. Proof of Theorem 9. Along the same lines as for the proof of Theorem 6, we start from inequality (48) derived in the proof of Theorem 8 and apply Lemma 15 only if Δ_x is big enough, taking into account the fact that:

$$\sum_{x \in \mathcal{B}} \mathbb{E}[n_{x,T}] \leq T.$$

Specifically, we have:

$$R_{B(1), \dots, B(C-1), T}$$

$$\begin{aligned}
&= \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq T}} \left\{ \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \min(\Delta_x \cdot n_x, 2^{10} \frac{(C+3)^3 \cdot \lambda^2}{\epsilon^6} \cdot \frac{\ln(T)}{\Delta_x} + 2^{11} \frac{(C+4)^2}{\epsilon^6} \cdot \Delta_x) \right\} \\
&+ O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)^4}{\epsilon^6} \cdot \ln(T)\right) \\
&= \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq T}} \left\{ \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \min(\Delta_x \cdot n_x, 2^{10} \frac{(C+3)^3 \cdot \lambda^2}{\epsilon^6} \cdot \frac{\ln(T)}{\Delta_x}) \right\} + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)^4}{\epsilon^6} \cdot \ln(T)\right) \\
&\leq \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq T}} \left\{ \sum_{x \in \mathcal{B}} \sqrt{2^{10} \frac{(C+3)^3 \cdot \lambda^2}{\epsilon^6} \cdot \ln(T) \cdot n_x} \right\} + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)^4}{\epsilon^6} \cdot \ln(T)\right) \\
&\leq 2^5 \frac{(C+3)^2 \cdot \lambda}{\epsilon^3} \cdot \sqrt{\ln(T)} \cdot \sup_{\substack{(n_x)_{x \in \mathcal{B}} \geq 0 \\ \sum_{x \in \mathcal{B}} n_x \leq T}} \left\{ \sum_{x \in \mathcal{B}} \sqrt{n_x} \right\} + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)^4}{\epsilon^6} \cdot \ln(T)\right) \\
&\leq 2^5 \frac{(C+3)^2 \cdot \lambda}{\epsilon^3} \cdot \sqrt{\sigma \cdot |\mathcal{B}| \cdot T \cdot \ln(T)} + O\left(\frac{\sigma \cdot |\mathcal{B}| \cdot (C+3)^4}{\epsilon^6} \cdot \ln(T)\right),
\end{aligned}$$

where we use the fact that $\Delta_x \leq 1$ (see the proof of Theorem 8) for the second equality, we maximize over each $\Delta_x \geq 0$ to derive the first inequality, and we use Cauchy-Schwartz for the last inequality.

F.7. Proof of Theorem 10. Define $\tilde{b}(i) = B(i)/\tilde{T}$ for any $i \in \{1, \dots, C\}$. If the decision maker stops pulling arms at round \tilde{T} at the latest, all the results derived in Section 7 hold as long as we substitute T with \tilde{T} and we get:

$$\tilde{T} \cdot \tilde{\text{opt}} - \mathbb{E}\left[\sum_{t=1}^{\min(\tau^*, \tilde{T})} r_{a_t, t}\right] \leq X,$$

where X denotes the right-hand side of the regret bound derived in either Theorem 8 or Theorem 9 and $\tilde{\text{opt}}$ denotes the optimal value of (3) when $b(i)$ is substituted with $\tilde{b}(i)$ for any $i \in \{1, \dots, C\}$. The key observation is that $\tilde{T} \cdot \tilde{\text{opt}} = T \cdot \text{opt}$, where opt denotes the optimal value of (3), because the time constraint is redundant in (3) even when $b(i)$ is substituted with $\tilde{b}(i)$ for any $i \in \{1, \dots, C\}$. This is enough to show the claim as we get:

$$\begin{aligned}
X &\geq T \cdot \text{opt} - \mathbb{E}\left[\sum_{t=1}^{\tau^*} r_{a_t, t}\right] \\
&\geq R_{B(1), \dots, B(C-1), T},
\end{aligned}$$

where we use Lemma 2 for the last inequality.

F.8. Proof of Theorem 11. The only difference with the proofs of Theorems 8 and 9 lies in how we bound $\mathbb{E}[\sum_{t=\tau^*+1}^T r_{a_t, t}]$. We have:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=\tau^*+1}^T r_{a_t, t}\right] &\leq \mathbb{E}[(T - \tau^*)_+] \\
&= \sum_{t=0}^T \mathbb{P}[\tau^* \leq T - t] \\
&\leq \sum_{i=1}^C \sum_{t=0}^T \mathbb{P}\left[\sum_{\tau=1}^{T-t} c_{a_\tau, \tau}(i) \geq B(i)\right]
\end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^C \sum_{t=0}^T \mathbb{P}\left[\sum_{\tau=1}^{T-t} (c_{a_\tau, \tau}(i) - b(i)) \geq t \cdot b(i)\right] \\
&= \sum_{i=1}^C \sum_{t=0}^T \mathbb{P}\left[t_{\text{ini}} + \sum_{x \notin \mathcal{B}} n_{x,T} + \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x)=0}} n_{x,T} + \sum_{x \in \mathcal{B}} (b_{x,T-t}(i) - n_{x,T-t} \cdot b(i)) \geq t \cdot b(i)\right] \\
&= \sum_{i=1}^C \sum_{t=0}^T \mathbb{P}\left[t_{\text{ini}} + \sum_{x \notin \mathcal{B}} n_{x,T} + \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x)=0}} n_{x,T} \geq t \cdot \frac{b(i)}{2}\right] \\
&+ \sum_{i=1}^C \sum_{t=0}^T \sum_{x \in \mathcal{B}} \mathbb{P}\left[b_{x,T-t}(i) - n_{x,T-t} \cdot b(i) \geq t \cdot \frac{b(i)}{2 \cdot |\mathcal{B}|}\right] \\
&\leq 2 \sum_{i=1}^C \sum_{t=1}^T \frac{t_{\text{ini}} + \sum_{x \notin \mathcal{B}} \mathbb{E}[n_{x,T}] + \sum_{\substack{x \text{ pseudo-basis for (3)} \\ \text{with } \det(A_x)=0}} \mathbb{E}[n_{x,T}]}{t \cdot b(i)} \\
&+ 2|\mathcal{B}| \cdot \sum_{i=1}^C \sum_{t=1}^T \sum_{x \in \mathcal{B}} \frac{\mathbb{E}[|b_{x,T-t}(i) - n_{x,T-t} \cdot b(i)|]}{t \cdot b(i)} + O(1) \\
&= O\left(\frac{(C+4)!^4 \cdot |\mathcal{B}|^2}{b \cdot \epsilon^6} \cdot \ln^2(T)\right),
\end{aligned}$$

where we use Lemma 13 and we bound $\mathbb{E}[|b_{x,T-t}(i) - n_{x,T-t} \cdot b(i)|]$ in the same fashion as in the proof of Theorem 8 using Lemma 14.

Appendix G: Proofs for Section A.

G.1. Proof of Lemma 16. The proof follows the same steps as for Lemma 8. We use the shorthand notations $\beta_k = 8^{\rho \cdot (\sum_{i=1}^C b(i))^2} \cdot (\frac{1}{\Delta_k})^2$ and $n_{k,t}^{\neq x^*} = \sum_{x \in \mathcal{B} \mid k \in \mathcal{K}_x, x \neq x^*} n_{k,t}^x$. Along the same lines as in Fact 1, we have:

$$\mathbb{E}[n_{k,\tau^*}] \leq 2\beta_k \cdot \mathbb{E}[\ln(\tau^*)] + \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{x_t \neq x^*} \cdot I_{a_t=k} \cdot I_{n_{k,t}^{\neq x^*} \geq \beta_k \ln(t)}\right],$$

and we can focus on bounding the second term, which can be broken down as follows:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{x_t \neq x^*} \cdot I_{a_t=k} \cdot I_{n_{k,t}^{\neq x^*} \geq \beta_k \ln(t)}\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x_t,t} + E_{x_t,t} \geq \text{obj}_{x^*,t} + E_{x^*,t}} \cdot I_{x_t \neq x^*} \cdot I_{a_t=k} \cdot I_{n_{k,t}^{\neq x^*} \geq \beta_k \ln(t)}\right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x_t,t} \geq \text{obj}_{x^*,t} + E_{x_t,t}}\right] \\
&+ \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x^*,t} \leq \text{obj}_{x^*,t} - E_{x^*,t}}\right] \\
&+ \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x^*} < \text{obj}_{x_t} + 2E_{x_t,t}} \cdot I_{x_t \neq x^*} \cdot I_{a_t=k} \cdot I_{n_{k,t}^{\neq x^*} \geq \beta_k \ln(t)}\right].
\end{aligned}$$

We study each term separately, just like in Lemma 8.

Fact 15

$$\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x_t,t} \geq \text{obj}_{x^*,t} + E_{x_t,t}}\right] \leq K \cdot \frac{\pi^2}{6}.$$

Proof. If $\text{obj}_{x_t,t} \geq \text{obj}_{x_t} + E_{x_t,t}$, there must exist $l \in \mathcal{K}_{x_t}$ such that $\bar{r}_{l,t} \geq \mu_l^r + \epsilon_{l,t}$, otherwise:

$$\begin{aligned} \text{obj}_{x_t,t} - \text{obj}_{x_t} &= \sum_{l \in \mathcal{K}_{x_t}} (\bar{r}_{l,t} - \mu_l^r) \cdot \xi_l^{x_t} \\ &< \sum_{l \in \mathcal{K}_{x_t}} \epsilon_{l,t} \cdot \xi_l^{x_t} \\ &= E_{x_t,t}, \end{aligned}$$

where the inequality is strict because there must exist $l \in \mathcal{K}_{x_t}$ such that $\xi_l^{x_t} > 0$ (at least one resource constraint is binding for a feasible basis to (3) aside from the basis \hat{x} associated with $\mathcal{K}_{\hat{x}} = \emptyset$). We obtain:

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x_t,t} \geq \text{obj}_{x_t} + E_{x_t,t}}\right] &\leq \mathbb{E}\left[\sum_{t=1}^{\tau^*} \sum_{l=1}^K I_{\bar{r}_{l,t} \geq \mu_l^r + \epsilon_{l,t}}\right] \\ &\leq \sum_{l=1}^K \sum_{t=1}^{\infty} \mathbb{P}[\bar{r}_{l,t} \geq \mu_l^r + \epsilon_{l,t}] \\ &\leq K \cdot \frac{\pi^2}{6}, \end{aligned}$$

where the last inequality is derived along the same lines as in the proof of Fact 3.

Similarly, we can show that:

$$\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x^*,t} \leq \text{obj}_{x^*} - E_{x^*,t}}\right] \leq K \cdot \frac{\pi^2}{6}.$$

We move on to study the last term.

Fact 16

$$\mathbb{E}\left[\sum_{t=1}^{\tau^*} I_{\text{obj}_{x^*} < \text{obj}_{x_t} + 2E_{x_t,t}} \cdot I_{x_t \neq x^*} \cdot I_{a_t = k} \cdot I_{n_{k,t}^{\neq x^*} \geq \beta_k \ln(t)}\right] = 0.$$

Proof. If $\text{obj}_{x^*} < \text{obj}_{x_t} + 2E_{x_t,t}$, $x_t \neq x^*$, and $a_t = k$, we have:

$$\begin{aligned} \frac{\Delta_k}{2} &\leq \frac{\Delta_{x_t}}{2} \\ &< \sum_{l \in \mathcal{K}_{x_t}} \xi_l^{x_t} \cdot \sqrt{\frac{2 \ln(t)}{n_{l,t}}} \\ &\leq \sum_{l \in \mathcal{K}_{x_t}} \sqrt{\frac{2 \xi_l^{x_t} \cdot \xi_k^{x_t} \ln(t)}{n_{k,t}}}, \end{aligned}$$

where we use the fact that, by definition of the load balancing algorithm and since $a_t = k$, $\xi_k^{x_t} \neq 0$ (otherwise arm k would not have been selected) and:

$$n_{l,t} \geq \frac{\xi_l^{x_t}}{\xi_k^{x_t}} n_{k,t}, \quad (49)$$

for all arms $l \in \mathcal{K}_{x_t}$. We get:

$$\begin{aligned} n_{k,t} &< \frac{8}{(\Delta_k)^2} \cdot \xi_k^{x_t} \cdot \left(\sum_{l \in \mathcal{K}_{x_t}} \sqrt{\xi_l^{x_t}}\right)^2 \cdot \ln(t) \\ &\leq \frac{8}{(\Delta_k)^2} \cdot \xi_k^{x_t} \cdot \rho \cdot \sum_{l \in \mathcal{K}_{x_t}} \xi_l^{x_t} \cdot \ln(t) \\ &\leq \frac{8}{(\Delta_k)^2} \cdot \rho \cdot \left(\sum_{l \in \mathcal{K}_{x_t}} \xi_l^{x_t}\right)^2 \cdot \ln(t), \end{aligned}$$

using the Cauchy–Schwarz inequality and the fact that a basis involves at most ρ arms. Now observe that:

$$\begin{aligned} \sum_{l \in \mathcal{K}_{x_t}} \xi_l^{x_t} &\leq \sum_{l \in \mathcal{K}_{x_t}} \frac{\sum_{i=1}^C c_l(i)}{\epsilon} \cdot \xi_l^{x_t} \\ &\leq \frac{\sum_{i=1}^C b(i)}{\epsilon} \end{aligned}$$

as x_t is a feasible basis to (8) and using Assumption 2. We obtain:

$$\begin{aligned} n_{k,t}^{\neq x^*} &\leq n_{k,t} \\ &< 8 \cdot \frac{\rho \cdot (\sum_{i=1}^C b(i))^2}{\epsilon^2 \cdot (\Delta_k)^2} \cdot \ln(t) \\ &= \beta_k \cdot \ln(t). \end{aligned}$$

G.2. Proof of Lemma 17. We first show (16) by induction on t . The base case is straightforward. Suppose that the inequality holds at time $t-1$. There are three cases:

- arm k is not pulled at time $t-1$, in which case the left-hand side of the inequality remains unchanged while the right-hand side can only increase, hence the inequality still holds at time t ,
- arm k is pulled at time $t-1$ after selecting $x_{t-1} \neq x^*$, in which case both sides of the inequality increase by one and the inequality still holds at time t ,
- arm k is pulled at time $t-1$ after selecting $x_{t-1} = x^*$. First observe that there must exist $l \in \mathcal{K}_{x^*}$ such that $n_{l,t-1} \leq (t-1) \cdot \frac{\xi_l^{x^*}}{\sum_{r=1}^K \xi_r^{x^*}}$. Suppose otherwise, we have:

$$\begin{aligned} t-1 &= \sum_{l=1}^K n_{l,t} \\ &\geq \sum_{l \in \mathcal{K}_{x^*}} n_{l,t} \\ &> \sum_{l \in \mathcal{K}_{x^*}} (t-1) \cdot \frac{\xi_l^{x^*}}{\sum_{r=1}^K \xi_r^{x^*}} \\ &= t-1, \end{aligned}$$

a contradiction. Suppose now by contradiction that inequality (16) no longer holds at time t , we have:

$$\begin{aligned} n_{k,t-1} &= n_{k,t} - 1 \\ &> n_{x^*,t} \cdot \frac{\xi_k^{x^*}}{\sum_{l=1}^K \xi_l^{x^*}} + \sum_{x \in \mathcal{B}, x \neq x^*} n_{x,t} \\ &\geq (n_{x^*,t} + \sum_{x \in \mathcal{B}, x \neq x^*} n_{x,t}) \cdot \frac{\xi_k^{x^*}}{\sum_{l=1}^K \xi_l^{x^*}} \\ &= (t-1) \cdot \frac{\xi_k^{x^*}}{\sum_{l=1}^K \xi_l^{x^*}}, \end{aligned}$$

which implies, using the preliminary remark above, that $\frac{\xi_k^{x^*}}{n_{k,t-1}} < \max_{l \in \mathcal{K}_{x^*}} \frac{\xi_l^{x^*}}{n_{l,t-1}}$, a contradiction given the definition of the load balancing algorithm.

We conclude that inequality (16) holds for all times t and arms $k \in \mathcal{K}_{x^*}$. We also derive inequality (15) as a byproduct, since, at any time t and for any arm $k \in \mathcal{K}_{x^*}$:

$$\begin{aligned} n_{k,t} &\geq n_{x^*,t} - \sum_{l \in \mathcal{K}_{x^*}, l \neq k} n_{l,t} \\ &\geq n_{x^*,t} \cdot \left(1 - \frac{\sum_{l \in \mathcal{K}_{x^*}, l \neq k} \xi_l^{x^*}}{\sum_{l=1}^K \xi_l^{x^*}}\right) - \rho \cdot \left(\sum_{x \in \mathcal{B}, x \neq x^*} n_{x,t} + 1\right) \\ &= n_{x^*,t} \cdot \frac{\xi_k^{x^*}}{\sum_{l=1}^K \xi_l^{x^*}} - \rho \cdot \left(\sum_{x \in \mathcal{B}, x \neq x^*} n_{x,t} + 1\right), \end{aligned}$$

as a basis involves at most ρ arms.

G.3. Proof of Theorem 12. The proof proceeds along the same lines as for Theorem 3. We build upon (4):

$$\begin{aligned} R_{B(1), \dots, B(C)} &\leq B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{k=1}^K \mu_k^r \cdot \mathbb{E}[n_{k,\tau^*}] + O(1) \\ &\leq B \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} - \sum_{k \in \mathcal{K}_{x^*}} \mu_k^r \cdot \mathbb{E}[n_{k,\tau^*}] + O(1) \\ &\leq \left(B - \frac{\mathbb{E}[n_{x^*,\tau^*}]}{\sum_{k=1}^K \xi_k^{x^*}}\right) \cdot \sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*} + \rho^2 \cdot \sum_{x \in \mathcal{B}, x \neq x^*} \mathbb{E}[n_{x,\tau^*}] + O(1), \end{aligned}$$

where we use (15) to derive the third inequality. Now observe that, by definition, at least one resource is exhausted at time τ^* . Hence, there exists $i \in \{1, \dots, C\}$ such that the following holds almost surely:

$$\begin{aligned} B(i) &\leq \sum_{k=1}^K c_k(i) \cdot n_{k,\tau^*} \\ &\leq \sum_{k \notin \mathcal{K}_{x^*}} n_{k,\tau^*} + \sum_{k \in \mathcal{K}_{x^*}} c_k(i) \cdot n_{k,\tau^*} \\ &\leq \sum_{x \in \mathcal{B}, x \neq x^*} n_{x,\tau^*} + \sum_{k \in \mathcal{K}_{x^*}} c_k(i) \cdot n_{k,\tau^*} \\ &\leq \rho \cdot \left(\sum_{x \in \mathcal{B}, x \neq x^*} n_{x,\tau^*} + 2\right) + n_{x^*,\tau^*} \cdot \sum_{k \in \mathcal{K}_{x^*}} c_k(i) \cdot \frac{\xi_k^{x^*}}{\sum_{l=1}^K \xi_l^{x^*}} \\ &\leq \rho \cdot \left(\sum_{x \in \mathcal{B}, x \neq x^*} n_{x,\tau^*} + 2\right) + b(i) \cdot \frac{n_{x^*,\tau^*}}{\sum_{k=1}^K \xi_k^{x^*}}, \end{aligned}$$

where we use (16) and the fact that x^* is a feasible basis to (3). Rearranging yields:

$$\frac{n_{x^*,\tau^*}}{\sum_{k=1}^K \xi_k^{x^*}} \geq B - \frac{\rho}{b} \cdot \left(\sum_{x \in \mathcal{B}, x \neq x^*} n_{x,\tau^*} + 2\right),$$

almost surely. Plugging this last inequality back into the regret bound, we get:

$$\begin{aligned} R_{B(1), \dots, B(C)} &\leq \rho \cdot \sum_{x \in \mathcal{B}, x \neq x^*} \mathbb{E}[n_{x,t}] \cdot \left(\frac{\sum_{k=1}^K \mu_k^r \cdot \xi_k^{x^*}}{b} + \rho\right) + O(1) \\ &\leq \rho \cdot \sum_{x \in \mathcal{B}, x \neq x^*} \mathbb{E}[n_{x,t}] \cdot \left(\frac{\sum_{k=1}^K \sum_{i=1}^C c_k(i) \cdot \xi_k^{x^*}}{\epsilon \cdot b} + \rho\right) + O(1) \end{aligned}$$

$$\begin{aligned}
&\leq \left(\frac{\rho \cdot \sum_{i=1}^C b(i)}{\epsilon \cdot b} + (\rho)^2 \right) \cdot \sum_{x \in \mathcal{B}, x \neq x^*} \mathbb{E}[n_{x,t}] + O(1) \\
&= \left(\frac{\rho \cdot \sum_{i=1}^C b(i)}{\epsilon \cdot b} + (\rho)^2 \right) \cdot \sum_{k=1}^K \mathbb{E} \left[\sum_{x \in \mathcal{B} \mid k \in \mathcal{K}_x, x \neq x^*} n_{k,\tau^*}^x \right] + O(1) \\
&\leq 32 \frac{\rho^3 \cdot (\sum_{i=1}^C b(i))^3}{\epsilon^3 \cdot b} \cdot \left(\sum_{k=1}^K \frac{1}{(\Delta_k)^2} \right) \cdot \mathbb{E}[\ln(\tau^*)] + O(1) \\
&\leq 32 \frac{\rho^3 \cdot (\sum_{i=1}^C b(i))^3}{\epsilon^3 \cdot b} \cdot \left(\sum_{k=1}^K \frac{1}{(\Delta_k)^2} \right) \cdot \ln \left(\frac{\sum_{i=1}^C b(i) \cdot B}{\epsilon} + 1 \right) + O(1),
\end{aligned}$$

where we use the fact that x^* is a feasible basis to (3) for the third inequality, Lemma 16 for the fourth inequality, the concavity of the logarithmic function along with Lemma 7 for the last inequality.

G.4. Proof of Lemma 18. We use the shorthand notations $\beta_k = 8C \cdot (\frac{\lambda}{\Delta_k})^2$ and, for any round t , $n_{k,t}^{\notin \mathcal{O}} = \sum_{x \in \mathcal{B} \mid k \in \mathcal{K}_x, x \notin \mathcal{O}} n_{k,t}^x$. Similarly as in Fact 11, we have:

$$\begin{aligned}
\mathbb{E}[n_{k,T}^{\notin \mathcal{O}}] &\leq 2\beta_k \cdot \ln(T) + 2^9 \frac{(C+3)!^4}{\epsilon^6} \\
&\quad + \mathbb{E} \left[\sum_{t=t_{\text{ini}}}^T I_{x_t \notin \mathcal{O}} \cdot I_{a_t=k} \cdot I_{n_{k,t}^{\notin \mathcal{O}} \geq \beta_k \ln(t)} \cdot I_{x^* \in \mathcal{B}_t} \right],
\end{aligned}$$

and what remains to be done is to bound the second term, which we can break down as follows:

$$\begin{aligned}
&\mathbb{E} \left[\sum_{t=t_{\text{ini}}}^T I_{x_t \notin \mathcal{O}} \cdot I_{a_t=k} \cdot I_{n_{k,t}^{\notin \mathcal{O}} \geq \beta_k \ln(t)} \cdot I_{x^* \in \mathcal{B}_t} \right] \\
&\leq \mathbb{E} \left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x_t,t} + E_{x_t,t} \geq \text{obj}_{x^*,t} + E_{x^*,t}} \cdot I_{x_t \notin \mathcal{O}} \cdot I_{a_t=k} \cdot I_{n_{k,t}^{\notin \mathcal{O}} \geq \beta_k \ln(t)} \cdot I_{x^* \in \mathcal{B}_t} \right] \\
&\leq \mathbb{E} \left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x_t,t} \geq \text{obj}_{x^*,t} + E_{x_t,t}} \cdot I_{x_t \in \mathcal{B}_t} \right] \\
&\quad + \mathbb{E} \left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*,t} \leq \text{obj}_{x^*,t} - E_{x^*,t}} \cdot I_{x^* \in \mathcal{B}_t} \right] \\
&\quad + \mathbb{E} \left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*} < \text{obj}_{x_t} + 2E_{x_t,t}} \cdot I_{x_t \notin \mathcal{O}} \cdot I_{a_t=k} \cdot I_{n_{k,t}^{\notin \mathcal{O}} \geq \beta_k \ln(t)} \right].
\end{aligned}$$

The study of the second term is the same as in the proof of Lemma 15. We can also bound the first term in the same fashion as in the proof of Lemma 15 since there is no reference to the load balancing algorithm in the proof of Fact 13. The major difference with the proof of Lemma 15 lies in the study of the last term.

Fact 17

$$\mathbb{E} \left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x_t,t} \geq \text{obj}_{x_t} + E_{x_t,t}} \cdot I_{x \in \mathcal{B}_t} \right] \leq 2^{10} \frac{K \cdot (C+3)!^2}{\epsilon^6}.$$

Proof. The only difference with the proof of Fact 13 is that the number of arms that belong to \mathcal{K}_x for x ranging in $\{\tilde{x} \in \mathcal{B} \mid k \in \mathcal{K}_{\tilde{x}}, \tilde{x} \notin \mathcal{O}\}$ can be as big as K , as opposed to C when we are considering one basis at a time. This increases the bound by a multiplicative factor K .

Fact 18

$$\mathbb{E}\left[\sum_{t=t_{\text{ini}}}^T I_{\text{obj}_{x^*} < \text{obj}_{x_t} + 2E_{x_t,t}} \cdot I_{x_t \notin \mathcal{O}} \cdot I_{a_t = k} \cdot I_{n_{k,t}^{\notin \mathcal{O}} \geq \beta_k \ln(t)}\right] = 0.$$

Proof. Assume that $\text{obj}_{x^*} < \text{obj}_{x_t} + 2E_{x_t,t}$, $x_t \notin \mathcal{O}$, and $a_t = k$. We have:

$$\begin{aligned} \frac{\Delta_k}{2} &\leq \frac{\Delta_{x_t}}{2} \\ &< \lambda \cdot \sum_{l \in \mathcal{K}_{x_t}} \xi_{l,t}^{x_t} \cdot \sqrt{\frac{2 \ln(t)}{n_{l,t}}} \\ &\leq \lambda \cdot \sum_{l \in \mathcal{K}_{x_t}} \sqrt{\frac{2 \xi_{l,t}^{x_t} \cdot \xi_{k,t}^{x_t} \ln(t)}{n_{k,t}}}, \end{aligned}$$

where we use the fact that, by definition of the load balancing algorithm and since $a_t = k$, $\xi_{k,t}^{x_t} \neq 0$ (otherwise arm k would not have been selected) and:

$$n_{l,t} \geq \frac{\xi_{l,t}^{x_t}}{\xi_{k,t}^{x_t}} \cdot n_{k,t}, \quad (50)$$

for any arm $l \in \mathcal{K}_{x_t}$. We get:

$$\begin{aligned} n_{k,t}^{\notin \mathcal{O}} &\leq n_{k,t} \\ &< 8 \left(\frac{\lambda}{\Delta_k}\right)^2 \cdot \xi_{k,t}^{x_t} \cdot \left(\sum_{l \in \mathcal{K}_{x_t}} \sqrt{\xi_{l,t}^{x_t}}\right)^2 \cdot \ln(t) \\ &\leq 8 \left(\frac{\lambda}{\Delta_k}\right)^2 \cdot \xi_{k,t}^{x_t} \cdot C \cdot \sum_{l \in \mathcal{K}_{x_t}} \xi_{l,t}^{x_t} \cdot \ln(t) \\ &\leq 8 \left(\frac{\lambda}{\Delta_k}\right)^2 \cdot C \cdot \left(\sum_{l \in \mathcal{K}_{x_t}} \xi_{l,t}^{x_t}\right)^2 \cdot \ln(t) \\ &\leq 8C \cdot \left(\frac{\lambda}{\Delta_k}\right)^2 \cdot \ln(t), \end{aligned}$$

using the Cauchy–Schwarz inequality, the fact that a basis involves at most C arms, and the fact that x_t is feasible for (8) whose linear constraints include $\sum_{l=1}^K \xi_l \leq 1$ and $\xi_l \geq 0, \forall l \in \{1, \dots, K\}$. We get $n_{k,t}^{\notin \mathcal{O}} < \beta_k \ln(t)$ by definition of β_k .

G.5. Proof of Theorem 13. Substituting $b(i)$ with $B(i)/B(C)$ for every resource $i \in \{1, \dots, C\}$, the regret bound obtained in Theorem 3 turns into:

$$R_{B(1), \dots, B(C)} \leq 16 \frac{\rho}{\epsilon} \cdot \frac{\sum_{i=1}^C B(i)}{B(C)} \cdot \left(\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x} \right) \cdot \ln\left(\frac{\sum_{i=1}^C B(i)}{\epsilon} + 1\right) + O(1). \quad (51)$$

Observe that $\mathcal{B}^{\mathcal{C}}$ and $\mathcal{O}^{\mathcal{C}}$ are defined by strict inequalities that are linear in the vector $(B(1)/B(C), \dots, B(C-1)/B(C))$. Hence, for $B(C)$ large enough, $\mathcal{B}_{\infty}^{\mathcal{C}} \subset \mathcal{B}^{\mathcal{C}}$ and $\mathcal{O}_{\infty}^{\mathcal{C}} \subset \mathcal{O}^{\mathcal{C}}$ and thus $\mathcal{B} \subset \mathcal{B}_{\infty}$ and $\mathcal{O} \subset \mathcal{O}_{\infty}$. We now move on to prove each claim separately.

First claim. Suppose that there exists a unique optimal basis to (3), which we denote by x^* . Then, we must have $\mathcal{O} = \{x^*\} = \mathcal{O}_{\infty}$ for $B(C)$ large enough. Indeed, using the set inclusion relations shown above, we have $\mathcal{O} \subset \mathcal{O}_{\infty} = \{x^*\}$ and \mathcal{O} can never be empty as there exists at least one optimal basis to (3) (this linear program is feasible and bounded). We get $\mathcal{O}^{\mathcal{C}} \cap \mathcal{B} \subset \mathcal{O}_{\infty}^{\mathcal{C}} \cap \mathcal{B}_{\infty}$ for $B(C)$ large enough. Note moreover that,

for any $x \in \mathcal{B}$, Δ_x converges to Δ_x^∞ (because both the objective value of a feasible basis and the optimal value of a linear program are Lipschitz in the right-hand side of the inequality constraint), which implies that $\Delta_x > \frac{\Delta_x^\infty}{2} > 0$ when $x \in \mathcal{B} \cap \mathcal{O}^b$ for $B(C)$ large enough. We conclude with (51) that:

$$R_{B(1), \dots, B(C)} \leq 32 \frac{\rho}{\epsilon} \cdot \frac{\sum_{i=1}^C B(i)}{B(C)} \cdot \left(\sum_{x \in \mathcal{B}_\infty \mid \Delta_x^\infty > 0} \frac{1}{\Delta_x} \right) \cdot \ln \left(\frac{\sum_{i=1}^C B(i)}{\epsilon} + 1 \right) + O(1),$$

for $B(C)$ large enough. This yields the result since $B(i)/B(C) \rightarrow b(i) > 0$ for any resource $i = 1, \dots, C-1$.

Second claim. Suppose that $\frac{B(i)}{B(C)} - b(i) = O(\frac{\ln(B(C))}{B(C)})$ for any resource $i \in \{1, \dots, C-1\}$. Starting from (30) derived in the proof of Theorem 3 and applying Lemma 8 only if Δ_x is big enough, we have:

$$\begin{aligned} & R_{B(1), \dots, B(C)} \\ & \leq \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \min \left(\frac{\Delta_x}{\sum_{k=1}^K \xi_k^x} \cdot \mathbb{E}[n_{x, \tau^*}], 16\rho \cdot \frac{\sum_{k=1}^K \xi_k^x}{\Delta_x} \cdot \ln \left(\frac{\sum_{i=1}^C B(i)}{\epsilon} + 1 \right) + \frac{\pi^2}{3} \rho \cdot \frac{\Delta_x}{\sum_{k=1}^K \xi_k^x} \right) + O(1) \\ & \leq \sum_{x \in \mathcal{B} \mid \Delta_x > 0} \min \left(\Delta_x \cdot \frac{B(C)}{\min_{i=1, \dots, C} B(i)} \cdot \frac{\sum_{i=1}^C B(i)}{\epsilon}, 16 \frac{\rho \cdot \sum_{i=1}^C B(i)/B(C)}{\epsilon} \cdot \frac{1}{\Delta_x} \cdot \ln \left(\frac{\sum_{i=1}^C B(i)}{\epsilon} + 1 \right) \right) \\ & \quad + O(1) \\ & \leq 16 \frac{\rho \cdot \sum_{i=1}^C B(i)/B(C)}{\epsilon} \cdot \left(\sum_{x \in \mathcal{B} \cap \mathcal{O}^b \cap \mathcal{O}_\infty^b} \frac{1}{\Delta_x} \right) \cdot \ln \left(\frac{\sum_{i=1}^C B(i)}{\epsilon} + 1 \right) \\ & \quad + \left(\sum_{x \in \mathcal{B} \cap \mathcal{O}^b \cap \mathcal{O}_\infty} \Delta_x \right) \cdot \frac{B(C)}{\min_{i=1, \dots, C} B(i)} \cdot \frac{\sum_{i=1}^C B(i)}{\epsilon} + O(1), \end{aligned}$$

where we use:

$$\sum_{k=1}^K \xi_k^x \in \left[\min_{i=1, \dots, C} B(i)/B(C), \frac{\sum_{i=1}^C B(i)/B(C)}{\epsilon} \right]$$

and

$$\Delta_x \leq \frac{\sum_{i=1}^C B(i)/B(C)}{\epsilon},$$

as shown in the proof of Theorem 3 (substituting b with $\min_{i=1, \dots, C} B(i)/B(C)$). For $x \in \mathcal{B} \cap \mathcal{O}^b \cap \mathcal{O}_\infty^b$, we have $x \in \mathcal{B}_\infty$ and $\Delta_x > \frac{\Delta_x^\infty}{2} > 0$ for $B(C)$ large enough, as shown for the first claim. For $x \in \mathcal{B} \cap \mathcal{O}^b \cap \mathcal{O}_\infty$, we have $\Delta_x = O(\ln(B(C))/B(C))$ as both the objective value of a feasible basis and the optimal value of a linear program are Lipschitz in the right-hand side of the inequality constraints. We conclude that, for $B(C)$ large enough:

$$\begin{aligned} & R_{B(1), \dots, B(C)} \\ & \leq 32 \frac{\rho \cdot \sum_{i=1}^C B(i)/B(C)}{\epsilon} \cdot \left(\sum_{x \in \mathcal{B}_\infty \cap \mathcal{O}_\infty^b} \frac{1}{\Delta_x} \right) \cdot \ln \left(\frac{\sum_{i=1}^C B(i)}{\epsilon} + 1 \right) \\ & \quad + \frac{1}{\epsilon} \cdot \frac{B(C)}{\min_{i=1, \dots, C} B(i)} \cdot \sum_{x \in \mathcal{B} \cap \mathcal{O}^b \cap \mathcal{O}_\infty} O(\ln(B(C))) + O(1). \end{aligned}$$

This yields the result since $|\mathcal{B} \cap \mathcal{O}^b \cap \mathcal{O}_\infty| \leq |\mathcal{O}_\infty|$ and $B(i)/B(C) \rightarrow b(i) > 0$ for any resource $i = 1, \dots, C-1$.

G.6. Proof of Theorem 14. The proof is along the same lines as for Theorem 13. Specifically, in a first step, we observe that all the proofs of Section 6 remain valid (up to universal constant factors) for T large enough as long as we substitute b with B/T . Indeed, for T large enough, we have $\frac{B}{T} \leq 2$ and $|\mu_k^c - \frac{B}{T}| > \frac{\epsilon}{2}$ for all arms $k \in \{1, \dots, K\}$ under Assumption 6. In a second step, just like in the proof of Theorem 13, we show that we can substitute $\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x}$ with $\sum_{x \in \mathcal{B}_\infty \mid \Delta_x^\infty > 0} \frac{1}{\Delta_x^\infty}$ in the regret bound up to universal constant factors.

G.7. Proof of Theorem 15. The proof is along the same lines as for Theorem 13. Specifically, in a first step, we observe that all the proofs of Section 7 remain valid (up to universal constant factors) for T large enough as long as we substitute b with $\min_{i=1, \dots, C-1} B(i)/T$. Indeed, for T large enough, we have $\min_{i=1, \dots, C-1} B(i)/T \leq 2$ and, under Assumption 8, any basis to (17) has determinant larger than $\epsilon/2$ in absolute value and is $\epsilon/2$ -non-degenerate by continuity of linear functions. In a second step, just like in the proof of Theorem 13, we show that we can substitute $\sum_{x \in \mathcal{B} \mid \Delta_x > 0} \frac{1}{\Delta_x}$ with $\sum_{x \in \mathcal{B}_\infty \mid \Delta_x^\infty > 0} \frac{1}{\Delta_x^\infty}$ in the regret bound up to universal constant factors.

References

- [1] Agrawal S, Devanur NR (2014) Bandits with concave rewards and convex knapsacks. *Proc. 15th ACM Conf. Economics and Computation*, 989–1006.
- [2] Agrawal S, Devanur NR, Li L (2016) An efficient algorithm for contextual bandits with knapsacks, and an extension to concave objectives. *Proc. 29th Annual Conf. Learning Theory*, 4–18.
- [3] Agrawal S, Goyal N (2012) Analysis of thompson sampling for the multi-armed bandit problem. *Proc. 25th Annual Conf. Learning Theory*, volume 23.
- [4] Auer P, Cesa-Bianchi N, Fischer P (2002) Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3):235–256.
- [5] Auer P, Cesa-Bianchi N, Freund Y, Schapire RE (2002) The nonstochastic multiarmed bandit problem. *SIAM J. Comput.* 32(1):48–77.
- [6] Babaioff M, Dughmi S, Kleinberg R, Slivkins A (2012) Dynamic pricing with limited supply. *Proc. 13th ACM Conf. Electronic Commerce*, 74–91.
- [7] Badanidiyuru A, Kleinberg R, Singer Y (2012) Learning on a budget: posted price mechanisms for online procurement. *Proc. 13th ACM Conf. Electronic Commerce*, 128–145.
- [8] Badanidiyuru A, Kleinberg R, Slivkins A (2013) Bandits with knapsacks. *Proc. 54th IEEE Annual Symp. Foundations of Comput. Sci.*, 207–216.
- [9] Badanidiyuru A, Langford J, Slivkins A (2014) Resourceful contextual bandits. *Proc. 27th Annual Conf. Learning Theory*, volume 35, 1109–1134.
- [10] Bertsimas D, Tsitsiklis JN (1997) *Introduction to linear optimization*, volume 6 (Athena Scientific).
- [11] Besbes O, Zeevi A (2012) Blind network revenue management. *Oper. Res.* 60(6):1537–1550.
- [12] Combes R, Jian C, Srikant R (2015) Bandits with budgets: Regret lower bounds and optimal algorithms. *Proc. 2015 ACM SIGMETRICS Int. Conf. Measurement and Modeling Computer Systems*, 245–257.
- [13] Ding W, Qin T, Zhang X, Liu T (2013) Multi-armed bandit with budget constraint and variable costs. *Proc. 27th AAAI Conf. Artificial Intelligence*, 232–238.
- [14] Ghosh A, Rubinstein BIP, Vassilvitskii S, Zinkevich M (2009) Adaptive bidding for display advertising. *Proc. 18th Int. Conf. World Wide Web*, 251–260.
- [15] Gracová D, Jacko P (2014) Generalized restless bandits and the knapsack problem for perishable inventories. *Oper. Res.* 62(3):696–711.
- [16] Grötschel M, Lovász L, Schrijver A (2012) *Geometric algorithms and combinatorial optimization*, volume 2 (Springer).
- [17] Johnson K, Simchi-Levi D, Wang H (2015) Online network revenue management using thompson sampling. *Working Paper*.

-
- [18] Kleinberg R, Leighton T (2003) The value of knowing a demand curve: Bounds on regret for online posted-price auctions. *Proc. 44th IEEE Annual Symp. Foundations of Comput. Sci.*, 594–605.
- [19] Lai TL, Robbins H (1985) Asymptotically efficient adaptive allocation rules. *Adv. Applied Math.* 6(1):4–22.
- [20] Papadimitriou CH, Tsitsiklis JN (1999) The complexity of optimal queuing network control. *Math. Oper. Res.* 24(2):293–305.
- [21] Robbins H (1952) Some aspects of the sequential design of experiments. *Bulleting American Math. Society* 58(5):527–535.
- [22] Slivkins A (2013) Dynamic ad allocation: Bandits with budgets. *arXiv preprint arXiv:1306.0155*.
- [23] Tran-Thanh L, Chapman A, Rogers A, Jennings NR (2010) ϵ -first policies for budget-limited multi-armed bandits. *Proc. 24th AAAI Conf. Artificial Intelligence*, 1211–1216.
- [24] Tran-Thanh L, Chapman A, Rogers A, Jennings NR (2012) Knapsack based optimal policies for budget-limited multi-armed bandits. *Proc. 26th AAAI Conf. Artificial Intelligence*, 1134–1140.
- [25] Tran-Thanh L, Rogers A, Jennings NR (2012) Long-term information collection with energy harvesting wireless sensors: a multi-armed bandit based approach. *Autonomous Agents and Multi-Agent Systems* 25(2):352–394.
- [26] Tran-Thanh L, Stavrogiannis C, Naroditskiy V, Robu V, Jennings NR, Key P (2014) Efficient regret bounds for online bid optimisation in budget-limited sponsored search auctions. *Proc. 30th Conf. Uncertainty in Artificial Intelligence*, 809–818.
- [27] Trovo F, Paladino S, Restelli M, Gatti N (2016) Budgeted multi-armed bandit in continuous action space. *Working Paper*.
- [28] Weed J, Perchet V, Rigollet P (2016) Online learning in repeated auctions. *Proc. 29th Annual Conf. Learning Theory*, volume 49, 1562–1583.
- [29] Wu H, Srikant S, Liu X, Jiang C (2015) Algorithms with logarithmic or sublinear regret for constrained contextual bandits. *Adv. Neural Inform. Processing Systems*, 433–441.
- [30] Xia Y, Ding W, Zhang X, Yu N, Qin T (2015) Budgeted bandit problems with continuous random costs. *Proc. 7th Asian Conf. Machine Learning*, 317–332.