
Confidence-Budget Matching for Sequential Budgeted Learning

Yonathan Efroni^{* 1} Nadav Merlis^{* 2} Aadirupa Saha¹ Shie Mannor^{2 3}

Abstract

A core element in decision-making under uncertainty is the feedback on the quality of the performed actions. However, in many applications, such feedback is restricted. For example, in recommendation systems, repeatedly asking the user to provide feedback on the quality of recommendations will annoy them. In this work, we formalize decision-making problems with querying budget, where there is a (possibly time-dependent) hard limit on the number of reward queries allowed. Specifically, we consider multi-armed bandits, linear bandits, and reinforcement learning problems. We start by analyzing the performance of ‘greedy’ algorithms that query a reward whenever they can. We show that in fully stochastic settings, doing so performs surprisingly well, but in the presence of any adversity, this might lead to linear regret. To overcome this issue, we propose the Confidence-Budget Matching (CBM) principle that queries rewards when the confidence intervals are wider than the inverse square root of the available budget. We analyze the performance of CBM based algorithms in different settings and show that they perform well in the presence of adversity in the contexts, initial states, and budgets.

1. Introduction

In the past few decades, there have been great advances in the field of sequential decision making under uncertainty. From a practical perspective, recent algorithms achieve superhuman performance in problems that had been considered unsolvable (Mnih et al., 2015; Silver et al., 2017). From a theoretical perspective, algorithms with order-optimal performance were presented to various important settings (Garivier & Cappé, 2011; Azar et al., 2017, and others).

^{*}Equal contribution ¹Microsoft Research, New York ²Technion, Israel ³Nvidia Research, Israel. Correspondence to: Yonathan Efroni <jonathan.efroni@gmail.com>, Nadav Merlis <merlis.nadav@gmail.com>.

To solve such problems, most works share the same abstract interaction model. At each round, an agent (i) observes some information on the state of the environment, (ii) decides how to act, based on previous interactions, and, (iii) observes new feedback on the effect of its action. Finally, the environment changes its state based on the agent’s action, and the cycle begins anew. Much effort had been devoted to study specific instances of this abstract model, e.g., multi-armed bandits (MABs) (Auer et al., 2002; Garivier & Cappé, 2011; Kaufmann et al., 2012; Agrawal & Goyal, 2012), linear bandits (Dani et al., 2008; Abbasi-Yadkori et al., 2011; Agrawal & Goyal, 2013; Abeille et al., 2017) and reinforcement learning (RL) settings (Azar et al., 2017; Jin et al., 2018; Dann et al., 2019; Zanette & Brunskill, 2019; Efroni et al., 2019; Simchowitz & Jamieson, 2019; Tarbouriech et al., 2020; Cohen et al., 2020; Zhang et al., 2020). However, there are (still) several gaps between theory and practice that hinder the application of these models in real-world problems.

One such evident gap is the need to act under a budget constraint that limits the amount of feedback from the environment. That is, receiving feedback on the quality of the agent’s actions has an inherent cost. Consider, for example, an online recommendation system. There, asking for feedback from users negatively affects their experience, and feedback should be requested sparingly. Another example can be found in most large-scale RL domains, including autonomous driving. In many such cases, the reward should be labeled manually, and the resources for doing so are limited. Motivated by these problems, in this work, we aim to tackle the following question:

How should an agent trade-off exploration and exploitation when the feedback is limited by a budget?

In our efforts to answer this question, we study the effect of time-varying observation budget in various decision-making problems. Formally, we assume that at each round, the agent observes a non-decreasing, possibly adversarial, budget $B(t)$, which limits the number of queries for the reward of the problem. We first show that when the problem is stochastic and the budget is oblivious, greedily using any available budget leads to good performance. However, as soon as adversarial elements appear in the problem, or when the budget is controlled by an adaptive adversary, such an algorithm miserably fails. To tackle this problem, we suggest

a simple, generic, scheme, that only samples rewards for actions with high uncertainty, in comparison to the budget. We call such a mechanism *confidence-budget matching* (CBM). We show how to apply CBM to MAB, linear bandit and RL problems. In all cases, the mechanism can be applied in the presence of adaptive adversarial budgets. For linear bandits and RL, we show that CBM can be applied even when the contexts and initial states are adversarial. Finally, we present lower bounds for MABs and linear bandits, which show that CBM leads to order-optimal regret bounds.

2. Preliminaries

We start by defining a general model for sequential decision-making under uncertainty. Then, we will explain its realization in each individual model. In the most general model, at each round t , the environment supplies the agent with a context u_t that may either be stochastic or adversarially chosen. Then, the agent selects a policy $\pi_t \in \Pi(u_t)$ that can depend on u_t and past observations. Finally, the environment generates two stochastic feedback variables, from fixed distributions conditioned on u_t and π_t : feedback on the interaction with the environment Z_t and reward feedback R_t . In RL, for example, Z_t is the visited state-actions while R_t is their respective rewards. We also assume that there exists a reward function f such that the agent aims to maximize $f(R_t)$ throughout the interaction. Alternatively, algorithms aim to minimize its *pseudo-regret* (or regret), which is defined as

$$\text{Reg}(T) = \sum_{t=1}^T \left(\max_{\pi \in \Pi(u_t)} \mathbb{E}[f(R_t)|u_t, \pi] - \mathbb{E}[f(R_t)|u_t, \pi_t] \right).$$

Note that the pseudo-regret is random, as the policy depend on random feedback from the environment and contexts might be stochastic. Thus, regret bounds for different algorithms hold either with expectation or with high probability.

To illustrate the generality of this model, we explain how it encompasses both MAB, linear bandit and RL problems:

Contextual Multi-Armed Bandits (CMABs). At the beginning of each round, a context $u_t \in \{1, \dots, S\}$ is chosen, either stochastically or adversarially. Then, the agent chooses an action (arm) from a finite set of cardinality A , $\pi_t \triangleq a_t \in \mathcal{A}$ and the environment generates a reward $R_t \in [0, 1]$ with an expectation $\mathbb{E}[R_t|u_t = u, a_t = a] = r(u, a)$. An optimal arm is denoted by $a^*(u) \in \arg \max_a r(u, a)$ and its value by $r^*(u) = \max_a r(u, a)$. The reward function is $f(R_t) = R_t$ and there is no additional feedback ($Z_t = \phi$). A specific case of interest is where a single context exists, which is the well known MAB problem. Then, we denote $r(a) \triangleq r(1, a)$.

Linear Contextual Bandits. In the stochastic setting, u_t contains a set of A vectors in \mathbb{R}^d , generated independently

from a fixed distribution. In the adversarial case, u_t is an arbitrary set of vectors in \mathbb{R}^d . At each round t , the agent selects a single vector $\pi_t \triangleq x_t \in u_t$. Then, the environment generates a reward $R_t = \langle x_t, \theta \rangle + \eta_t$, where η_t is zero-mean subgaussian noise and $\theta \in \mathbb{R}^d$ is unknown. As in the CMAB problem, the reward function is $f(R_t) = R_t$ and there are no additional observations ($Z_t = \phi$).

Episodic Reinforcement Learning. Let \mathcal{S}, \mathcal{A} be finite state and action sets with cardinalities of S, A , respectively. Before each episode t , an initial state $s_{t,1}$ is generated either stochastically or adversarially (and serves as a context u_t). Then, an agent selects a nonstationary policy $\pi_t : \mathcal{S} \times [H] \rightarrow \mathcal{A}$, for some $H \in \mathbb{N}$. The policy is evaluated for H steps, and states are generated according to a transition kernel P ; namely, for any $s' \in \mathcal{S}$ and $h \in \{1, \dots, H\}$, $\Pr(s_{t,h+1} = s' | s_{t,h}, \pi_{t,h}) = P_h(s' | s_{t,h}, \pi_{t,h}(s_{t,h}))$. For brevity, we denote $a_{t,h} = \pi_{t,h}(s_{t,h})$. The agent observes the trajectory $Z_t = \{(s_{t,h}, a_{t,h})\}_{h=1}^H$ and for each visited state, a reward $R_t = \{R_{t,h}\}_{h=1}^H \in [0, 1]^H$ is generated such that $\mathbb{E}[R_{t,h} | s_{t,h} = s, a_{t,h} = a] = r(s, a)$. The reward function is then $f(R_t) = \sum_{h=1}^{H-1} R_{t,h}$.

Sequential Budgeted Learning. In most cases, it is natural to observe the effect of the policy on the environment; for example, it is reasonable to assume that the agent observes the visited states in RL, as it acts according to them. Thus, we assume that the agent always observes Z_t . On the other hand, many applications require specifically querying or labeling the reward. Then, oftentimes, such feedback is limited. Formally, let $\{B(t)\}_{t \geq 1}$ be a non-negative budget sequence that might be adversarially chosen. We also assume that the budget is non-decreasing, that is, a budget that is given cannot be taken. At each round t , the agent observes $B(t)$ and selects whether to query R_t or not, which we denote by $q_t = 1$ and $q_t = 0$, respectively. However, the agent can choose $q_t = 1$ only if its budget was not exhausted. Throughout most of the paper, we assume that querying a reward incurs unit cost. Then, an agent can select $q_t = 1$ only if $n_{t-1}^q \triangleq \sum_{k=1}^{t-1} 1\{q_k = 1\} \leq B(t) - 1$. In some cases, we extend the cost to be action-dependent. Then, a reward can only be queried if

$$B^q(t-1) \triangleq \sum_{k=1}^{t-1} c(\pi_k) 1\{q_k = 1\} \leq B(t) - c(\pi_t).$$

Notice that when queries have unit costs, then $n_t^q = B^q(t)$. For the RL setting, we give access to more refined feedback from specific time steps, to avoid confusion we only discuss it in Section 5.3. In all cases, we allow q_t to also depend on Z_t . Finally, and for ease of notations, we assume that the agent always observes $Y_t = R_t \cdot q_t$.

General Notations We let $\{F_t\}_{t \geq 0}$ be a filtration, where F_t is the σ -algebra that contains the random variables

$\{(u_k, \pi_k, Z_k, q_k, Y_k, B(k))_{k=0}^t, B(t+1), u_{t+1}\}$. In words, it contains the information on all *observed* rewards, actions, budget until the t^{th} episode, the budget at the $(t+1)^{\text{th}}$ episode, and the context at the $(t+1)^{\text{th}}$ episode. We denote $[n] = \{1, \dots, n\}$ for $n \in \mathbb{N}$ and also $x \vee 1 = \max\{x, 1\}$ for any $x \in \mathbb{R}$. We use $\mathcal{O}(X)$ and $\tilde{\mathcal{O}}(X)$ to refer to a quantity that depends on X up to constants and poly-log and constant expressions in problem parameters, respectively. Lastly, \lesssim, \gtrsim denote inequalities that hold up to poly-log and constant expressions in problem parameters.

3. Lower Bounds for Budgeted Problems

Before suggesting algorithms to the budgeted setting, it is of importance to understand how the new constraint affects the best-achievable regret. To this end, we study problem-independent lower bounds for budgeted MAB. By the end of the section, we also shortly discuss lower bounds for budgeted linear bandits. To derive the lower bounds, we require a more detailed description of the MAB model and additional notations. Moreover, we need to adapt the fundamental inequality of [Garivier et al. \(2019\)](#) to the case where the agent does not query all samples (Lemma 7). We refer the reader to Appendix A.1 for more details on the model and to Appendix A.2 for Lemma 7. Other proofs for this section can be found at Appendix A.3. Using Lemma 7, we can prove a lower bound for the following scenario in which (i) sampling an arm requires a unit cost, (ii) the budget constraint holds in expectation, and, (iii) the budget is given to the learner at the initial interaction, i.e. $\forall t \in [T], B(t) = B$:

Proposition 1. *Let T be some time horizon and let π be some bandit strategy such that for any bandit instance, it holds that $\mathbb{E}[n_T^q] \leq B$. Then, for $A \geq 2$, there exists a bandit instance for which*

$$\mathbb{E}[\text{Reg}(T)] \geq \frac{1}{140} \min \left\{ T \sqrt{\frac{A}{B}}, T \right\}.$$

As expected, when the budget is linear ($B = T$), we get the standard $\Omega(\sqrt{AT})$ lower bound. However, as we decrease the budget, the lower bound increases, up to the point of linear regret when the budget is not time-dependent. We also remark that the lower bound holds even if the budget constraint is only met *in expectation*. We will later present algorithms whose regret match these bounds, without *ever* violating the budget constraint. This implies that relaxing the budget requirement to hold in expectation cannot improve performance, from a worst-case perspective. Finally, note that when the budget is polynomial in T , e.g., $B = T^\beta$, we get a lower bound of $\Omega(\sqrt{AT}^{1-\beta})$. We will later prove upper bounds that match this rate.

Next, it is of interest to generalize the bound to the case of

arm-dependent costs. In this case, we require a more subtle analysis that also costs in a $\sqrt{\log A}$ factor:

Proposition 2. *Let T be the time horizon, and let $c(1), \dots, c(A) \geq 0$ be arm-dependent querying costs. Also, let π be some bandit strategy such that for any bandit instance, it holds that $\mathbb{E}[B^q(T)] \leq B$. Then, for $A \geq 2$, there exists a bandit instance for which*

$$\mathbb{E}[\text{Reg}(T)] \geq \frac{1}{140} \min \left\{ T \sqrt{\frac{\sum_{a=1}^A c(a)}{B(1 + \log A)}}, T \right\}.$$

While both bounds deal with fixed budget, $B(t) = B$ for all rounds, one can easily reduce them to lower bounds for time-dependent budgets, by reducing the lower bound only at a logarithmic factor. This is done by lower bounding the regret by the bound of the ‘worst-case’ time horizon $\Omega\left(\max_{t \in [T]} \left\{ \frac{t}{\sqrt{B(t)}} \right\}\right)$. We demonstrate how to do so for the case of arm-dependent costs in the following corollary:

Corollary 3. *Let $c(1), \dots, c(A) \geq 0$ be arm-dependent querying costs and let $B(1), \dots, B(T) > 0$ be an arbitrary non-decreasing budget sequence. Also, let π be some bandit strategy such that for any bandit instance and any time index $t \in [T]$, it holds that $\mathbb{E}[B^q(t)] \leq B(t)$. Then, for any $A \geq 2$, there exists a bandit instance for which*

$$\begin{aligned} \mathbb{E}[\text{Reg}(T)] \\ \geq \frac{1}{140(1 + \log T)} \sum_{t=1}^T \min \left\{ \sqrt{\frac{\sum_{a=1}^A c(a)}{B(t)(1 + \log A)}}, 1 \right\}. \end{aligned}$$

Proof. By Proposition 2, for any $t \in [T]$, there exists an instance such that

$$\mathbb{E}[\text{Reg}(t)] \geq \frac{t}{140} \min \left\{ \sqrt{\frac{\sum_{a=1}^A c(a)}{B(t)(1 + \log A)}}, 1 \right\}.$$

Let t_m be the time index in which the r.h.s. is maximized and fix the bandit problem to the corresponding instance that leads to its lower bound. Then,

$$\begin{aligned} \mathbb{E}[\text{Reg}(T)] &\geq \mathbb{E}[\text{Reg}(t_m)] \\ &= \max_{t \in [T]} \left\{ \frac{t}{140} \min \left\{ \sqrt{\frac{\sum_{a=1}^A c(a)}{B(t)(1 + \log A)}}, 1 \right\} \right\}. \end{aligned}$$

Finally, by Hölder’s inequality, if $x, y \in \mathbb{R}^T$ are such that $x_t, y_t \geq 0$ for all $t \in [T]$, then

$$\max_t x_t = \|x\|_\infty \geq \frac{\sum_{t=1}^T x_t y_t}{\|y\|_1} = \frac{\sum_{t=1}^T x_t y_t}{\sum_{t=1}^T y_t}.$$

Taking $x_t = \frac{t}{140} \min \left\{ \sqrt{\frac{\sum_{a=1}^A c(a)}{B(t)(1+\log A)}}, 1 \right\}$ and $y_t = \frac{1}{t}$ and recalling that $\sum_{t=1}^T \frac{1}{t} \leq 1 + \log T$ concludes the proof. \square

In the following sections, we derive regret upper bounds of similar budget-dependence, e.g., $\tilde{\mathcal{O}}\left(\sum_{t=1}^T \sqrt{A/B(t)}\right)$, if $c(a) = 1$ for all $a \in \mathcal{A}$. It is therefore of interest to observe the behavior of such bounds as a function of different budget profiles.

Example 1. (Budget Profiles and Consequences).

- **Linear Budget:** if $B(t) = \epsilon t$ for some $\epsilon > 0$, then $\text{Reg}(T) \leq 2\sqrt{AT}/\epsilon$. Specifically, if $\epsilon = \Omega(1)$, then we get the standard rates of $\text{Reg}(T) = \Theta(\sqrt{AT})$.
- **Polynomial Budget:** if $B(t) = t^c$ for some $c \in (0, 1]$, then the regret is also polynomial, i.e., $\text{Reg}(T) = \Theta(\sqrt{AT}^{1-c/2})$.
- **Fixed Budget:** if $B(t) = B_0 > 0$ is an initial budget, then $\text{Reg}(T) = \Theta(\sqrt{AT}/\sqrt{B_0})$. However, if the budget is given at the end of the game, namely, $B(t) = 0$ for $t \leq T - B_0$ and $B(t) = B_0$ for $t > T - B_0$, then $\text{Reg}(T) = \Omega(T)$ for any $B_0 = o(T)$.
- **Periodically-replenished budget** if the budget is replenished by $B_0 > 0$ every $N \in \mathbb{N}$ steps, namely, $B(t) = B_0 \cdot (1 + \lfloor \frac{t}{N} \rfloor)$, then

$$\text{Reg}(T) = \mathcal{O}\left(\sum_{s=1}^{\lfloor T/N \rfloor} \frac{\sqrt{AN}}{\sqrt{B_0 s}}\right) = \mathcal{O}\left(\sqrt{\frac{ATN}{B_0}}\right).$$

3.1. Lower Bounds for Linear Contextual Bandits

We end this section by presenting a lower bound for linear bandits. Here, we assume that the budget constraint is never violated (as we assume in the upper bounds). Then, for fixed budget and context space, we derive the following bound:

Proposition 4. Let $T \in \mathbb{N}$ be some time horizon and let π be a linear bandit policy such that $n_T^q \leq B$ a.s. for some fixed $B \leq T$. Then, there exists a d -dimensional linear bandit instance with arm set $[-1, 1]^d$ for which the expected regret of π is lower bounded by $\frac{dT}{80\sqrt{B}}$.

See Appendix A.4 for a proof. Importantly, this bound can be generalized to time-varying budgets, as in Corollary 3.

4. The Greedy Reduction: Gap Between Adversarial and Stochastic Contexts

We start by tackling the simpler case where the contexts are stochastic and the budget is oblivious. Formally, before the game starts, a sequence of budgets $\{B(t)\}_{t \geq 1}$ is chosen, possibly adversarially. Later, at the beginning of each

Algorithm 1 Greedy Reduction

- 1: **Require:** Algorithm \mathbb{A} , initial budget $B(1) \geq 1$
- 2: **Initialize:** $l = 0$
- 3: **for** $t = 1, \dots, T$ **do**
- 4: Observe context $u_t \sim \mathcal{P}_u$, and current budget $B(t)$
- 5: **if** $B(t) \geq B^q(t) + 1$ **then**
- 6: // Query reward feedback, act with \mathbb{A}
- 7: Advance $l \leftarrow l + 1$ and calculate $\pi_t \leftarrow \mathbb{A}_l(u_t)$
- 8: Act with $\pi_t, q_t = 1$; observe Z_t and R_t
- 9: **else**
- 10: // Don't query feedback, act with 'average' policy
- 11: Sample $j \sim \text{Uniform}(\{1, \dots, l\})$
- 12: Act with $\pi_t \leftarrow \mathbb{A}_j(u_t)$ and $q_t = 0$; ignore Z_t
- 13: **end if**
- 14: **end for**

round t , a context u_t is generated from a distribution \mathcal{P}_u , independently at random of other rounds. Then, the model continues as in Section 2. For this section, we also assume that queries have unit costs.

For this model, we suggest a simple greedy reduction (see Algorithm 1). Take an algorithm \mathbb{A} . If there is enough budget, query reward feedback and ask \mathbb{A} for a policy π_t to act with. Otherwise, when there is no available budget, pick uniformly at random a policy from past policies returned by \mathbb{A} , $\{\pi_t\}$, and act with it. We remark that $\mathbb{A}_k(u)$ denotes an output-policy of the algorithm at its k^{th} iteration, with u as the input context. Albeit simple, this algorithm performs surprisingly well, as we show in the following theorem:

Theorem 1 (Black Box Reduction for Stochastic Contexts). Let \mathbb{A} be an anytime algorithm with bounded regret $\mathbb{E}[\text{Reg}(T)] \leq \alpha T^\beta + C$ for some $\alpha, C \in \mathbb{R}_+, \beta \in [0, 1]$ and any $T \in \mathbb{N}$. Moreover, assume the budget is chosen by an oblivious adversary such that it is non-decreasing, $B(1) \geq 1$ and $B(t) \in \mathbb{N}$ for all $t \geq 1$. Then, the expected regret of Algorithm 1 with base algorithm \mathbb{A} and budget sequence $\{B(t)\}_{t \geq 1}$ is upper bounded by $\alpha T^\beta + C + \sum_{t=1}^T \frac{\alpha}{B^{1-\beta}(t)} + \frac{C}{B(t)}$.

The proof of the theorem (and all other results in the section) can be found at Appendix B. One possible application of the theorem is in the MAB setting, combined with MOSS-anytime (Degenne & Perchet, 2016). This would result in a regret bound of $\mathcal{O}\left(\sum_{t=1}^T \sqrt{A/B(t)}\right)$, which matches the lower bound of Corollary 3 up to log-factors. For linear bandits, using OFUL (Abbasi-Yadkori et al., 2011) as the base algorithm implies a regret of $\tilde{\mathcal{O}}\left(\sum_{t=1}^T d/\sqrt{B(t)}\right)$, which matches the lower bound of Proposition 4 up to log-factors. In general, we believe that this reduction is tight in many interesting settings. One possible intuitive explanation for this can be found at the following proposition. In it,

we prove that in non-contextual problems, for any fixed horizon, any general algorithm can be converted to one that uses the budget at the beginning of the game. A reasonable adaptation for time-varying budget and anytime algorithm would be to use the budget whenever possible.

Proposition 5. *Assume that the decision-making problem is non-contextual ($u_t = \phi, \forall t$) with no environment feedback ($Z_t = \phi, \forall t$) and unit-querying costs. Then, for any $T, B \in \mathbb{N}$ such that $B \leq T$ and any policy π under which $n_T^q \leq B$, there exists a policy π' such that $q_t = 1$ for all $t \in [B]$ (and zero otherwise) and $\mathbb{E}[\text{Reg}(T)|\pi'] = \mathbb{E}[\text{Reg}(T)|\pi]$.*

We end this section by returning to its basic assumptions - stochastic contexts and oblivious budget. We show that when at least one of these assumptions do not hold, the greedy reduction suffers a linear regret in a very simple CMAB problem, even if the budget is linear in expectation:

Proposition 6 (Greedy Reduction Degrades in the Presence of Adversary). *If an adaptive adversary controls either (i) the contexts, or (ii) the budget, then for any base algorithm \mathbb{A} used in Algorithm 1, there exists a contextual MAB problem with two contexts and two arms such that $\mathbb{E}[\text{Reg}(T)] \geq \frac{T}{4}$, even if $\mathbb{E}[B(t)] = \frac{t}{2}$ for all $t \in [T]$.*

Proof Sketch. Consider a contextual multi-armed bandit instance with two contexts, $u \in \{1, 2\}$. Assume that querying a reward feedback costs 1 for all contexts and all arms. Furthermore, assume the budget increases in each episode by 1 with probability 1/2.

If the adversary is adaptive to the history, it can choose $u = 1$ every round the budget increases and otherwise choose $u = 2$. The greedy reduction then only queries for feedback for $u = 1$. Thus, the regret for $u = 2$ is linear in T , since no information is gathered for this context, and the number of rounds $u = 2$ is $\Omega(T)$. Lastly, it can be shown that $\mathbb{E}[B(t)] = t/2$ in this construction. Equivalently, the same result holds if the contexts are uniformly distributed and an adaptive adversarial budget increases by a single unit only when $u = 1$. \square

This emphasizes the need for developing non-greedy algorithms that store budget to face adversities in the problem.

5. The Confidence-Budget Matching Principle

In the previous section, we showed that a simple greedy query rule performs well for sequential budgeted learning with stochastic contexts and oblivious budget. That is, querying for feedback as long as a spare budget exists results in a well-performing approach. However, this ‘greedy’ approach can miserably fail in the presence of adversarial contexts or budget. In this section, we introduce an alternative approach

Algorithm 2 Confidence-Budget Matching (CBM) Scheme

- 1: **Require:** Optimistic algorithm \mathbb{A} , $\{\alpha_t\}_{t \geq 1}$
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: Observe context u_t
 - 4: Act with π_t , acquired from $\mathbb{A}(F_{t-1})$ and observe Z_t
 - 5: Observe current budget $B(t)$
 - 6: **if** $CI_t(u_t, \pi_t) \geq \alpha_t \sqrt{1/B(t)}$ **then**
 - 7: Ask for feedback ($q_t = 1$) and observe R_t
 - 8: **end if**
 - 9: **end for**
-

we refer to as the Confidence-Budget Matching (CBM) principle. Unlike the greedy approach, CBM works well in the presence of adversities as it adequately preserves budget.

CBM is a generic algorithmic scheme that converts an un-budgeted optimistic algorithm to an algorithm that can be utilized in sequential budgeted learning. As evident in Algorithm 2, the agent follows a policy calculated by the baseline algorithm \mathbb{A} . Then, feedback on the reward of π_t is queried if the confidence interval (CI) of the policy, given current context, $CI_t(u_t, \pi_t)$ is larger than $\alpha \sqrt{1/B(t)}$ for some $\alpha > 0$. As querying rewards decreases the CI, $CI_t(u_t, \pi_t)$ will gradually decrease. Then, if a policy is chosen frequently enough, reward querying will stop once its confidence matches $\alpha \sqrt{1/B(t)}$.

Unlike the greedy reduction, the performance of CBM does not degrade in the presence of adversarial contexts or budget, as we demonstrate later in this section. A crucial reason for this is that CBM stops querying rewards of policies with small CI. This somewhat conservative behavior leads to a more robust algorithm. To better understand the robustness of this querying rule, we consider the MAB problem. For this problem, we set $\alpha_t \sim \sqrt{A}$, thus, for the MAB problem, CBM queries reward feedback if $CI_t(a_t) \geq \tilde{O}(\sqrt{A/B(t)})$. Denoting the number of queries from action a before the t^{th} episode by $n_{t-1}^q(a)$ and setting $CI_t(a_t) \sim 1/\sqrt{n_{t-1}^q(a_t)}$ (Hoeffding-based CI) leads to the following equivalent condition to CBM query rule for MAB: *ask for reward feedback if $n_{t-1}^q(a_t) \lesssim B(t)/A$. Namely, query for feedback if a_t was queried less than $B(t)/A$ times so far. Thus, this rule implicitly allocates $1/A$ of the current budget to each of the arms for possible use. This immediately implies the budget constraint is never violated, since there are A arms in total.*

Remark 1. *Notice that the CBM scheme plays actions selected by the optimistic baseline algorithm \mathbb{A} , which do not depend on the current budget $B(t)$. In particular, all our results also hold even if the budget is revealed after the agent selects an action, as depicted in Algorithm 2.*

Next, we study the performance of the CBM principle applied to MAB, linear bandits and RL problems. Importantly, we show that for all these settings, it matches the perfor-

mance of the greedy reduction for stochastic environments, while being able to face adversarial contexts and budgets.

Remark 2 (Sufficient Initial Budget). *For simplicity, we assume the initial budget $B(1)$ is large enough such that Algorithm 2 queries at the first round, that is $CI_1(u_1, \pi_1) \geq \alpha_1 \sqrt{1/B(1)}$. If this condition does not hold, an extra term of T_I should be added to the regret bounds where T_I is the first time in which $CI_{T_I}(u_{T_I}, \pi_{T_I}) \geq \alpha_{T_I} \sqrt{1/B(T_I)}$.*

5.1. Multi-Armed Bandits

We start by studying the performance of CBM for the MAB problem, where the base algorithm is UCB1 (Auer et al., 2002). We call the resulting algorithm CBM-UCB, which follows Algorithm 2 with $\alpha_t = 4\sqrt{6 \sum_a c(a) \log(At)}$. Although this setting is extremely simple, it highlights the central analysis technique, which is extended in the rest of this section to more challenging decision-making problems.

Theorem 2 (Confidence Budget Matching for Multi Armed Bandits). *For any querying costs $c(1), \dots, c(A) \geq 0$, any adaptive non-decreasing adversarially chosen sequence $\{B(t)\}_{t \geq 1}$ and for any $T \geq 1$, the expected regret of CBM-UCB is upper bounded by $\tilde{\mathcal{O}}\left(\sqrt{AT} + \sqrt{\sum_a c(a) \sum_{t=1}^T \mathbb{E}\left[\sqrt{\frac{1}{B(t)}}\right]}\right)$.*

Full description of the algorithm, alongside the proof of Theorem 2, is supplied at Appendix C. We now present a proof sketch that highlights how the CBM principle affects the regret bounds.

Proof Sketch. We use UCB bonus of $b_t^r(a) \triangleq \sqrt{\frac{3 \log(At)}{2n_{t-1}^q(a) \vee 1}}$, where $n_t^q(a)$ is the number of times arm a was queried up to round t ; namely, if $\bar{r}_t(a)$ is the empirical mean of a then, $UCB_t(a) = \bar{r}_{t-1}(a) + b_t^r(a)$, $LCB_t(a) = \bar{r}_{t-1}(a) - b_t^r(a)$ and $CI_t(a) = UCB_t(a) - LCB_t(a) = 2b_t^r(a)$.

Budget analysis. We start the proof by establishing that the budget constraint is never violated, $B^q(T) \leq B(T)$ for all $T \geq 1$. For simplicity, we do so for unit querying costs (where $B^q(t) = n^q(t)$). By the CBM condition, if $q_t = 1$, then $CI_t(a_t) \geq \alpha_t / \sqrt{B(t)}$. Then, for any $T \geq 1$

$$\begin{aligned} n^q(T) &= \sum_{t=1}^T 1\{q_t = 1\} \leq \sum_{t=1}^T \frac{CI_t(a_t)}{\alpha_t / \sqrt{B(t)}} 1\{q_t = 1\} \\ &\lesssim \sqrt{B(T)} \sum_{t=1}^T \frac{1}{\sqrt{n_{t-1}^q(a) \vee 1}} 1\{q_t = 1\}, \end{aligned}$$

where in the last relation we substituted all parameters and used the fact that the budget is non-decreasing. Importantly, notice that when the reward of an arm is queried, its count

increases, up to $n_t^q(a)$. Therefore, for any $T \geq 1$

$$n^q(T) \lesssim \sqrt{B(T)} \sum_{a=1}^A \sum_{i=0}^{n_T^q(a)} \frac{1}{\sqrt{i \vee 1}} \lesssim \sqrt{B(T)} \sqrt{n^q(T)}.$$

Reorganizing and choosing the right constants leads to the relation $n^q(T) \leq B(T)$, which deterministically holds. Importantly, this implies that CBM-UCB never tries to query reward without sufficient budget, so $q_t = 1$ if and only if the CBM condition holds, or, equivalently, $q_t = 0$ if and only if the CBM condition does not hold.

Regret analysis. Using standard concentration arguments, the expected regret $\mathbb{E}[\text{Reg}(T)]$ is bounded by

$$\sum_{t=1}^T \mathbb{E}[(UCB_t(a_t) - LCB_t(a_t))1\{q_t = 1\}] \quad (1)$$

$$+ \sum_{t=1}^T \mathbb{E}[(UCB_t(a_t) - LCB_t(a_t))1\{q_t = 0\}]. \quad (2)$$

For term (1), reward is always queried; therefore, the analysis closely follows standard analysis for UCB, which results with a bound of $\mathcal{O}\left(\sqrt{AT \log(AT)}\right)$. For (2), we know that reward was not queried, i.e., $q_t = 0$. Since $q_t = 0$ if and only if the CBM condition is not met, it implies that the CI is lower than the CBM-threshold, namely

$$\begin{aligned} &\sum_{t=1}^T \mathbb{E}[(UCB_t(a_t) - LCB_t(a_t))1\{q_t = 0\}] \\ &\lesssim \sum_{t=1}^T \mathbb{E}\left[\frac{\alpha_t}{\sqrt{B(t)}}\right] = \tilde{\mathcal{O}}\left(\sum_a c(a)\right) \sum_{t=1}^T \mathbb{E}\left[\sqrt{\frac{1}{B(t)}}\right]. \end{aligned}$$

Combining both bounds leads to the desired regret bound. \square

5.2. Linear Bandits

Next, we focus on applying the CBM principle, i.e., Algorithm 2, for linear bandits. The base algorithm that we rely on is OFUL (Abbasi-Yadkori et al., 2011), and we set $\alpha_t = \tilde{\mathcal{O}}(d)$ (see Appendix D for the full description of the algorithm). We call the resulting algorithm CBM-OFUL. Importantly, and in contrast to the greedy reduction of Section 4, we allow both the contexts and the budget to be chosen by an adaptive adversary. Nonetheless, CBM-OFUL still achieve the same performance as the greedy reduction Theorem 1, while not suffering of performance degradation in the presence of adaptive adversary (for a complete proof see Appendix D):

Theorem 3 (Confidence Budget Matching for Linear Bandits). *For any adaptive adversarially chosen sequence of non-decreasing budget and context sets*

$\{B(t), u_t\}_{t \geq 1}$ the regret of CBM-OFUL is upper bounded by $\tilde{O}\left(d\left(\sqrt{T} + \sum_{t=1}^T \frac{1}{\sqrt{B(t)}}\right)\right)$ for any $T \geq 1$ with probability greater than $1 - \delta$.

Notice that this matches the lower bound of Proposition 4. Notably, the examples of Proposition 6 can be represented as a linear bandit problem with $d = 4$. Thus, in contrast to the greedy reduction, which suffers linear regret, the regret of CBM-OFUL is $\tilde{O}(\sqrt{T})$.

5.3. Reinforcement Learning

In this section we apply the CBM principle to RL. For this setting, we relax the budget model presented in Section 2 and allow agents to query specific state action pairs along the trajectory observed at the t^{th} episode $\{(s_{t,h}, a_{t,h})\}_{h \in [H]}$. Namely, at the t^{th} episode, the agent acts with π_t , observes a trajectory $\{(s_{t,h}, a_{t,h})\}_{h \in [H]}$ and is allowed to query for reward feedback from any state-action pair along the trajectory. If the agent queries reward feedback in the t^{th} episode at the h^{th} time step it receives $R_{t,h}(s_{t,h}, a_{t,h})$. We denote this event as choosing $q_{t,h} = 1$. For simplicity, we work with unit-budget costs, i.e., the total budget used by the agent is $B^q(t) = \sum_{k=1}^t \sum_{h=1}^H 1\{q_{k,h} = 1\}$ and must be smaller than $B(t)$. Observe that in the standard RL setting, the reward budget is $B(t) = Ht$ for all $t \geq 1$.

Notably, querying reward feedback from specific time steps allows us to derive regret bounds that depend on the *sparsity* of the reward function. Formally, let \mathcal{L}_R be the set of tuples (s, a, h) with $r_h(s, a) \neq 0$. Then, for any $(s, a, h) \notin \mathcal{L}_R$, $r_h(s, a) = 0$, and since $R_t \in [0, 1]$, it also implies that $R_{t,h} \equiv 0$. Assume that the algorithm knows the cardinality of this set $|\mathcal{L}_R|$ (or an upper bound on $|\mathcal{L}_R|$). Leveraging this knowledge, we set the CBM feedback query rule in Algorithm 2, line 6, as follows,

$$\begin{aligned} &\text{Ask for reward feedback on } (s_{t,h}, a_{t,h}) \text{ if} \\ &CI_{t,h}^R(s_{t,h}, a_{t,h}) \gtrsim \sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + \frac{SAH}{B(t)} (q_{t,h} = 1), \end{aligned}$$

where $CI_{t,h}^R(s_{t,h}, a_{t,h})$ is the CI of the reward estimation of $s_{t,h}, a_{t,h}$ in the h^{th} time step at the t^{th} episode. Setting the reward bonus of the ‘optimistic’ model as in UCBVI-CH (Azar et al., 2017) leads to the following bound (see Appendix E for more details on the algorithm and proofs).

Theorem 4 (CBM-UCBVI). *For any adaptive adversarially chosen sequence of non-decreasing budget and initial state, $\{B(t), s_{t,1}\}_{t \geq 1}$, the regret of CBM-UCBVI is upper bounded by*

$$\tilde{O}\left(\sqrt{SAH^4 T} + H^3 S^2 A + \sum_{t=1}^T \sqrt{\frac{|\mathcal{L}_R| H^2}{B(t)}} + \frac{SAH^2}{B(t)}\right)$$

for any $T \geq 1$ with probability greater than $1 - \delta$.

Notice that the last term of the regret is dominated by its first term when $B(t) = \Omega(\sqrt{T})$ and the remaining budget-dependent term only scales with the sparsity-level of the reward $|\mathcal{L}_R|$. Notably, this implies that when $B(t) \sim t \lceil |\mathcal{L}_R| / SAH \rceil$, the third term is of the same order as the first term. Differently put, if the query budget $B(t)$ increases by a single unit every $SAH / |\mathcal{L}_R|$ episodes, the worst case performance of CBM-UCBVI remains the same, while reducing the amount of reward feedback.

While CBM-UCBVI clearly demonstrates the analysis techniques and insights from applying the CBM principle to RL, it is of interest to combine it with an algorithm with order-optimal regret bounds of $\sqrt{SAH^3 T}$ (e.g., (Jin et al., 2018)) when $B(t) = Ht$, that is, in the standard RL setting (notice that T is the number of *episodes* and not the total number of time steps). We achieve this goal by performing a more refined analysis that uses tighter concentration results based on (Azar et al., 2017; Dann et al., 2019; Zanette & Brunskill, 2019). Indeed, doing so leads to tighter regret bounds by a \sqrt{H} factor in the leading term (Full details on the algorithm and proofs can be found at Appendix F).

Theorem 5 (CBM-ULCVI). *For any adaptive adversarially chosen sequence of non-decreasing budget and initial state, $\{B(t), s_{t,1}\}_{t \geq 1}$, the regret of CBM-ULCVI is upper bounded by*

$$\tilde{O}\left(\sqrt{SAH^3 T} + H^3 S^2 A + \sum_{t=1}^T \sqrt{\frac{|\mathcal{L}_R| H^2}{B(t)}} + \frac{SAH^2}{B(t)}\right)$$

for any $T \geq 1$ with probability greater than $1 - \delta$.

This bound results in an interesting conclusion for general RL problems, i.e., when $|\mathcal{L}_R| = SAH$. Plugging this into Theorem 5, we observe that a budget of $B(t) = t$ – instead a budget of $B(t) = Ht$ as used in standard RL – results in order optimal regret bound. That is, it suffices for CBM-ULCVI to query reward feedback once per episode, without causing for performance degradation in a minimax sense.

5.4. General View on CBM for Optimistic Algorithms

The CBM principle queries for reward feedback (Algorithm 2, line 6) if the CI of the applied context-action is larger than a threshold, $CI_t(u_t, \pi_t) \geq \alpha_t \sqrt{1/B(t)}$, or more generally, if $CI_t(u_t, \pi_t) \geq \alpha_t f(B(t))$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$. A natural question arises: how to choose α_t and f ?

A useful rule of thumb to guide the choice of α_t and f is the following: if the regret of the optimistic algorithm \mathbb{A} is bounded by $\tilde{O}(\alpha T^\beta)$ then set $\alpha_t = \tilde{O}(\alpha)$ and $f(x) = \tilde{O}(x^{\beta-1})$. This matches the parameters chosen for both MAB and linear bandits. In RL, we relied on this rule but used a more complex function f , due to the application of an empirical Bernstein concentration argument (Maurer & Pontil, 2009).

The logic behind this choice is simple; it guarantees that the budget constraint is never violated, $B^q(T) \leq B(T)$ for all $T \geq 1$. Differently put, for any episode, reward feedback is not queried *if and only if* $CI_t(u_t, \pi_t) \leq \alpha_t f(B(t))$. This property can be proved via similar technique as in the proof sketch of Theorem 2 for CBM-MAB. An informal proof for the correctness of this statement for the general case goes as follows (for unit feedback-costs),

$$\begin{aligned} B^q(T) &\leq \sum_{t=1}^T 1\{q_t = 1\} CI_t(u_t, x_t) / (\alpha B(t)^{\beta-1}) \\ &\stackrel{(a)}{\leq} (B(T)^{1-\beta} / \alpha) \sum_{t=1}^T 1\{q_t = 1\} CI_t(u_t, x_t) \\ &\stackrel{(b)}{\lesssim} (B(T)^{1-\beta} / \alpha) \alpha B^q(T)^\beta, \end{aligned}$$

where (a) holds since the budget is non-decreasing, and (b) since $\sum_{t=1}^T 1\{q_t = 1\} CI_t(u_t, x_t) \sim \text{Reg}(B^q(T))$ for optimistic algorithms. Rearranging yields that $B^q(T)^{1-\beta} \lesssim B(T)^{1-\beta}$ which implies that $B^q(T) \leq B(T)$ by the monotonicity of $x^{1-\beta}$. Although the analysis for CBM in linear bandits and RL is more subtle, the intuition supplied by this informal reasoning is of importance; we believe it can serve as a starting point for future analysis of CBM-based algorithms in sequential budgeted learning.

6. Related Work

Multi-Armed Bandits with Paid Observations (Seldin et al., 2014). Closely related to our work is the framework of MAB with paid observations. There, an agent plays with an arm a_t and is allowed to query reward feedback on any subset of arms. Unlike in our case, there is no strict budget for observations, but, rather, each query comes at a cost that is subtracted from the reward. Notably, this requires translating the query costs to the same units as the reward, which is oftentimes infeasible. For example, in online recommendations, there is no clear way to quantify user dissatisfaction from feedback requests. In such cases, it is much more natural to enforce a (possibly time-varying) hard constraint on the number of feedback queries. Furthermore, the work of Seldin et al. (2014) focus on the MAB problem, whereas in this work, we focus on more involved contextual problems (i.e., linear bandits and RL). It is important to note that the analysis in (Seldin et al., 2014) holds for the adversarial reward model, whereas in this work, we focused on the stochastic reward model (with adversarial contexts and budget). We believe it is an interesting question what type of guarantees can be derived for the fully adversarial setting, i.e., when the rewards, budget and contexts are adversarially chosen. Finally, when applied to the stochastic case, the algorithm of Seldin et al. (2014) requires $B(T) = \Omega(T^{2/3})$. In contrast, our results hold for lower budgets, while achiev-

ing similar bounds when $B(T) = \Omega(T^{2/3})$.

MABs with Additional Observations (Yun et al., 2018).

In this closely related MAB setting, observing the reward of arms that were not played is possible, at a certain cost, as long as a non-decreasing budget constraint is not violated. Nonetheless, a key difference from our work is that Yun et al. (2018) assume that the reward of the played arm is *always* observed and does not consume any budget. Therefore, there is no clear way to apply their results to our setting.

Bandits with Knapsacks (BwK) (Badanidiyuru et al., 2013).

In the BwK model, a sampling budget is given prior to the game. At each round, the agent selects an arm and observes noisy samples of both the reward and the cost of the selected arm. That is, the agent always receives feedback on its actions. This comes in stark contrast to our model, where the budget restricts the amount of feedback an agent can obtain. Furthermore, in the BwK model, the game stops as soon as the cumulative cost exceeds the initial budget. In our model, where the budget serves as a constraint on the reward feedback, interaction continues even without an observation budget. When the budget is exhausted, the agent can still utilize its past information on the system to perform reasonably good actions. Notably, this forces the agent to sufficiently explore actions, even if they are costly, to identify high-rewarding ones.

We remark that there are additional extensions of the MAB setting in which arms incur costs (e.g., Sinha et al., 2021). There, the objective of an agent is to minimize a relaxed notion of cumulative regret and the cumulative cost. Unlike this work, we do not attribute cost to applying an action, but attribute a cost to *receiving feedback* on the reward.

RL with trajectory feedback (Efroni et al., 2020).

Under this model, instead of observing a reward for each played state-action, the agent only observes the cumulative rewards of each episode. This serves two reasons: first, and similarly to our work, it aims to reduce the feedback that the algorithm requires (by a factor of H), and when rewards are manually labeled, reduce the labeling load. Second, for many applications, it is much more natural to label the reward for a full trajectory than to each state-action. However, this approach comes at a noticeable cost, both in performance and computational complexity. In contrast, by sampling specific state-action pairs, our approach allows reducing the amount of feedback while maintaining similar performance and computational complexity. Nonetheless, we believe that when trajectory feedback is more natural, our approach can also be applied to further reduce the feedback for this setting. We leave such an extension for future work.

7. Summary and Discussion

In this work, we presented a novel framework for sequential decision-making under time-varying budget constraints. We analyzed what can and cannot be achieved by greedily using querying whenever possible. Then, we presented the CBM principle, which only queries rewards for actions with high uncertainty, compared to the current budget. We demonstrated how to apply the principle to MAB, linear bandits and RL problems and proved that it performs well also in the presence of adversities. We believe that this model can be adapted to many real-world problems and leaves room for interesting extensions, which we leave for future work.

Is there a value in knowing the future budget? Throughout this work, we assume the agent only observes the current budget $B(t)$ at the beginning of each round and does not have knowledge on future values of the budget $B(t')$ for $t' > t$. Intuitively, one expects that knowing the future budget would result in an improved and less conservative behavior in terms of budget allocation. Surprisingly, our matching lower and upper bounds for MAB (Corollary 3 and Theorem 2) and linear bandits (Proposition 4 and Theorem 3) show that this intuition does not always hold. Nonetheless, understanding if or when information on future budget is of value remains an interesting open question.

Monotonicity of the budget. Throughout this work, we assume that the budget never decreases. Intuitively, it implies that once a budget is allocated, it does not matter when the algorithm decides to use it. Nonetheless, for some problems, different assumptions are sometimes more relevant. A budget might be given alongside an ‘expiration date’ or might expire probabilistically. Another possible assumption is that the spare (unused) budget is bounded. Finally, in some instances, the total budget might be characterized by a specific random process, e.g., a biased random walk.

Problem-dependent bounds. Throughout this work, we focused on problem-independent regret bounds, that is, bounds that do not depend on the specific problem instance. Bounds that depend on specific instances usually focus on sufficiently sampling suboptimal arms, while implicitly assuming that optimal arms are sufficiently sampled (Auer et al., 2002). In contrast, when rewards are not always observed, algorithms must also control the number of queries from optimal arms. This becomes much harder in the presence of multiple optimal arms; in this case, an algorithm can never know if an arm is optimal or has a small suboptimal gap and might ‘waste’ budget while trying to discern which is true. In some sense, we believe that the CBM principle is well-suited for this setting, as it prevents the agent from exhausting all budget on specific arms.

Adaptivity to structure. In Section 5.3, we proved that when rewards are sparse, our algorithm can query rewards

according to the sparsity level, while maintaining the same regret bounds as the unbudgeted case. However, to do so, we required an upper bound on the sparsity of the problem. Therefore, a natural extension is to devise an algorithm that can adapt to an unknown sparsity level. Moreover, it is well known that structural assumptions can lead to improved regret bounds, and previous works proposed algorithms whose regret depends on nontrivial structural properties of the problem (Maillard et al., 2014; Zanette & Brunskill, 2019; Foster et al., 2019; 2020; Merlis & Mannor, 2019; 2020). Thus, it is interesting to understand what structural properties (beyond sparsity) affect the budgeted performance and how to design algorithms that adapt to such properties.

Acknowledgments

This work was partially funded by the Israel Science Foundation under ISF grant number 2199/20. YE is partially supported by the Viterbi scholarship, Technion. NM is partially supported by the Gutwirth Scholarship.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pp. 2312–2320, 2011.
- Abeille, M., Lazaric, A., et al. Linear thompson sampling revisited. *Electronic Journal of Statistics*, 11(2):5165–5197, 2017.
- Agrawal, S. and Goyal, N. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, pp. 39–1, 2012.
- Agrawal, S. and Goyal, N. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pp. 127–135, 2013.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Azar, M. G., Osband, I., and Munos, R. Minimax regret bounds for reinforcement learning. *arXiv preprint arXiv:1703.05449*, 2017.
- Badanidiyuru, A., Kleinberg, R., and Slivkins, A. Bandits with knapsacks. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 207–216. IEEE, 2013.
- Beygelzimer, A., Langford, J., Li, L., Reyzin, L., and Schapire, R. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Four-*

- teenth International Conference on Artificial Intelligence and Statistics*, pp. 19–26, 2011.
- Bretagnolle, J. and Huber, C. Estimation des densités: risque minimax. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 47(2):119–137, 1979.
- Cohen, A., Kaplan, H., Mansour, Y., and Rosenberg, A. Near-optimal regret bounds for stochastic shortest path. *arXiv preprint arXiv:2002.09869*, 2020.
- Dani, V., Hayes, T. P., and Kakade, S. M. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*, pp. 355–366, 2008.
- Dann, C., Li, L., Wei, W., and Brunskill, E. Policy certificates: Towards accountable reinforcement learning. In *International Conference on Machine Learning*, pp. 1507–1516, 2019.
- Degenne, R. and Perchet, V. Anytime optimal algorithms in stochastic multi-armed bandits. In *International Conference on Machine Learning*, pp. 1587–1595. PMLR, 2016.
- Efroni, Y., Merlis, N., Ghavamzadeh, M., and Mannor, S. Tight regret bounds for model-based reinforcement learning with greedy policies. In *Advances in Neural Information Processing Systems*, pp. 12224–12234, 2019.
- Efroni, Y., Merlis, N., and Mannor, S. Reinforcement learning with trajectory feedback. *arXiv preprint arXiv:2008.06036*, 2020.
- Foster, D. J., Krishnamurthy, A., and Luo, H. Model selection for contextual bandits. *arXiv preprint arXiv:1906.00531*, 2019.
- Foster, D. J., Rakhlin, A., Simchi-Levi, D., and Xu, Y. Instance-dependent complexity of contextual bandits and reinforcement learning: A disagreement-based perspective. *arXiv preprint arXiv:2010.03104*, 2020.
- Garivier, A. and Cappé, O. The kl-ucb algorithm for bounded stochastic bandits and beyond. In *Conference on Learning Theory*, pp. 359–376, 2011.
- Garivier, A., Ménard, P., and Stoltz, G. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 44(2):377–399, 2019.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pp. 4863–4873, 2018.
- Jin, T. and Luo, H. Learning adversarial mdps with bandit feedback and unknown transition. *arXiv preprint arXiv:1912.01192*, 2019.
- Kaufmann, E., Korda, N., and Munos, R. Thompson sampling: An asymptotically optimal finite-time analysis. In *International conference on algorithmic learning theory*, pp. 199–213. Springer, 2012.
- Lattimore, T. and Szepesvári, C. *Bandit algorithms*. Cambridge University Press, 2020.
- Maillard, O.-A., Mann, T. A., and Mannor, S. How hard is my mdp? the distribution-norm to the rescue”. *Advances in Neural Information Processing Systems*, 27: 1835–1843, 2014.
- Maurer, A. and Pontil, M. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- Merlis, N. and Mannor, S. Batch-size independent regret bounds for the combinatorial multi-armed bandit problem. In *Conference on Learning Theory*, pp. 2465–2489. PMLR, 2019.
- Merlis, N. and Mannor, S. Tight lower bounds for combinatorial multi-armed bandits. In *Conference on Learning Theory*, pp. 2830–2857. PMLR, 2020.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Seldin, Y., Bartlett, P., Crammer, K., and Abbasi-Yadkori, Y. Prediction with limited advice and multiarmed bandits with paid observations. In *International Conference on Machine Learning*, pp. 280–287. PMLR, 2014.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354, 2017.
- Simchowitz, M. and Jamieson, K. G. Non-asymptotic gap-dependent regret bounds for tabular mdps. In *Advances in Neural Information Processing Systems*, pp. 1153–1162, 2019.
- Sinha, D., Sankararaman, K. A., Kazerouni, A., and Avadhanula, V. Multi-armed bandits with cost subsidy. In *International Conference on Artificial Intelligence and Statistics*, pp. 3016–3024. PMLR, 2021.
- Tarbouriech, J., Garcelon, E., Valko, M., Pirota, M., and Lazaric, A. No-regret exploration in goal-oriented reinforcement learning. In *International Conference on Machine Learning*, pp. 9428–9437. PMLR, 2020.

Yun, D., Proutiere, A., Ahn, S., Shin, J., and Yi, Y. Multi-armed bandit with additional observations. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2(1):1–22, 2018.

Zanette, A. and Brunskill, E. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. *arXiv preprint arXiv:1901.00210*, 2019.

Zhang, Z., Ji, X., and Du, S. S. Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. *arXiv preprint arXiv:2009.13503*, 2020.

A. Lower bounds for Budgeted Problems

A.1. Detailed Decision-Making MAB Model and Relevant Notations

At each round t , the agent choose an arm a_t , which generates a reward R_t from a distribution ν_{a_t} , independently at random from other rounds. Throughout this section, we assume that all distributions ν_a are on $[0, 1]$. A bandit problem is characterized by its arm distributions, which we denote by $\underline{\nu} = \{\nu_a\}_{a=1}^A$. When we want to emphasize the arm distribution when taking an expectation, we denote it $\mathbb{E}_{\underline{\nu}}$. If the agent chooses to query this reward, we say that $q_t = 1$ and, otherwise, $q_t = 0$. Then, for ease of notations, we say that the agent observes $Y_t = R_t \cdot q_t$. A bandit strategy π maps all previous information and, possibly, internal randomization, into actions. Formally, let U_0, U_1, \dots be independent and identically distributed random variables with uniform distribution over $[0, 1]$. We denote the information known at time t by $I_t = (U_0, Y_1, U_1, \dots, Y_t, U_t)$, where $I_0 = U_0$. Then, a strategy maps the current information to actions, i.e., $\pi_t^a(I_t) = a_{t+1}$ and $\pi_t^q(I_t) = q_{t+1}$.

Next, we denote $n_t(a) = \sum_{k=1}^t 1\{a_k = a\}$, the number of times that an arm a was sampled up to time t when playing according to strategy π , and similarly let $n_t^q(a) = \sum_{k=1}^t 1\{a_k = a, q_k = 1\}$ be the total number of queries from arm a . Finally, let $\text{KL}(\cdot, \cdot)$ be the Kullback-Leibler between two probability measure, and for any $x, y \in [0, 1]$, we denote the KL divergence between Bernoulli random variables with expectations x, y by

$$\text{kl}(x, y) = x \log \frac{x}{y} + (1 - x) \log \frac{1 - x}{1 - y} .$$

A.2. Basic Inequalities

Lemma 7. For any $T \geq 1$, any $\sigma(I_T)$ measurable random variable Z with values in $[0, 1]$ and any two bandit problems $\underline{\nu}$ and $\underline{\nu}'$, it holds that

$$\sum_{a=1}^A \mathbb{E}_{\underline{\nu}}[n_T^q(a)] \text{KL}(\nu_a, \nu'_a) \geq \text{kl}(\mathbb{E}_{\underline{\nu}}[Z], \mathbb{E}_{\underline{\nu}'}[Z]) \quad (3)$$

Proof. The proof closely follows the one of Inequality (6) in (Garivier et al., 2019) and only differs by ignoring rounds where $q_t = 0$. Formally, let ν_{a_t, q_t} be the distribution of Y_t when playing a_t and querying the reward according to q_t . Specifically, if $q_t = 1$, then $\nu_{a_t, q_t} = \nu_{a_t}$, and if $q_t = 0$, it deterministically outputs $Y_t = 0$. Following the notations of (Garivier et al., 2019), for any two bandit instances $\underline{\nu}$ and $\underline{\nu}'$, we let $\mathbb{P}_{\underline{\nu}}$ and $\mathbb{P}_{\underline{\nu}'}$ be their associated probability measures, defined in some common measurable space (Ω, \mathcal{F}) (which exists by Kolmogorov's extension theorem). Then, for any $t \geq 0$ and any Borel sets $G \subset \mathbb{R}$ and $G' \subset [0, 1]$, our model implies that

$$\mathbb{P}_{\underline{\nu}}(Y_{t+1} \in G, U_{t+1} \in G' | I_t) = \nu_{\pi_t^a(I_t), \pi_t^q(I_t)}(B) \lambda(B')$$

where λ is the Lasbesgue measure on $[0, 1]$. Next, for any $t \geq 0$, let $\mathbb{P}_{\underline{\nu}}^{I_t}$ and $\mathbb{P}_{\underline{\nu}'}^{I_t}$ be the respective distributions of I_t w.r.t. $\underline{\nu}$ and $\underline{\nu}'$, and similarly use this notation for Y_t and U_t . Then, the previous relation can be written as

$$\mathbb{P}_{\underline{\nu}}^{(Y_{t+1}, U_{t+1}) | I_t} = \nu_{\pi_t^a(I_t), \pi_t^q(I_t)} \otimes \lambda,$$

where \otimes denotes the product of measures. Moreover, by the chain-rule for KL divergences, for any $t \geq 0$, we can write

$$\begin{aligned} \text{KL}\left(\mathbb{P}_{\underline{\nu}}^{I_{t+1}}, \mathbb{P}_{\underline{\nu}'}^{I_{t+1}}\right) &= \text{KL}\left(\mathbb{P}_{\underline{\nu}}^{(I_t, Y_{t+1}, U_{t+1})}, \mathbb{P}_{\underline{\nu}'}^{(I_t, Y_{t+1}, U_{t+1})}\right) \\ &= \text{KL}\left(\mathbb{P}_{\underline{\nu}}^{I_t}, \mathbb{P}_{\underline{\nu}'}^{I_t}\right) + \text{KL}\left(\mathbb{P}_{\underline{\nu}}^{(Y_{t+1}, U_{t+1}) | I_t}, \mathbb{P}_{\underline{\nu}'}^{(Y_{t+1}, U_{t+1}) | I_t}\right). \end{aligned} \quad (4)$$

Notably, the second term can be simplified to

$$\begin{aligned} \text{KL}\left(\mathbb{P}_{\underline{\nu}}^{(Y_{t+1}, U_{t+1}) | I_t}, \mathbb{P}_{\underline{\nu}'}^{(Y_{t+1}, U_{t+1}) | I_t}\right) &= \mathbb{E}_{\underline{\nu}}\left[\mathbb{E}_{\underline{\nu}}\left[\text{KL}\left(\nu_{\pi_t^a(I_t), \pi_t^q(I_t)} \otimes \lambda, \nu'_{\pi_t^a(I_t), \pi_t^q(I_t)} \otimes \lambda\right) | I_t\right]\right] \\ &= \mathbb{E}_{\underline{\nu}}\left[\mathbb{E}_{\underline{\nu}}\left[\text{KL}\left(\nu_{\pi_t^a(I_t), \pi_t^q(I_t)}, \nu'_{\pi_t^a(I_t), \pi_t^q(I_t)}\right) | I_t\right]\right] \\ &\stackrel{(*)}{=} \mathbb{E}_{\underline{\nu}}\left[\mathbb{E}_{\underline{\nu}}\left[1\{\pi_{t+1}^q(I_t) = 1\} \text{KL}\left(\nu_{\pi_t^a(I_t)}, \nu'_{\pi_t^a(I_t)}\right) | I_t\right]\right] \\ &= \mathbb{E}_{\underline{\nu}}\left[1\{\pi_{t+1}^q(I_t) = 1\} \sum_{a=1}^A \text{KL}(\nu_a, \nu'_a) 1\{\pi_{t+1}^a(I_t) = a\}\right] \\ &= \mathbb{E}_{\underline{\nu}}\left[\sum_{a=1}^A \text{KL}(\nu_a, \nu'_a) 1\{a_{t+1} = a, q_{t+1} = 1\}\right] \end{aligned}$$

where in (*) we used the fact that if $q_{t+1} = \pi_{t+1}^q(I_t) = 0$, then both distributions deterministically output 0, so the KL divergence is also 0. For the last relation, also recall that $\pi_t^a(I_t) = a_{t+1}$ and $\pi_{t+1}^q(I_t) = q_{t+1}$. Substituting back into Equation (4) and applying the same argument recursively, we get

$$\begin{aligned} \text{KL}\left(\mathbb{P}_{\underline{\nu}}^{I_T}, \mathbb{P}_{\underline{\nu}'}^{I_T}\right) &= \sum_{t=1}^T \mathbb{E}_{\underline{\nu}}\left[\sum_{a=1}^A \text{KL}(\nu_a, \nu'_a) 1\{a_t = a, q_t = 1\}\right] \\ &= \sum_{a=1}^A \mathbb{E}_{\underline{\nu}}[n_T^q(a)] \text{KL}(\nu_a, \nu'_a). \end{aligned}$$

Notice that the recursion stops at $\text{KL}\left(\mathbb{P}_{\underline{\nu}}^{I_0}, \mathbb{P}_{\underline{\nu}'}^{I_0}\right) = \text{KL}\left(\mathbb{P}_{\underline{\nu}}^{U_0}, \mathbb{P}_{\underline{\nu}'}^{U_0}\right) = \text{KL}(\lambda, \lambda) = 0$. To conclude the proof, we apply Lemma 1 of (Garivier et al., 2019), which implies that for any $\sigma(I_T)$ measurable random variable Z over $[0, 1]$, it holds that

$$\text{KL}\left(\mathbb{P}_{\underline{\nu}}^{I_T}, \mathbb{P}_{\underline{\nu}'}^{I_T}\right) \geq \text{kl}(\mathbb{E}_{\underline{\nu}}[Z], \mathbb{E}_{\underline{\nu}'}[Z]).$$

□

In the following, we present a simplified version for the bound of Lemma 7, that will be of use in our proofs:

Lemma 8. *Let $\underline{\nu}$ be a bandit instance and let $k \in [A]$ be some arm. Furthermore, let $\underline{\nu}'$ be a bandit instance that differs from $\underline{\nu}$ only at arm k , i.e., $\nu'_a = \nu_a$ for all $a \neq k$. Then,*

$$\mathbb{E}_{\underline{\nu}}[n_T(k)] - T\sqrt{\frac{1}{2}\mathbb{E}_{\underline{\nu}}[n_T^q(k)]\text{KL}(\nu_k, \nu'_k)} \leq \mathbb{E}_{\underline{\nu}'}[n_T(k)] \leq \mathbb{E}_{\underline{\nu}}[n_T(k)] + T\sqrt{\frac{1}{2}\mathbb{E}_{\underline{\nu}}[n_T^q(k)]\text{KL}(\nu_k, \nu'_k)} .$$

Moreover, if ν_k and ν'_k are Bernoulli distributions with parameters $\frac{1}{2}$ and $\frac{1}{2} + \epsilon$, respectively, for $\epsilon \in (0, \frac{1}{4})$, then

$$\mathbb{E}_{\underline{\nu}}[n_T(k)] - 2\epsilon T\sqrt{\log \frac{4}{3}\sqrt{\mathbb{E}_{\underline{\nu}}[n_T^q(k)]}} \leq \mathbb{E}_{\underline{\nu}'}[n_T(k)] \leq \mathbb{E}_{\underline{\nu}}[n_T(k)] + 2\epsilon T\sqrt{\log \frac{4}{3}\sqrt{\mathbb{E}_{\underline{\nu}}[n_T^q(k)]}} .$$

Proof. Our proof closely follows the proof of Theorem 6 in (Garivier et al., 2019), with small modifications due to the budget. We start by applying Lemma 7 on $\underline{\nu}$ and $\underline{\nu}'$ with $Z = \frac{n_T(k)}{T}$, while noticing that $\text{KL}(\nu_a, \nu'_a) = 0$ for all $a \neq k$:

$$\mathbb{E}_{\underline{\nu}}[n_T^q(k)]\text{KL}(\nu_k, \nu'_k) \geq \text{kl}\left(\frac{\mathbb{E}_{\underline{\nu}}[n_T(k)]}{T}, \frac{\mathbb{E}_{\underline{\nu}'}[n_T(k)]}{T}\right) \geq 2\left(\frac{\mathbb{E}_{\underline{\nu}}[n_T(k)]}{T} - \frac{\mathbb{E}_{\underline{\nu}'}[n_T(k)]}{T}\right)^2 ,$$

where the last inequality is by Pinsker's inequality. Alternatively, we can write

$$|\mathbb{E}_{\underline{\nu}'}[n_T(k)] - \mathbb{E}_{\underline{\nu}}[n_T(k)]| \leq T\sqrt{\frac{1}{2}\mathbb{E}_{\underline{\nu}}[n_T^q(k)]\text{KL}(\nu_k, \nu'_k)} ,$$

which leads to the first result of the lemma. For the second result, we directly upper bound $\text{KL}(\nu_k, \nu'_k)$ for any $\epsilon \in (0, \frac{1}{4})$ by:

$$\text{KL}(\nu_k, \nu'_k) = \text{kl}\left(\frac{1}{2}, \frac{1}{2} + \epsilon\right) = \frac{1}{2} \log \frac{1}{1 - 4\epsilon^2} \leq 8\epsilon^2 \log \frac{4}{3} ,$$

where the last inequality holds since $\log \frac{1}{1-u} \leq 4u \log \frac{4}{3}$ for any $u \in (0, \frac{1}{4})$. Substituting back to the first result of the lemma concludes the proof. \square

A.3. Proofs for lower bounds

Proposition 1. *Let T be some time horizon and let π be some bandit strategy such that for any bandit instance, it holds that $\mathbb{E}[n_T^q] \leq B$. Then, for $A \geq 2$, there exists a bandit instance for which*

$$\mathbb{E}[\text{Reg}(T)] \geq \frac{1}{140} \min\left\{T\sqrt{\frac{A}{B}}, T\right\} .$$

Proof. Let $\underline{\nu}$ be an arm distribution such that all arms are Bernoulli-distributed with parameter $\frac{1}{2}$. By the pigeonhole principle, one can prove that for any $A \geq 2$, there must exist some arm t such that both $\mathbb{E}_{\underline{\nu}}[n_T(k)] \leq \frac{5T}{4A}$ and $\mathbb{E}_{\underline{\nu}}[n_T^q(k)] \leq \frac{15B}{A}$ (see prove in Lemma 9 for $\alpha = \frac{5}{4}$ and $\beta = 15$).

Then, we define a new bandit instance $\underline{\nu}'$ such that $\nu'_a = \nu_a$ for all $a \neq k$ and ν'_k is Bernoulli-distributed with parameter $\frac{1}{2} + \epsilon$, for some $\epsilon \in (0, \frac{1}{4})$ that will be determined later. Specifically, it implies that the regret of π on $\underline{\nu}'$ is

$$\mathbb{E}_{\underline{\nu}'}[\text{Reg}(T)] = \sum_{a \neq k} \epsilon \mathbb{E}_{\underline{\nu}'}[n_T(a)] = \epsilon T \left(1 - \frac{\mathbb{E}_{\underline{\nu}'}[n_T(k)]}{T}\right) . \quad (5)$$

Thus, to lower bound the regret, we need to upper bound $\mathbb{E}_{\underline{\nu}'}[n_T(k)]$. By Lemma 8, we have

$$\mathbb{E}_{\underline{\nu}'}[n_T(k)] \leq \mathbb{E}_{\underline{\nu}}[n_T(k)] + 2\epsilon T\sqrt{\log \frac{4}{3}\sqrt{\mathbb{E}_{\underline{\nu}}[n_T^q(k)]}} ,$$

and recalling that $\mathbb{E}_{\underline{\nu}}[n_T(k)] \leq \frac{5T}{4A}$ and $\mathbb{E}_{\underline{\nu}}[n_T^q(k)] \leq \frac{15B}{A}$, we get

$$\mathbb{E}_{\underline{\nu}'}[n_T(k)] \leq \frac{5T}{4A} + 2\epsilon T \sqrt{15 \left(\log \frac{4}{3} \right) \frac{B}{A}} .$$

Finally, fix $\epsilon = \frac{1}{35} \min \left\{ \sqrt{\frac{A}{B}}, 1 \right\}$ (and, specifically, if $B = 0$, then $\epsilon = \frac{1}{35}$). Since $A \geq 2$, this value of ϵ yields

$$\mathbb{E}_{\underline{\nu}'}[n_T(k)] \leq \frac{5T}{4A} + \frac{1}{8}T \leq \frac{3}{4}T .$$

Substituting this bound and ϵ back to Equation (5) leads to the desired result and concludes the proof. \square

Proposition 2. *Let T be the time horizon, and let $c(1), \dots, c(A) \geq 0$ be arm-dependent querying costs. Also, let π be some bandit strategy such that for any bandit instance, it holds that $\mathbb{E}[B^q(T)] \leq B$. Then, for $A \geq 2$, there exists a bandit instance for which*

$$\mathbb{E}[\text{Reg}(T)] \geq \frac{1}{140} \min \left\{ T \sqrt{\frac{\sum_{a=1}^A c(a)}{B(1 + \log A)}}, T \right\} .$$

Proof. Without loss of generality, assume that $c(1) \geq c(2) \geq \dots \geq c(A)$. To prove the lemma, we will show that for any $i \in [A]$, there exists a bandit instance $\underline{\nu}'$ such that

$$\mathbb{E}_{\underline{\nu}'}[\text{Reg}(T)] \geq \frac{1}{140} \min \left\{ T \sqrt{\frac{ic(i)}{B(1 + \log A)}}, T \right\} . \quad (6)$$

In turn, this will imply that there exists an instance such that

$$\mathbb{E}_{\underline{\nu}'}[\text{Reg}(T)] \geq \frac{1}{140} \min \left\{ T \sqrt{\frac{\max_{i \in [A]} \{ic(i)\}}{B(1 + \log A)}}, T \right\} . \quad (7)$$

Finally, we apply Hölder's inequality to lower bound the maximum; for vectors x, y such that $x_i = ic(i)$ and $y_i = \frac{1}{i}$, it holds that

$$\max_{i \in [A]} \{ic(i)\} = \|x\|_{\infty} \geq \frac{\sum_{i=1}^A x_i y_i}{\|y\|_1} = \frac{\sum_{i=1}^A x_i y_i}{\|y\|_1} = \frac{\sum_{i=1}^A c(i)}{\sum_{i=1}^A \frac{1}{i}} \geq \frac{\sum_{i=1}^A c(i)}{1 + \log A} .$$

Substituting this bound to Equation (7) will then conclude the proof.

We start by proving Equation (6) when $i \geq 2$, a proof that greatly resembles the one of Proposition 1. Fix some $2 \leq i \leq A$ and assume that $c(i) > 0$, as otherwise, the bound trivially holds. Also, let $\underline{\nu}$ be an arm distribution such that all arms are Bernoulli-distributed with parameter $\frac{1}{2}$. Notice that it holds that $\sum_{a=1}^i \mathbb{E}_{\underline{\nu}}[n_T(a)] \leq T$, and since $c(j) \geq c(i)$ for all $j \leq i$, it also holds that $\sum_{a=1}^i \mathbb{E}_{\underline{\nu}}[n_T^q(a)] \leq \frac{B}{c(i)}$, or otherwise, the budget constraint is violated. Then, by Lemma 9, there must exist some arm $k \in [i]$ such that both $\mathbb{E}_{\underline{\nu}}[n_T(k)] \leq \frac{5T}{4i}$ and $\mathbb{E}_{\underline{\nu}}[n_T^q(k)] \leq \frac{15B}{ic(i)}$.

Next, we define a new bandit instance $\underline{\nu}'$ such that $\nu'_a = \nu_a$ for all $a \neq k$ and ν'_k is Bernoulli-distributed with parameter $\frac{1}{2} + \epsilon$, for some $\epsilon \in (0, \frac{1}{4})$. Then, by Lemma 8, we have

$$\mathbb{E}_{\underline{\nu}'}[n_T(k)] \leq \mathbb{E}_{\underline{\nu}}[n_T(k)] + 2\epsilon T \sqrt{\log \frac{4}{3}} \sqrt{\mathbb{E}_{\underline{\nu}}[n_T^q(k)]} ,$$

and recalling that $\mathbb{E}_{\underline{\nu}}[n_T(k)] \leq \frac{5T}{4i}$ and $\mathbb{E}_{\underline{\nu}}[n_T^q(k)] \leq \frac{15B}{ic(i)}$, we get

$$\mathbb{E}_{\underline{\nu}'}[n_T(k)] \leq \frac{5T}{4i} + 2\epsilon T \sqrt{15 \left(\log \frac{4}{3} \right) \frac{B}{ic(i)}} .$$

Finally, fix $\epsilon = \frac{1}{35} \min\left\{\sqrt{\frac{ic(i)}{B}}, 1\right\}$ (and, specifically, if $B = 0$, then $\epsilon = \frac{1}{35}$). Since $i \geq 2$, this value of ϵ yields

$$\mathbb{E}_{\nu'}[n_T(k)] \leq \frac{5T}{4A} + \frac{1}{8}T \leq \frac{3}{4}T .$$

Finally, as in the proof of Proposition 1, we lower bound the regret by

$$\mathbb{E}_{\nu'}[\text{Reg}(T)] = \sum_{a \neq k} \epsilon \mathbb{E}_{\nu'}[n_T(a)] = \epsilon T \left(1 - \frac{\mathbb{E}_{\nu'}[n_T(k)]}{T}\right) \geq \frac{1}{140} \min\left\{T\sqrt{\frac{ic(i)}{B}}, T\right\} ,$$

which leads to Equation (6) for any $2 \leq i \leq A$.

Finally, we prove that Equation (6) holds when $i = 1$. If $c(1) = 0$, the result trivially holds. Otherwise, by our assumptions, we have that $\mathbb{E}_{\nu}[n_T^q(k)] \leq \frac{B}{c(i)}$. For $\mathbb{E}_{\nu}[n_T^q(1)]$, we divide the proof into to cases:

- If $\mathbb{E}_{\nu}[n_T(1)] \leq \frac{T}{2}$, we set ν' such that $\nu'_a = \nu_a$ for all $a > 1$ and ν'_1 is Bernoulli-distributed with parameter $\frac{1}{2} + \epsilon$. Then, by Lemma 8, we have

$$\mathbb{E}_{\nu'}[n_T(k)] \leq \mathbb{E}_{\nu}[n_T(1)] + 2\epsilon T \sqrt{\log \frac{4}{3}} \sqrt{\mathbb{E}_{\nu}[n_T^q(1)]} \leq \frac{T}{2} + 2\epsilon T \sqrt{\log \frac{4}{3}} \sqrt{\frac{B}{c(1)}} ,$$

and fixing $\epsilon = \frac{1}{5} \min\left\{\sqrt{\frac{c(1)}{B}}, 1\right\}$ (or $\epsilon = \frac{1}{5}$ when $B = 0$) leads to

$$\mathbb{E}_{\nu'}[n_T(k)] \leq \frac{3T}{4} .$$

Then, the regret for this instance is lower bounded by

$$\mathbb{E}_{\nu'}[\text{Reg}(T)] = \sum_{a \neq 1} \epsilon \mathbb{E}_{\nu'}[n_T(a)] = \epsilon T \left(1 - \frac{\mathbb{E}_{\nu'}[n_T(1)]}{T}\right) \geq \frac{1}{20} \min\left\{T\sqrt{\frac{c(1)}{B}}, T\right\} ,$$

- If $\mathbb{E}_{\nu}[n_T(1)] \geq \frac{T}{2}$, we set ν' such that $\nu'_a = \nu_a$ for all $a > 1$ and ν'_1 is Bernoulli-distributed with parameter $\frac{1}{2} - \epsilon$. Then, by Lemma 8, we have

$$\mathbb{E}_{\nu'}[n_T(k)] \geq \mathbb{E}_{\nu}[n_T(1)] - 2\epsilon T \sqrt{\log \frac{4}{3}} \sqrt{\mathbb{E}_{\nu}[n_T^q(1)]} \geq \frac{T}{2} - 2\epsilon T \sqrt{\log \frac{4}{3}} \sqrt{\frac{B}{c(1)}} ,$$

and fixing $\epsilon = \frac{1}{5} \min\left\{\sqrt{\frac{c(1)}{B}}, 1\right\}$ (or $\epsilon = \frac{1}{5}$ when $B = 0$) leads to

$$\mathbb{E}_{\nu'}[n_T(k)] \geq \frac{T}{4} .$$

Then, the regret for this instance is lower bounded by

$$\mathbb{E}_{\nu'}[\text{Reg}(T)] = \epsilon \mathbb{E}_{\nu'}[n_T(1)] \geq \frac{1}{20} \min\left\{T\sqrt{\frac{c(1)}{B}}, T\right\} ,$$

Combining both cases leads to Equation (6) (with a better constant) and concludes the proof. \square

Lemma 9. Let $x, y \in \mathbb{R}_+^n$ for some $n \geq 2$, and assume that $\sum_{i=1}^n x_i \leq X$ and $\sum_{i=1}^n y_i \leq Y$. Then, for any $\alpha \in (1, 2)$ and $\beta \geq \frac{3\alpha}{\alpha-1}$, there exists an index t such that $x_t \leq \frac{\alpha X}{n}$ and $y_t \leq \frac{\beta Y}{n}$ simultaneously.

Proof. We divide the proof into the case where $n \leq \beta$ and $n > \beta$.

If $n \leq \beta$, recall that $y_i \geq 0$ and $\sum_{i=1}^n y_i \leq Y$. Therefore, for any $i \in [n]$, $y_i \leq Y \leq \frac{\beta}{n}Y$, and the required condition holds for all coordinates. Moreover, as $\alpha > 1$, by the pigeonhole principle, there exists at least one coordinate t such that $x_t \leq \frac{X}{n} \leq \frac{\alpha X}{n}$. Then, for this coordinate, both required conditions hold.

Next, we analyze the case where $n > \beta$. First notice that by the pigeonhole principle, there are at least $\lfloor (1 - \frac{1}{\alpha})n \rfloor$ coordinates such that $x_i \leq \frac{\alpha X}{n}$. To see this, assume in contradiction that this does not hold. Then, there are at least $n - \lfloor (1 - \frac{1}{\alpha})n \rfloor \geq \frac{n}{\alpha}$ coordinates such that $x_i > \frac{\alpha X}{n}$, and

$$\sum_{i=1}^n x_i \geq \sum_{i=1}^n x_i \cdot 1_{\left\{x_i > \frac{\alpha X}{n}\right\}} > \frac{\alpha X}{n} \cdot \left| \left\{ i : x_i > \frac{\alpha X}{n} \right\} \right| \geq \frac{\alpha X}{n} \cdot \frac{n}{\alpha} = X$$

which violates the assumption that $\sum_{i=1}^n x_i \leq X$. Similarly, there are at least $\lfloor (1 - \frac{1}{\beta})n \rfloor$ coordinates such that $y_i \leq \frac{\beta Y}{n}$. The number of coordinates for which at least one of the condition holds is then

$$\begin{aligned} \sum_{i=1}^n 1_{\left\{x_i \leq \frac{\alpha X}{n}\right\}} + 1_{\left\{y_i \leq \frac{\beta Y}{n}\right\}} &\geq \left\lfloor \left(1 - \frac{1}{\alpha}\right)n \right\rfloor + \left\lfloor \left(1 - \frac{1}{\beta}\right)n \right\rfloor \\ &\geq \left(1 - \frac{1}{\alpha}\right)n - 1 + \left(1 - \frac{1}{\beta}\right)n - 1 \\ &= \left(1 - \frac{1}{\alpha} - \frac{1}{\beta}\right)n + n - 2 \end{aligned}$$

We then use the fact that $n > \beta$ to bound the first term as

$$\left(1 - \frac{1}{\alpha} - \frac{1}{\beta}\right)n > \left(1 - \frac{1}{\alpha} - \frac{1}{\beta}\right)\beta = \beta \left(1 - \frac{1}{\alpha}\right) - 1 \geq 2,$$

where the last inequality is since $\alpha > 1$ and $\beta \geq \frac{3\alpha}{\alpha-1} = 3\frac{1}{1-1/\alpha}$. Substituting back leads to

$$\sum_{i=1}^n 1_{\left\{x_i \leq \frac{\alpha X}{n}\right\}} + 1_{\left\{y_i \leq \frac{\beta Y}{n}\right\}} > 2 + n - 2 = n.$$

Thus, there are strictly more than n coordinates where at least one condition holds, which implies that there is at least one coordinate t for which both conditions holds, i.e., $x_t \leq \frac{\alpha X}{n}$ and $y_t \leq \frac{\beta Y}{n}$. \square

A.4. Lower Bounds for Linear Bandits

Proposition 4. *Let $T \in \mathbb{N}$ be some time horizon and let π be a linear bandit policy such that $n_T^q \leq B$ a.s. for some fixed $B \leq T$. Then, there exists a d -dimensional linear bandit instance with arm set $[-1, 1]^d$ for which the expected regret of π is lower bounded by $\frac{dT}{80\sqrt{B}}$.*

Proof. Our proof closely follows the standard lower bound techniques used for proving the fundamental performance limit for linear bandits (e.g., Lattimore & Szepesvári, 2020, Theorem 24.1). Notably, we prove the lower bound for a problem with a fixed context space $\mathcal{X} = [-1, 1]^d$. Therefore, the conditions of Proposition 5 hold; namely, for any fixed problem, there exists a modified policy π' with the same expected regret as π that queries all rewards at the first B rounds. Thus, and without loss of generality, we assume that all rewards are queried at the B initial rounds.

As described in Section 2, we assume that the reward when choosing a context x_t is $R_t = \langle x_t, \theta \rangle + \eta_t$ for some unknown θ . In this section, we limit $\theta \in \Theta \triangleq \left\{ -\frac{1}{\sqrt{B}}, \frac{1}{\sqrt{B}} \right\}^d$. Moreover, we assume that the noises $\eta_t \sim \mathcal{N}(0, 1)$ are i.i.d. standard Gaussian variables. Notably, for a given θ , one can easily observe that the optimal context $x^* = \arg \max_{x \in \mathcal{X}} \langle x, \theta \rangle$ is such that $x^*(i) = \text{sign}(\theta(i)), \forall i \in [d]$. We also use the notation $Y_t = R_t \cdot q_t$ to denote the reward observed by the agent and write $R_t(\theta)$ when we want to emphasize the parameter that governed the reward generation.

Suppose that x_1, \dots, x_T are the contexts played by π at rounds $1, \dots, T$ and for any fixed problem with parameter θ , denote the measure of contexts and queried rewards induced by the interaction of π with the problem by \mathbb{P}_θ . Then, we lower bound the regret by

$$\begin{aligned}
 \mathbb{E}[\text{Reg}(T)] &= \mathbb{E} \left[\sum_{t=1}^T \langle x^*, \theta \rangle - \langle x_t, \theta \rangle \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^d |\theta(i)| - x_t(i)\theta(i) \right] \\
 &\geq \sum_{i=1}^d |\theta(i)| \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}\{\text{sign}(\theta(i)) \neq \text{sign}(x_t(i))\} \right] \quad (x_t(i) \in [-1, 1]) \\
 &\geq \sum_{i=1}^d |\theta(i)| \frac{T}{2} \mathbb{E} \left[\mathbf{1} \left\{ \sum_{t=1}^T \mathbf{1}\{\text{sign}(\theta(i)) \neq \text{sign}(x_t(i))\} \geq \frac{T}{2} \right\} \right] \\
 &= \frac{T}{2\sqrt{B}} \sum_{i=1}^d \mathbb{P}_\theta(i), \tag{8}
 \end{aligned}$$

where we denoted $\mathbb{P}_\theta(i) := \mathbb{P}_\theta \left(\sum_{t=1}^T \mathbf{1}\{\text{sign}(\theta(i)) \neq \text{sign}(x_t(i))\} \geq \frac{T}{2} \right)$. Next, we lower bound $\sum_{i=1}^d \mathbb{P}_\theta(i)$.

For any choice of $\theta \in \Theta$ and $i \in [d]$, let θ^i be a vector such that $\theta^i(j) = \theta(j), \forall j \in [d] \setminus \{i\}$ and $\theta^i(i) = -\theta(i)$. Now let us define $\mathbb{P}_{\theta^i}^c(i) := \mathbb{P}_{\theta^i} \left(\sum_{t=1}^T (\text{sign}(\theta^i(i)) \neq \text{sign}(x_{ti})) \leq \frac{T}{2} \right)$. By applying *Bretagnolle-Huber inequality* (Bretagnolle & Huber, 1979; Lattimore & Szepesvári, 2020, Theorem 14.2) we further get that for any $\theta \in \mathbb{R}^d$:

$$\mathbb{P}_\theta(i) + \mathbb{P}_{\theta^i}^c(i) \geq \frac{1}{2} \exp\{-\text{KL}(\mathbb{P}_\theta, \mathbb{P}_{\theta^i})\},$$

where $\text{KL}(\cdot, \cdot)$ denotes the Kullback–Leibler divergence. Applying the chain rule of KL-divergence while recalling that the

policy for both bandit instances is the same, we get:

$$\begin{aligned}
 \mathbb{P}_\theta(i) + \mathbb{P}_{\theta^i}(i) &\geq \frac{1}{2} \exp \left\{ - \sum_{t=1}^T \text{KL}(\mathbb{P}_\theta(Y_t(\theta)), \mathbb{P}_{\theta^i}(Y_t(\theta^i)) | x_1, Y_1, \dots, x_{t-1}, Y_{t-1}, x_t) \right\} \\
 &\stackrel{(a)}{=} \frac{1}{2} \exp \left\{ - \sum_{t=1}^B \text{KL}(\mathbb{P}_\theta(Y_t(\theta)), \mathbb{P}_{\theta^i}(Y_t(\theta^i)) | x_t) \right\} \\
 &\stackrel{(b)}{=} \frac{1}{2} \exp \left\{ - \frac{1}{2} \sum_{t=1}^B (\langle x_t, \theta - \theta^i \rangle)^2 \right\} \\
 &\stackrel{(c)}{\geq} \frac{1}{2} \exp(-2),
 \end{aligned}$$

Relation (a) follows since we assume the algorithm only queries for feedback for the first B rounds, so $Y_t \neq 0$ only generated for the first B rounds, and the algorithm does not observe any feedback for the last $T - B$ rounds. Moreover, this relation also requires the assumption that the noise η_t is i.i.d, and therefore, the reward only depends on the chosen context x_t . Furthermore, since we assume that $\eta_t \sim \mathcal{N}(0, 1)$, (b) follows by direct calculation:

$$\text{KL}(\mathbb{P}_\theta(R_t(\theta)), \mathbb{P}_{\theta^i}(R_t(\theta^i)) | x_t) = \text{KL}(\mathcal{N}(\langle x_t, \theta \rangle, 1), \mathcal{N}(\langle x_t, \theta^i \rangle, 1) | x_t) = (x_t^T(\theta - \theta^i))^2.$$

Relation (c) holds since $\theta(j) = \theta^i(j)$ for all $j \neq i$, $|\theta(i) - \theta^i(i)| = 2/\sqrt{B}$ and $x_t(i) \in [-1, 1]$; thus, we have that $(x_t^T(\theta - \theta^i))^2 = x_t(i)^2(\theta(i) - \theta^i(i))^2 \leq 4/B$.

Finally, notice that this holds for any $\theta \in \Theta$ and averaging over all $\theta \in \Theta$ we get:

$$\frac{1}{|\Theta|} \sum_{\theta \in \Theta} \sum_{i=1}^d \mathbb{P}_\theta(i) \sum_{i=1}^d \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{P}_\theta(i) \geq \frac{d \exp(-2)}{4}.$$

This implies there exists at least one $\theta \in \Theta$, say $\tilde{\theta}$, for which $\sum_{i=1}^d \mathbb{P}_{\tilde{\theta}}(i) > \frac{d \exp(-2)}{4}$. The claim now follows from eq. (8) and setting $\theta = \tilde{\theta}$.

□

B. Gap Between Adversarial and Stochastic Contexts

Theorem 1 (Black Box Reduction for Stochastic Contexts). *Let \mathbb{A} be an anytime algorithm with bounded regret $\mathbb{E}[\text{Reg}(T)] \leq \alpha T^\beta + C$ for some $\alpha, C \in \mathbb{R}_+, \beta \in [0, 1]$ and any $T \in \mathbb{N}$. Moreover, assume the budget is chosen by an oblivious adversary such that it is non-decreasing, $B(1) \geq 1$ and $B(t) \in \mathbb{N}$ for all $t \geq 1$. Then, the expected regret of Algorithm 1 with base algorithm \mathbb{A} and budget sequence $\{B(t)\}_{t \geq 1}$ is upper bounded by $\alpha T^\beta + C + \sum_{t=1}^T \frac{\alpha}{B(t)^{1-\beta}} + \frac{C}{B(t)}$.*

Remark (Budget is increasing in integer units). *For simplicity, we assume the budget is an integer and is increased by an integer number. If such assumption does not hold, the analysis should be modified by rounding the budget to the closest smallest integer number.*

Proof. Let $\{B(t)\}_{t \geq 1}$ be any fixed budget sequence and let $T \in \mathbb{N}$ be some arbitrary time. For brevity, and with slight abuse of notation, we denote $f(u, \pi) = \mathbb{E}[f(R_t)|u, \pi]$ and $\pi^*(u) = \arg \max_{\pi \in \Pi(u)} \mathbb{E}[f(u, \pi)]$. Under these notations, the expected regret can also be written as

$$\mathbb{E}[\text{Reg}(T)] = \mathbb{E} \left[\sum_{t=1}^T f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right].$$

Also, let $1\{Q_t\} = 1\{q_t = 1\}$ be the indicator function of the event a reward was queried at the t^{th} time step. Importantly, the greedy query rule of Algorithm 1 does not depend on the observations, and for this reason, the sequence $\{1\{Q_t\}\}_{t \geq 1}$ is fixed given $\{B(t)\}_{t \geq 1}$. Let $B^q(t) = \sum_{s=1}^{t-1} 1\{Q_s\}$ be number of times a budget was used until time step t . Let $\{E_1(t), E_2(t)\}_{t \geq 0}$ be sequence of events defined as

$$\begin{aligned} E_1(t) &= \{B(t) \geq B^q(t) + 1, B(t-1) < B^q(t-1) + 1\} \\ E_2(t) &= \{B(t) < B^q(t) + 1, B(t-1) \geq B^q(t-1) + 1\}. \end{aligned}$$

In words, $E_1(t)$ represents the event that at time step t budget became available whereas at time step $t-1$ it was not available. Furthermore, $E_2(t)$, represents the event budget became unavailable at time step t whereas it was available at time step $t-1$. Furthermore, define the following time steps (which are deterministic for a fixed budget sequence)

$$\tau_1(k) = \inf \left\{ t : \sum_{i=1}^t 1\{E_1(i)\} \geq k \right\} \quad \text{and} \quad \tau_2(k) = \inf \left\{ t : \sum_{i=1}^t 1\{E_2(i)\} \geq k \right\}.$$

That is, $\tau_1(k)$ is the k^{th} time the budget became available and $\tau_2(k)$ is the k^{th} time the budget became unavailable. We now prove the following claim. For any time step T the expected regret of Algorithm 1 is bounded by

$$\mathbb{E}[\text{Reg}(T)] \leq \mathbb{E} \left[\sum_{t \in \mathcal{I}_T^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] + \sum_{t=1}^T \frac{\alpha}{B(t)^{1-\beta}} + \frac{C}{B(t)}, \quad (9)$$

where $\mathcal{I}_t^q = \{k \in [t] : 1\{Q_k\} = 1\}$, i.e., all time steps until time step t in which a reward was queried. Notice that these are the only time steps where algorithm \mathbb{A} advances, so the algorithm has been effectively applied only on these time steps. Then, by the anytime regret assumption on algorithm \mathbb{A} it holds that the first term in (9) is bounded by $\alpha |\mathcal{I}_T^q|^\beta = \alpha (B^q(T))^\beta \leq \alpha \min\{B(T), T\}^\beta + C$ which completes the proof.

We establish Equation (9) via an induction on time steps in which the budget becomes unavailable, that is on $\tau_2(1), \tau_2(2), \dots$. Assume that at $t = 1$, there is an available budget, that is, $B(1) \geq 1$.

Base case. The claim holds for all $t \in [1, \tau_2(k=1) - 1]$. By assumption and definition, for all $t \in [1, \tau_2(1) - 1]$ Algorithm 1 query rewards. For this reason, the regret is bounded by the regret of algorithm \mathbb{A} . Thus, for all $t \in [1, \tau_2(1) - 1]$

$$\text{Reg}(t) = \mathbb{E} \left[\sum_{t \in \mathcal{I}_t^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] \leq \mathbb{E} \left[\sum_{t \in \mathcal{I}_t^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] + \sum_{l=1}^t \frac{\alpha}{B(l)^{1-\beta}} + \frac{C}{B(l)}.$$

Induction Step: Proving that the claim holds for all $t \in [1, \tau_2(k+1) - 1]$. Assume the claim holds until time step $t_k = \tau_2(k) - 1$, that is, until the k^{th} time the budget becomes unavailable. We now prove the regret bound holds until the $(k+1)^{\text{th}}$ time the budget becomes unavailable for $k \geq 1$ (if no such time exists then the exact same proof holds by replacing $\tau_2(k+1)$ by $\min\{\tau_2(k+1), T+1\}$ and $\tau_1(k+1)$ by $\min\{\tau_1(k+1), T+1\}$). Consider the time steps between the k^{th} and $(k+1)^{\text{th}}$ time the budget becomes unavailable. We can partition the time steps into two parts $[\tau_2(k), \tau_1(k+1) - 1]$, $\{\tau_1(k+1), \dots, \tau_2(k+1) - 1\}$, that is, between the k^{th} and $(k+1)^{\text{th}}$ time the budget becomes unavailable there exists a time step $\tau_1(k+1)$ in which the budget becomes available for the $(k+1)^{\text{th}}$ time (since we assumed at the first time step a budget is available, so $\tau_1(1) = 1$ and $\tau_2(1) > 1$). Furthermore, let $\text{Reg}(t_1 : t_2) = \sum_{t=t_1}^{t_2} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t)$ be partial sum of the cumulative regret. The following relations hold.

$$\begin{aligned}
 & \mathbb{E}[\text{Reg}(\tau_2(k+1) - 1)] \\
 &= \mathbb{E}[\text{Reg}(\tau_2(k) - 1)] + \mathbb{E}[\text{Reg}(\tau_2(k) : \tau_2(k+1) - 1)] \quad (\text{Additive form of regret}) \\
 &\leq \mathbb{E} \left[\sum_{t \in \mathcal{I}_{\tau_2(k)}^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] + \sum_{t=1}^{\tau_2(k)-1} \left(\frac{\alpha}{B(t)^{1-\beta}} + \frac{C}{B(t)} \right) + \mathbb{E}[\text{Reg}(\tau_2(k) : \tau_2(k+1) - 1)] \\
 &\quad (\text{Induction hypothesis}) \\
 &= \mathbb{E} \left[\sum_{t \in \mathcal{I}_{\tau_2(k)}^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] + \sum_{t=1}^{\tau_2(k)-1} \left(\frac{\alpha}{B(t)^{1-\beta}} + \frac{C}{B(t)} \right) + \mathbb{E}[\text{Reg}(\tau_2(k) : \tau_1(k+1) - 1)] \\
 &\quad + \mathbb{E} \left[\sum_{t=\tau_1(k+1)}^{\tau_2(k+1)-1} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] \\
 &= \mathbb{E} \left[\sum_{t \in \mathcal{I}_{\tau_2(k+1)-1}^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] + \sum_{t=1}^{\tau_2(k)-1} \left(\frac{\alpha}{B(t)^{1-\beta}} + \frac{C}{B(t)} \right) + \mathbb{E}[\text{Reg}(\tau_2(k) : \tau_1(k+1) - 1)]. \quad (10)
 \end{aligned}$$

where the equalities are by the regret definition and since in the interval $[\tau_2(k), \tau_2(k+1) - 1]$, rewards were queried only at $[\tau_1(k), \tau_2(k+1) - 1]$; therefore, $\mathcal{I}_{\tau_2(k+1)-1}^q = \mathcal{I}_{\tau_2(k)-1}^q \cup \{\tau_1(k), \dots, \tau_2(k+1) - 1\}$. All that remains is to bound $\mathbb{E}[\text{Reg}(\tau_2(k) : \tau_1(k+1) - 1)]$. In time steps $t \in [\tau_2(k) : \tau_1(k+1) - 1]$ there is no available budget and Algorithm 1 acts in accordance to one of the previous queried rounds; i.e., $j \sim \text{Uniform}(I_t^q)$ and $\pi_t = A_j(u_t)$. Notably, in these time steps, we have $I_t^q = I_{\tau_2(k)-1}^q$. Also recall that $\mathcal{I}_{\tau_2(k)-1}^q$ is deterministic given $\{B(t)\}_{t \geq 1}$. Thus, the expected value of π_t is given by

$$\begin{aligned}
 \mathbb{E}[f(u_t, \pi^*(u_t)) - f(u_t, \pi_t)] &= \mathbb{E}[\mathbb{E}[f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) | F_{\tau_2(k)-1}]] \\
 &= \mathbb{E} \left[\frac{1}{|I_{\tau_2(k)-1}^q|} \sum_{j \in I_{\tau_2(k)-1}^q} \mathbb{E}[f(u_t, \pi^*(u_t)) - f(u_t, A_j(u_t)) | F_{\tau_2(k)-1}] \right] \\
 &\stackrel{(1)}{=} \mathbb{E} \left[\frac{1}{|I_{\tau_2(k)-1}^q|} \sum_{j \in I_{\tau_2(k)-1}^q} \mathbb{E}[f(u_t, \pi^*(u_t)) - f(u_t, A_j(u_t)) | F_{\tau_2(j)-1}] \right] \\
 &\stackrel{(2)}{=} \mathbb{E} \left[\frac{1}{|I_{\tau_2(k)-1}^q|} \sum_{j \in I_{\tau_2(k)-1}^q} \mathbb{E}[f(u_j, \pi^*(u_j)) - f(u_t, A_j(u_j)) | F_{\tau_2(j)-1}] \right] \\
 &= \frac{1}{|I_{\tau_2(k)-1}^q|} \mathbb{E} \left[\sum_{j \in I_{\tau_2(k)-1}^q} f(u_j, \pi^*(u_j)) - f(u_t, A_j(u_j)) \right]
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(3)}{\leq} \frac{\alpha |I_{\tau_2(k)-1}^q|^\beta + C}{|I_{\tau_2(k)-1}^q|} \\
 &\stackrel{(4)}{=} \frac{\alpha}{B(\tau_2(k))^{1-\beta}} + \frac{C}{B(\tau_2(k))}.
 \end{aligned}$$

Relation (1) is since rewards and contexts are i.i.d. between rounds and \mathbb{A}_j only depends on samples in $\mathcal{I}_{\tau(j)-1}^q$. Equality (2) is since the contexts are i.i.d. and \mathbb{A}_j only depends on samples in $F_{\tau_2(j)-1}$ and internal randomness; therefore, for any $t \geq j$,

$$\mathbb{E}[f(u_t, \pi^*(u_t)) - f(u_t, \mathbb{A}_j(u_t)) | F_{\tau_2(j)-1}] = \mathbb{E}[f(u_j, \pi^*(u_j)) - f(u_j, \mathbb{A}_j(u_j)) | F_{\tau_2(j)-1}].$$

Next, (3) is by the regret bound of algorithm \mathbb{A} on samples in $\mathcal{I}_{\tau_2(k)-1}^q$, the only rounds where the algorithm advanced. Finally, (4) holds since the budget at time step $\tau_2(k)$ equals to the number of queries at time step $\tau_2(k) - 1$ (if not, then there is an available budget at time step $\tau_2(k)$, in contradiction to its definition). Plugging this back into (10) we get

$$\begin{aligned}
 (10) &\leq \mathbb{E} \left[\sum_{t \in \mathcal{I}_{\tau_2(k+1)-1}^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] + \sum_{t=1}^{\tau_2(k)} \left(\frac{\alpha}{B(t)^{1-\beta}} + \frac{C}{B(t)} \right) + (\tau_1(k+1) - 1 - \tau_2(k)) \alpha B(\tau_2(k))^{\beta-1} \\
 &= \mathbb{E} \left[\sum_{t \in \mathcal{I}_{\tau_2(k+1)-1}^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] + \sum_{t=1}^{\tau_1(k+1)-1} \left(\frac{\alpha}{B(t)^{1-\beta}} + \frac{C}{B(t)} \right) \\
 &\hspace{20em} (B(t) \text{ is fixed for } t \in [\tau_2(k), \tau_1(k+1)]) \\
 &\leq \mathbb{E} \left[\sum_{t \in \mathcal{I}_{\tau_2(k+1)-1}^q} f(u_t, \pi^*(u_t)) - f(u_t, \pi_t) \right] + \sum_{t=1}^{\tau_2(k+1)-1} \left(\frac{\alpha}{B(t)^{1-\beta}} + \frac{C}{B(t)} \right), \quad (\text{Budget is positive})
 \end{aligned}$$

which proves the induction hypothesis. □

Proposition 5. *Assume that the decision-making problem is non-contextual ($u_t = \phi, \forall t$) with no environment feedback ($Z_t = \phi, \forall t$) and unit-querying costs. Then, for any $T, B \in \mathbb{N}$ such that $B \leq T$ and any policy π under which $n_T^q \leq B$, there exists a policy π' such that $q_t = 1$ for all $t \in [B]$ (and zero otherwise) and $\mathbb{E}[\text{Reg}(T)|\pi'] = \mathbb{E}[\text{Reg}(T)|\pi]$.*

Proof. We prove the claim iteratively; specifically, we prove that for any $b \in \{0, \dots, B-1\}$ and any policy π^b for which $q_t = 1, \forall t \in [b]$, there exists a policy π^{b+1} such that $q_t = 1, \forall t \in [b+1]$ and $\mathbb{E}[\text{Reg}(T)|\pi^{b+1}] = \mathbb{E}[\text{Reg}(T)|\pi^b]$. Then, we can choose $\pi^0 = \pi$ and apply this result B times to obtain $\pi' = \pi^B$ for which $q_t = 1$ for all $t \in [B]$ and $\mathbb{E}[\text{Reg}(T)|\pi'] = \mathbb{E}[\text{Reg}(T)|\pi]$.

To prove this claim, we further delve into the probabilistic model of the decision-making problem, which we carefully choose for this prove (see, e.g., [Lattimore & Szepesvári 2020](#), Chapters 4.6,4.7, for more details). Notice that under our assumptions, the action set is fixed for all time steps and only reward is generated. Also recall that under our model, and given the action, the reward is generated independently of other rounds. Denote by ν_a , the distribution of R_t given an action a was taken, and let $B^q(t)$ be the number of queries taken up to time t . Then, we describe the decision-process through the following random variables:

- U is a uniform random that is generated prior to the game and represents all randomness of the agent
- a_1, \dots, a_T and q_1, \dots, q_T are the (random) actions and queries taken by the agent
- We let R_1^q, \dots, R_T^q and R_1^u, \dots, R_T^u be two reward sequences, one for queried actions and one for unqueried ones. Formally, if $q_t = 1$, then $R_{B^q(t)}^q$ is sampled independently at random from ν_{a_t} and $R_t = R_{B^q(t)}^q$. Similarly, if $q_t = 0$, then $R_{t-B^q(t)}^u$ is sampled independently at random from ν_{a_t} and $R_t = R_{t-B^q(t)}^u$. We brevity, we denote $Y_t = R_t \cdot q_t$.
- We emphasize that under this model, a policy π_t deterministically maps U, t and the sequence $\{Y_k\}_{k:q_k=1}$ to action and queries a_t^π and q_t^π (given the internal randomization, the decision rule is deterministic, so previous actions are not needed to describe the mapping).

Now, let $b \in \{0, \dots, B-1\}$ and let $\pi = \pi^b$ be a policy such that $q_t = 1$ for all $n \in [b]$, and let U, Y_1, \dots, Y_b an instantiation of the internal randomness and the rewards. Notably, under our model, $Y_t = R_t^q$ for all $t \in [b]$. Also note that until a new query is taken, the policy receives no new input. This implies that the action sequence up to the $(b+1)^{\text{th}}$ -query is deterministic given U, Y_1, \dots, Y_b . We denote this time by $\tau(b+1)$ and say that $\tau(b+1) = \phi$ if no additional query is taken.

Next, we define a new policy $\pi' = \pi^{b+1}$ as follows:

- For any $t \leq b$ we fix $\pi' = \pi$.
- For any U, Y_1, \dots, Y_b , if $\tau(b+1) = \phi$, we set $q_{b+1}^{\pi'} = 1$ but continue choosing $a_t^{\pi'} = a_t^\pi$ for all $t \geq b+1$ and $q_{b+1}^{\pi'} = 0$ for any $t \geq b+2$.
- For any U, Y_1, \dots, Y_b , if $\tau(b+1) \neq \phi$, we permute the action at time $\tau(b+1)$ to time $b+1$ and delay the actions of π at times $b+1, \dots, \tau(b+1)-1$ by a single time step. Formally:
 - $a_{b+1}^{\pi'} = a_{\tau(b+1)}^{\pi^b}$ and $q_{b+1} = 1$.
 - $a_t^{\pi'} = a_{t-1}^{\pi^b}$ and $q_t = 0$ for all $t \in \{b+2, \dots, \tau(b+1)\}$.
 - $\pi'_t = \pi_t$ for all $t > \tau(b+1)$.

Clearly, $q_t^{\pi'} = 1, \forall t \in [b+1]$. Thus, it remains to prove that $\mathbb{E}[\text{Reg}(T)|\pi'] = \mathbb{E}[\text{Reg}(T)|\pi]$. Denote the instantaneous regret of algorithm π by $r_t^\pi = \max_{\pi \in \Pi} \mathbb{E}[f(R_t)|\pi] - f(R_t)$, where R_t is generated when playing according to π . We use a coupling argument, where as long as the policies agree, the model simultaneously generates the same reward for both process. As soon as either R_k^q or R_k^u is supposed to be generated from a different action, then the processes split into two independent process and continue separately. Since $\pi'_t = \pi_t$ for all $t \leq b$, it also implies that the action and reward processes for both polices are identical up to time b and $r_t^{\pi'} = r_t^\pi$ for all $t \leq b$.

For the rest of the time steps, for any instantiation of U, Y_1, \dots, Y_b , we divide the analysis into two cases:

1. $\tau(b+1) \neq \phi$. In this case, notice that during time steps $b+1, \dots, \tau(b+1)$, both policies generate R_{b+1}^q from $a_{\tau(b+1)}^\pi$ and generate $R_1^u, \dots, R_{\tau(b+1)-b-1}$ from $a_{b+1}^\pi, \dots, a_{\tau(b+1)-1}^\pi$. Therefore, for these time steps, both policies agree on the actions and generate the same rewards, albeit in a different order. Specifically, this implies that $\sum_{t=b+1}^{\tau(b+1)} r_t^{\pi'} = \sum_{t=b+1}^{\tau(b+1)} r_t^\pi$. Moreover, both policies generated the same queried reward R_{b+1}^q and continues the same for any $t > \tau(b+1)$, which implies that they will generate the same actions and rewards until the end of the interactions. Thus, in this case, we have that $\sum_{t=b+1}^T r_t^{\pi'} = \sum_{t=b+1}^T r_t^\pi$.
2. $\tau(b+1) = \phi$. In this case, no additional reward is queried, and the same sequence of actions a_{b+1}, \dots, a_T is deterministically chosen as a function of U, Y_1, \dots, Y_b , for both π and π' . Thus, as the reward is generated from the same distribution for both policies, we have for any $t > b+1$

$$E[r_t^{\pi'} | U, Y_1, \dots, Y_b] = E[r_t^\pi | U, Y_1, \dots, Y_b]$$

Combining both parts, and using the tower property we have

$$\begin{aligned} \mathbb{E} \left[\sum_{t=b+1}^T r_t^{\pi'} \right] &= \mathbb{E} \left[\mathbb{E} \left[\sum_{t=b+1}^T r_t^{\pi'} \middle| U, Y_1, \dots, Y_b \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\underbrace{1\{\tau(b+1) \neq \phi\} \sum_{t=b+1}^T r_t^{\pi'} \middle| U, Y_1, \dots, Y_b}_{=\sum_{t=b+1}^T r_t^\pi} + \sum_{t=b+1}^T 1\{\tau(b+1) = \phi\} \underbrace{\mathbb{E} \left[r_t^{\pi'} \middle| U, Y_1, \dots, Y_b \right]}_{=\mathbb{E} \left[r_t^\pi \middle| U, Y_1, \dots, Y_b \right]} \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[1\{\tau(b+1) \neq \phi\} \sum_{t=b+1}^T r_t^\pi \middle| U, Y_1, \dots, Y_b \right] + \sum_{t=b+1}^T 1\{\tau(b+1) = \phi\} \mathbb{E} \left[r_t^\pi \middle| U, Y_1, \dots, Y_b \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\sum_{t=b+1}^T r_t^\pi \middle| U, Y_1, \dots, Y_b \right] \right] \\ &= \mathbb{E} \left[\sum_{t=b+1}^T r_t^\pi \right]. \end{aligned}$$

Finally, recalling that the polices are identical for $t \leq b$, we get the desired result for $\pi' = \pi^{b+1}$ and $\pi = \pi^b$:

$$\mathbb{E}[\text{Reg}(T) | \pi'] = \mathbb{E} \left[\sum_{t=1}^b r_t^{\pi'} \right] + \mathbb{E} \left[\sum_{t=b+1}^T r_t^{\pi'} \right] = \mathbb{E} \left[\sum_{t=1}^b r_t^\pi \right] + \mathbb{E} \left[\sum_{t=b+1}^T r_t^\pi \right] = \mathbb{E}[\text{Reg}(T) | \pi] .$$

□

Proposition 6 (Greedy Reduction Degrades in the Presence of Adversary). *If an adaptive adversary controls either (i) the contexts, or (ii) the budget, then for any base algorithm \mathbb{A} used in Algorithm 1, there exists a contextual MAB problem with two contexts and two arms such that $\mathbb{E}[\text{Reg}(T)] \geq \frac{T}{4}$, even if $\mathbb{E}[B(t)] = \frac{t}{2}$ for all $t \in [T]$.*

Proof. Consider a contextual multi-armed bandit instance with two contexts. That is, the environment has two contexts $u = 1, 2$; if we observe $u = 1$, then we interact with the first MAB problem, and if $u = 2$, we interact with the second MAB problem. We assume there is no relation between the first and second MAB problems.

Greedy Reduction Algorithm in the Presence of Adversarial Contexts.

Assume the budget sequence increases at each episode w.p. $\frac{1}{2}$ by a unit. Assume the adversary picks $u = 1$ if the budget increases by one and $u = 2$ if the budget does not increase. Applying the Greedy Reduction, the algorithm will only query information from MAB $u = 1$ and will have no information on the rewards for MAB $u = 2$. Let the bandit problems of both contexts be two-armed problems. Then, there exists an arm a^* of context $u = 2$ that the algorithm samples, in expectation, at most $\frac{T}{4}$ times, namely,

$$\mathbb{E} \left[\sum_{t=1}^T 1\{u_t = 2, a_t = a^*\} \right] \leq \frac{1}{2} \mathbb{E} \left[\sum_{t=1}^T 1\{u_t = 2\} \right] = \frac{T}{4},$$

where we used the fact that $\Pr(u_t = 2) = \Pr(B(t) = B(t-1) + 1) = \frac{1}{2}$. For this arm, we fix the reward to be equal $R = 1$, and for the other arm, we let $R = 0$. For simplicity, we fix the means of all arms in context $u = 1$ to be identical. Then, the regret is lower bounded by

$$\mathbb{E}[\text{Reg}(T)] = \mathbb{E} \left[\sum_{t=1}^T 1\{u_t = 2, a_t \neq a^*\} \right] = \mathbb{E} \left[\sum_{t=1}^T 1\{u_t = 2\} \right] - \mathbb{E} \left[\sum_{t=1}^T 1\{u_t = 2, a_t = a^*\} \right] \geq \frac{T}{2} - \frac{T}{4} = \frac{T}{4}$$

Notice that by definition, $\mathbb{E}[B(t)] = \frac{t}{2}$ for all $t \in [T]$.

Greedy Reduction Algorithm in the Presence of Adversarial Budget.

The example is symmetric to the previous one, i.e., we exchange the roles of the context and the budget. Assume that the context are chosen stochastically such that where $\Pr(u_t = 1) = \Pr(u_t = 2) = \frac{1}{2}$. If the budget increases by one each time $u_t = 1$ then Greedy Reduction will not acquire any information on the MAB with $u_t = 2$. Repeating the same analysis as the previous case results in $\mathbb{E}[\text{Reg}(T)] \geq \Omega(T)$ as well as $\mathbb{E}[B(t)] = \frac{t}{2}$ for all $t \in [T]$. \square

C. Confidence-Budget Matching for Multi Armed Bandits

Algorithm 3 CBM-UCB

- 1: **Initialize:** $n_1^q(a) = 0, \bar{r}_1(a) = 0$
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: Observe current budget $B(t)$
 - 4: Act with $a_t \in \arg \max_a UCB_t(a)$
 - 5: **if** $CI_t(a_t) \geq 4\sqrt{\frac{6 \log(At) \sum_{a=1}^A c(a)}{B(t)}}$ (or, alternatively, $n_{t-1}^q(a_t) \leq \frac{B(t)}{4 \sum_{a=1}^A c(a)}$) **then**
 - 6: Ask for feedback ($q_t = 1$)
 - 7: Observe R_t and update $n_t^q(a_t), \bar{r}_t(a_t)$
 - 8: **end if**
 - 9: **end for**
-

We start with some notations: let $n_t^q(a) = \sum_{k=1}^t 1\{a_k = a, q_k = 1\}$ be the number of times arm a was queried up to time t and let $\bar{r}_t(a) = \frac{1}{n_t^q(a) \vee 1} \sum_{k=1}^t R_k 1\{a_k = a, q_k = 1\}$ be its empirical mean. We use Hoeffding-based CI, i.e., if $b_t^r(a) \triangleq \sqrt{\frac{3 \log(At)}{2n_{t-1}^q(a) \vee 1}}$, then

$$UCB_t(a) = \bar{r}_{t-1}(a) + b_t^r(a) \quad \text{and} \quad LCB_t(a) = \bar{r}_{t-1}(a) - b_t^r(a),$$

which leads to $CI_t(a) = UCB_t(a) - LCB_t(a) = 2b_t^r(a) = \sqrt{6 \frac{\log At}{n_t^q(a) \vee 1}}$.

Remark. The observant reader might find the CBM condition wasteful; a better condition, for example, would be to sample arms if $n_{t-1}^q(a_t) \leq \frac{B(t)}{\sum_{a=1}^A c(a)} + c(a_t)$. Indeed, doing so will improve the constants of the regret bounds. Nonetheless, the main goal of this section is to demonstrate the techniques we use for the more complex settings (linear bandits and RL), where the CBM principle is not equivalent to count-thresholding and there is no clear way to tune the querying condition to be tighter.

We now prove the regret bound for CBM-UCB:

Theorem 2 (Confidence Budget Matching for Multi Armed Bandits). *For any querying costs $c(1), \dots, c(A) \geq 0$, any adaptive non-decreasing adversarially chosen sequence $\{B(t)\}_{t \geq 1}$ and for any $T \geq 1$, the expected regret of CBM-UCB is upper bounded by*

$$\mathbb{E}[\text{Reg}(T)] \leq \sqrt{24AT \log(AT)} + A\sqrt{6 \log(AT)} + 4\sqrt{6 \log(AT) \sum_{a=1}^A c(a) \sum_{t=1}^T \mathbb{E} \left[\sqrt{\frac{1}{B(t)}} \right]} + 2$$

Proof. Define the filtration $\{F_t\}_{t \geq 0}$ where F_{t-1} contains $\{B(1), a_1, R_1, \dots, B(t-1), a_{t-1}, R_{t-1}, B(t)\}$, i.e., all past actions, rewards and budgets, combined with the value of total budget at time step t . We also assume w.l.o.g. that $A \geq 2$, otherwise the regret is always zero. The good event is defined as

$$E^r(t) = \left\{ a \in A : |\bar{r}_{t-1}(a) - r(a)| \leq \sqrt{\frac{3 \log(At)}{2n_{t-1}^q(a) \vee 1}} \stackrel{\text{def}}{=} b_t^r(a) \right\}.$$

Importantly, notice that when $E^r(t)$ holds, then $r(a^*) \leq UCB_t(a^*) \leq UCB_t(a_t)$, and thus, the UCB of the chosen arm is optimistic. Moreover, the event directly implies that $r(a_t) \geq LCB_t(a_t)$.

Meeting the budget constraint We prove that the CBM principle of Algorithm 3 never violates the budget constraint in Lemma 11, that is, $B^q(t) \leq B(t)$ for all $t \in [T]$. Thus, throughout the proof, we assume that $q_t = 1$ if and only if the CBM query rule decides so (i.e., the case where the algorithm wants to set $q_t = 1$ and does not have enough budget to do so cannot happen).

Regret analysis. We decouple the regret as follows.

$$\begin{aligned}
 \mathbb{E}[\text{Reg}(T)] &= \sum_{t=1}^T \mathbb{E}[(r^* - r(a_t))1\{E^r(t)\}] + \sum_{t=1}^T \mathbb{E}[(r^* - r(a_t))1\{\overline{E^r(t)}\}] \\
 &\leq \sum_{t=1}^T \mathbb{E}[(r^* - r(a_t))1\{E^r(t)\}] + \sum_{t=1}^T \mathbb{E}[1\{\overline{E^r(t)}\}] \quad ((r^* - r(a)) \in 1 \text{ for all } a \in [A]) \\
 &\leq \sum_{t=1}^T \mathbb{E}[(r^* - r(a_t))1\{E^r(t)\}] + \sum_{t=1}^T \frac{1}{t^2} \quad (\text{Lemma 10}) \\
 &\leq \underbrace{\sum_{t=1}^T \mathbb{E}[1\{q_t = 1\}(r^* - r(a_t))1\{E^r(t)\}]}_{(i)} + \underbrace{\sum_{t=1}^T \mathbb{E}[1\{q_t = 0\}(r^* - r(a_t))1\{E^r(t)\}]}_{(ii)} + 2.
 \end{aligned}$$

Terms (i) and (ii) represent the regret over episodes in which feedback was queried and not queried, respectively. We bound each of the terms separately.

Bound on term (i), episodes in which feedback is queried, under the good event. The following relations hold

$$\begin{aligned}
 (i) &= \sum_{t=1}^T \mathbb{E}[1\{q_t = 1\}(r^* - r(a_t))1\{E^r(t)\}] \\
 &\leq \sum_{t=1}^T \mathbb{E}[1\{q_t = 1\}1\{E^r(t)\}(UCB(a_t) - r(a_t))] \quad (\text{Optimism in the event } E^r(t)) \\
 &\leq \sum_{t=1}^T \mathbb{E}[1\{q_t = 1\}1\{E^r(t)\}(UCB(a_t) - LCB(a_t))] \quad (\text{In the event } E^r(t)) \\
 &\leq \mathbb{E}\left[\sum_{t=1}^T 1\{q_t = 1\}\sqrt{\frac{6 \log(At)}{n_{t-1}^q(a_t) \vee 1}}\right].
 \end{aligned}$$

We bound the term in the expectation as follows (for every history sequence).

$$\begin{aligned}
 \sum_{t=1}^T 1\{q_t = 1\}\sqrt{\frac{6 \log(At)}{n_{t-1}^q(a_t) \vee 1}} &\leq \sqrt{6 \log(AT)} \sum_{t=1}^T \frac{1\{q_t = 1\}}{\sqrt{n_{k-1}^q(a_t) \vee 1}} \\
 &\stackrel{(*)}{=} \sqrt{6 \log(AT)} \sum_{a=1}^A \sum_{i=0}^{n_T^q(a)} \frac{1}{\sqrt{i \vee 1}} \\
 &\leq \sqrt{6 \log(AT)} \sum_a \left(2\sqrt{n_T^q(a)} + 1\right) \quad \left(\sum_{i=1}^T \frac{1}{\sqrt{i}} \leq 2\sqrt{T}\right) \\
 &\leq \sqrt{24 \log(AT)} \sqrt{An_T^q} + A\sqrt{6 \log(AT)}, \quad (\text{Jensen's inequality and } \sum_a n_T^q(a) = n_T^q)
 \end{aligned}$$

where (*) holds since every time an action a was queried, its counter advanced by 1. In the second relation, the summation is performed over time steps a reward is queried. Lastly, using $n_T^q \leq T$, i.e., the number of times the algorithm queried feedback is smaller than the total number of round, we get,

$$(i) \leq \sqrt{24AT \log(AT)} + A\sqrt{6 \log(AT)}.$$

Bound on term (ii), episodes in which feedback is not queried. To bound this term, we use the query rule, that is,

$$\begin{aligned}
 (ii) &= \sum_{t=1}^T \mathbb{E}[1\{q_t = 0\}1\{E^r(t)\}(r^* - r(a_t))] \\
 &\leq \sum_{t=1}^T \mathbb{E}[1\{q_t = 0\}1\{E^r(t)\}(UCB(a_t) - LCB(a_t))] && \text{(Optimism)} \\
 &\leq \sum_{t=1}^T \mathbb{E} \left[4 \sqrt{\frac{6 \log(At) \sum_{a=1}^A c(a)}{B(t)}} \right] && \text{(CBM query rule)} \\
 &\leq 4 \sqrt{6 \log(AT) \sum_{a=1}^A c(a) \sum_{t=1}^T \mathbb{E} \left[\sqrt{\frac{1}{B(t)}} \right]}.
 \end{aligned}$$

Combining the bounds. Combining the two bounds we conclude the proof,

$$\mathbb{E}[\text{Reg}(T)] \leq \sqrt{24AT \log(AT)} + A \sqrt{6 \log(AT)} + 4 \sqrt{6 \log(AT) \sum_{a=1}^A c(a) \sum_{t=1}^T \mathbb{E} \left[\sqrt{\frac{1}{B(t)}} \right]} + 2.$$

□

C.1. The Good Event

Lemma 10 (The Good Event). *For any $A \geq 2$ and $t \geq 1$, it holds that $\Pr(\overline{E^r(t)}) \leq \frac{1}{t^2}$.*

Proof. Fix $t \geq 1$. The following relations hold by applying Hoeffding's inequality and the union bound. Denote by $\hat{r}_n(a)$, the empirical mean of n i.i.d random variables over $[0, 1]$ with an expectation $r(a)$. Then, we have

$$\begin{aligned}
 &\Pr \left(|\bar{r}_{t-1}(a_t) - r(a_t)| \geq \sqrt{\frac{3 \log(At)}{2n_{t-1}^q(a_t) \vee 1}} \right) \\
 &= \sum_{a=1}^A \Pr \left(|\bar{r}_{t-1}(a) - r(a)| \geq \sqrt{\frac{3 \log(At)}{2n_{t-1}^q(a) \vee 1}}, a_t = a \right) \\
 &\leq \sum_{a=1}^A \Pr \left(|\bar{r}_{t-1}(a) - r(a)| \geq \sqrt{\frac{3 \log(At)}{2n_{t-1}^q(a) \vee 1}} \right) \\
 &= \sum_{a=1}^A \Pr \left(\cup_{n=0}^t \left\{ |\bar{r}_{t-1}(a) - r(a)| \geq \sqrt{\frac{3 \log(At)}{2n \vee 1}}, n_{t-1}^q(a) = n \right\} \right) \\
 &= \sum_{a=1}^A \Pr \left(\cup_{n=1}^t \left\{ |\hat{r}_n(a) - r(a)| \geq \sqrt{\frac{3 \log(At)}{2n}}, n_{t-1}^q(a) = n \right\} \right) && \text{(Holds trivially for } n = 0 \text{)} \\
 &\leq \sum_{a=1}^A \sum_{n=1}^t \Pr \left(|\hat{r}_n(a) - r(a)| \geq \sqrt{\frac{3 \log(At)}{2n}} \right) && \text{(Union bound)} \\
 &\leq \sum_{a=1}^A \sum_{n=1}^t \frac{1}{At^3} && \text{(Hoeffding's inequality)} \\
 &= \frac{1}{t^2}.
 \end{aligned}$$

□

C.2. The Budget Constraint is not Violated

Lemma 11 (CBM-UCB: Budget Constraint is Satisfied). *For any $t \geq 1$ the budget constraint is not violated (a.s.), $B^q(t) \leq B(t)$.*

Proof. Assume w.l.o.g. that $c(a) > 0$ for at least one action, as otherwise, the budget constraint can never be violated. Similarly, assume that $B(1) > 0$; otherwise, until a budget becomes available, no reward will be queried, and the same analysis would hold, starting from the first time querying became available. The following relations hold for any history.

$$\begin{aligned}
 B^q(t) &= \sum_{k=1}^t c(a_k) 1\{q_k = 1\} \\
 &= \sum_{a=1}^A \sum_{k=1}^t c(a) 1\{a_k = a, q_k = 1\} \\
 &\leq \sum_{a=1}^A c(a) \sum_{k=1}^t 1\{a_k = a, q_k = 1\} \frac{CI_t(a_t)}{4\sqrt{6 \log(Ak) (\sum_{a'=1}^A c(a'))/B(k)}} \quad (\text{By the update rule, when } q_k = 1) \\
 &= \sum_{a=1}^A c(a) \sum_{k=1}^t 1\{a_k = a, q_k = 1\} \frac{\sqrt{6 \log(Ak) / (n_{k-1}^q(a) \vee 1)}}{4\sqrt{6 \log(Ak) (\sum_{a'=1}^A c(a'))/B(k)}} \\
 &= \frac{1}{4\sqrt{\sum_{a'=1}^A c(a')}} \sum_{a=1}^A c(a) \sum_{k=1}^t 1\{a_k = a, q_k = 1\} \sqrt{\frac{1}{(n_{k-1}^q(a) \vee 1)}} \sqrt{B(k)} \\
 &\leq \frac{1}{4\sqrt{\sum_{a'=1}^A c(a')}} \sqrt{B(t)} \sum_{a=1}^A c(a) \sum_{i=0}^{n_t^q(a)} \sqrt{\frac{1}{i \vee 1}} \quad (B(t) \text{ is increasing}) \\
 &\leq \frac{1}{2\sqrt{\sum_{a'=1}^A c(a')}} \sqrt{B(t)} \sum_{a=1}^A c(a) \sum_{i=1}^{n_t^q(a)} \sqrt{\frac{1}{i}} \\
 &\leq \frac{1}{\sqrt{\sum_{a'=1}^A c(a')}} \sqrt{B(t)} \sum_{a=1}^A c(a) \sqrt{n_t^q(a)} \quad (\sum_{i=1}^T \frac{1}{\sqrt{i}} \leq 2\sqrt{T}) \\
 &\leq \frac{1}{\sqrt{\sum_{a'=1}^A c(a')}} \sqrt{B(t)} \sqrt{\sum_{a'=1}^A c(a')} \sqrt{\sum_{a=1}^A c(a) n_t^q(a)} \quad (\text{Cauchy-Schwartz inequality}) \\
 &= \sqrt{B(t)} \sqrt{B^q(t)}.
 \end{aligned}$$

Rearranging we get that $B^q(t) \leq B(t)$ for any $t \geq 0$. □

D. Confidence-Budget Matching for Linear Bandits

Algorithm 4 CBM-OFUL

- 1: **Require:** $\delta \in (0, 1), \lambda, d, \sigma, L, D > 0$
 - 2: **Set:** $V_0 = \lambda I_d, \hat{\theta}_0 = \mathbf{0}_d, v_t = \sqrt{2d \log(1 + \frac{tL^2}{d\lambda})}, l_t = \max\left\{1, \sigma \sqrt{2d \log\left\{\frac{1+tL^2/\lambda}{\delta}\right\}} + \lambda^{1/2}D\right\}$
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: Observe current budget $B(t)$ and context space \mathcal{X}_t
 - 5: Act with $x_t \in \arg \max_{x \in \mathcal{X}_t} \max_{\theta \in C_{t-1}} \langle x, \theta \rangle$ for $C_t = \left\{\theta \in \mathbb{R}^d : \|\hat{\theta}_t - \theta\|_{V_t} \leq l_t\right\}$
 - 6: Calculate $CI_t(x_t) = 2l_{t-1} \min\left\{\|x_t\|_{V_{t-1}^{-1}}, 1\right\}$
 - 7: **if** $CI_t(x_t) \geq \frac{l_{t-1}v_{B(t)}}{2\sqrt{B(t)}}$ (or, alternatively, $\|x_t\|_{V_{t-1}^{-1}} \geq \frac{v_{B(t)}}{\sqrt{B(t)}}$) **then**
 - 8: Ask for feedback ($q_t = 1$)
 - 9: Observe R_t and update $V_t = V_{t-1} + x_t x_t^T$ and $\hat{\theta}_t$ according to (11)
 - 10: **end if**
 - 11: **end for**
-

We start by more formally define the linear bandit model with budget constraints. At the beginning of each round t , an adaptive adversary reveals to the learner a budget $B(t)$ (such that $B(t) \geq B(t-1)$) and a context space $\mathcal{X}_t \subset \mathbb{R}^d$ (which serves as the ‘context’ $u_t = \mathcal{X}_t$). Then the learner selects a context $x_t \in \mathcal{X}_t$ and, if enough budget is available, she can choose to query for a noisy reward feedback $R_t = \langle x_t, \theta^* \rangle + \eta_t$, for some unknown $\theta^* \in \mathbb{R}^d$ such that $\langle x, \theta^* \rangle \in [-1, 1]$ for all $x \in \mathcal{X}_t, t \geq 1$. In this section, we assume that querying rewards incur unit costs. the noise is assumed to be zero-meaned and conditionally σ^2 -subgaussian. We also assume that for all $\|x\|_2 \leq L, \forall x \in \mathcal{X}_t$ and that $\|\theta^*\|_2 \leq D$. The (pseudo) regret in this setting is defined as $\text{Reg}(T) = \sum_{t=1}^T \max_{x \in \mathcal{X}_t} \langle x, \theta^* \rangle - \langle x_t, \theta^* \rangle$.

Define $V_t = \lambda I_d + \sum_{k=1}^t x_k x_k^T \mathbf{1}\{q_k = 1\}$ for some $\lambda > 0$. At the end of each round, the algorithm calculates the regularized least-squares estimator for θ^* over the queried rewards, namely

$$\hat{\theta}_t = V_t^{-1} \sum_{k=1}^t x_k R_k \mathbf{1}\{q_k = 1\}. \quad (11)$$

Then, we define the confidence set $C_t = \left\{\theta \in \mathbb{R}^d : \|\hat{\theta}_t - \theta\|_{V_t} \leq l_t\right\}$, where $\|x\|_A = \sqrt{x^T A x}$ and

$$l_t = \max\left\{1, \sigma \sqrt{2d \log\left(\frac{1 + tL^2/\lambda}{\delta}\right)} + \lambda^{1/2}D\right\}.$$

As a result, and since $\langle x_t, \theta^* \rangle \in [-1, 1]$, the CI at the beginning of each round is

$$\langle x_t, \theta^* \rangle \in \left[\max\left\{\min_{\theta \in C_{t-1}} \langle x_t, \theta \rangle, -1\right\}, \min\left\{\max_{\theta \in C_{t-1}} \langle x_t, \theta \rangle, 1\right\} \right],$$

and its width is upper bounded by

$$\begin{aligned} \min\left\{\max_{\theta \in C_{t-1}} \langle x_t, \theta \rangle - \min_{\theta \in C_{t-1}} \langle x_t, \theta \rangle, 2\right\} &= \min\left\{\max_{\theta \in C_{t-1}} \langle x_t, \theta - \hat{\theta}_t \rangle - \min_{\theta \in C_{t-1}} \langle x_t, \theta - \hat{\theta}_t \rangle, 2\right\} \\ &= 2 \min\left\{l_{t-1} \|x_t\|_{V_{t-1}^{-1}}, 1\right\} && \text{(By Lemma 33)} \\ &\leq 2l_{t-1} \min\left\{\|x_t\|_{V_{t-1}^{-1}}, 1\right\} && (l_t \geq 1) \\ &\triangleq CI_t(x_t). \end{aligned}$$

Finally, we say that the CBM condition queries reward if $CI_t(x_t) \geq \frac{l_{t-1}v_{B(t)}}{2\sqrt{B(t)}}$ for

$$v_t = \sqrt{2d \log\left(1 + \frac{tL^2}{d\lambda}\right)}.$$

We now prove the regret bound of Theorem 3 when setting $\lambda = \max\{D^{-1/2}, 1\}$:

Theorem 3 (Confidence Budget Matching for Linear Bandits). *For any adaptive adversarially chosen sequence of non-decreasing budget and context set $\{B(t), u_t\}_{t \geq 1}$ the regret of CBM-OFUL is upper bounded by*

$$\begin{aligned} \text{Reg}(T) &\leq 2l_T v_T \left(\sum_{t=1}^T \frac{1}{\sqrt{B(t)}} + \sqrt{T} \right) \\ &= \mathcal{O}\left((d\sigma + \sqrt{d\lambda}D) \log\left(\frac{1+tL^2/\lambda}{\delta}\right) \left(\sum_{t=1}^T \frac{1}{\sqrt{B(t)}} + \sqrt{T} \right) \right). \end{aligned}$$

for any $T \geq 1$ with probability greater than $1 - \delta$, where the \mathcal{O} -notation holds when $d\sigma + \sqrt{d\lambda}D \geq 1$.

Proof. With probability greater than $1 - \delta$ it holds that $\theta^* \in C_t = \{\theta \in \mathbb{R}^d : \|\hat{\theta}_t - \theta\|_{V_t} \leq l_t\}$, uniformly for all $t \geq 0$ (see Theorem 7 which generalizes (Abbasi-Yadkori et al., 2011, Theorem 2)). We define this event as the good event and denote it by \mathbb{G} . Also, we denote the parameter vector that maximizes the UCB by $\tilde{\theta}_t$, i.e., $(x_t, \tilde{\theta}_t) \in \arg \max_{x \times \theta \in \mathcal{X}_t \times C_{t-1}} \langle x, \theta \rangle$. Then, under the good event we have that $\max_{x \in \mathcal{X}_t} \langle x, \theta^* \rangle \leq \langle x_t, \tilde{\theta}_t \rangle$ ('optimism').

As in the MAB setting (Theorem 2), we prove that when CBM-OFUM asks for a query, it always has sufficient budget (Lemma 12). Therefore, throughout the proof, we assume that the algorithm observes reward iff it sets $q_t = 1$ and that $B^q(t) \leq \min\{T, B(t)\}$ for all t .

We now derive the performance bound by bounding the regret of the algorithm. The following relations hold for any $T \geq 1$:

$$\text{Reg}(T) = \underbrace{\sum_{t=1}^T 1\{q_t = 1\} \left(\max_{x \in \mathcal{X}_t} \langle x, \theta^* \rangle - \langle x_t, \theta^* \rangle \right)}_{(i)} + \underbrace{\sum_{t=1}^T 1\{q_t = 0\} \left(\max_{x \in \mathcal{X}_t} \langle x, \theta^* \rangle - \langle x_t, \theta^* \rangle \right)}_{(ii)}.$$

Terms (i) and (ii) represent the regret over episodes in which feedback was queried and not queried, respectively. We bound each of the terms.

Bound on term (i), episodes in which feedback is queried. The following relations hold.

$$\begin{aligned} (i) &= \sum_{t=1}^T 1\{q_t = 1\} \left(\max_{x \in \mathcal{X}_t} \langle x, \theta^* \rangle - \langle x_t, \theta^* \rangle \right) \\ &= \sum_{t=1}^T 1\{q_t = 1\} \min\left\{ \left(\max_{x \in \mathcal{X}_t} \langle x, \theta^* \rangle - \langle x_t, \theta^* \rangle \right), 2 \right\} && \text{(Bounded expected reward)} \\ &\leq \sum_{t=1}^T 1\{q_t = 1\} \min\left\{ \left(\langle x_t, \tilde{\theta}_t \rangle - \langle x_t, \theta^* \rangle \right), 2 \right\} && \text{(Optimism)} \\ &= \sum_{t=1}^T 1\{q_t = 1\} \min\left\{ \left(\langle x_t, \tilde{\theta}_t - \hat{\theta}_t \rangle + \langle x_t, \hat{\theta}_t - \theta^* \rangle \right), 2 \right\} \\ &= \sum_{t=1}^T 1\{q_t = 1\} \min\left\{ \|x_t\|_{V_{t-1}^{-1}} \left(\|\tilde{\theta}_t - \hat{\theta}_t\|_{V_{t-1}} + \|\hat{\theta}_t - \theta^*\|_{V_{t-1}} \right), 2 \right\} && \text{(Cauchy-Schwartz inequality)} \\ &\leq 2l_T \sum_{t=1}^T 1\{q_t = 1\} \min\left\{ \|x_t\|_{V_{t-1}^{-1}}, 1 \right\} && \text{(Conditioned on } \mathbb{G} \text{ \& } l_t \geq 1 \text{ is increasing in } t) \end{aligned}$$

Next, denote the k^{th} time that CBM-OFUL queried a reward by τ_k . Then, we can write $V_t = \lambda I_d + \sum_{k=1}^{B^q(t)} x_{\tau_k} x_{\tau_k}^T$. Rewriting the bound on (i) and then using the the elliptical potential lemma (Lemma 30) and get

$$\begin{aligned}
 (i) &\leq 2l_T \sum_{k=1}^{B^q(T)} \min\left\{\|x_{\tau_k}\|_{V_{\tau_k}^{-1}}, 1\right\} \\
 &\leq 2l_T \sqrt{B^q(T)} \sqrt{\sum_{k=1}^{B^q(T)} \min\left\{\|x_{\tau_k}\|_{V_{\tau_k}^{-1}}^2, 1\right\}} && ((\sum_{i=1}^n x_i)^2 \leq n \sum_{i=1}^n x_i^2) \\
 &\leq 2l_T \sqrt{B^q(T)} \sqrt{2d \log\left(1 + \frac{B^q(T)L^2}{d\lambda}\right)} && \text{(By Lemma 30)} \\
 &\leq 2l_T \sqrt{T} \sqrt{2d \log\left(1 + \frac{TL^2}{d\lambda}\right)}.
 \end{aligned}$$

Bound on term (ii), episodes in which feedback is not queried. The following relations hold.

$$\begin{aligned}
 (ii) &= \sum_{t=1}^T 1\{q_t = 0\} \left(\max_{x \in \mathcal{X}_t} \langle x, \theta^* \rangle - \langle x_t, \theta^* \rangle\right) \\
 &\leq \sum_{t=1}^T 1\{q_t = 0\} \min\left\{\langle x_t, \tilde{\theta}_t \rangle - \langle x_t, \theta^* \rangle, 2\right\} && \text{(Optimism)} \\
 &\leq \sum_{t=1}^T 1\{q_t = 0\} \underbrace{\min\left\{\max_{\theta \in C_t} \langle x_t, \theta \rangle - \min_{\theta \in C_t} \langle x_t, \theta \rangle, 2\right\}}_{\leq CI_t(x_t)} \\
 &\leq \sum_{t=1}^T 1\{q_t = 0\} \frac{2l_{t-1}v_{B(t)}}{\sqrt{B(t)}} && \text{(By update rule)} \\
 &\leq 2l_T v_T \sum_{t=1}^T \frac{1}{\sqrt{B(t)}}. && (l_t, v_t \text{ are increasing in } t)
 \end{aligned}$$

The forth relation holds since the algorithm does not query for feedback *only* when $CI_t(x_t) \leq 2l_{t-1}v_{B(t)}/\sqrt{B(t)}$ (a case where $q_t = 0$ since the algorithm ran out of budget cannot happen).

Combining the bounds. Combining the bounds over (i) and (ii) we get that conditioned on the good event (which holds with probability greater than $1 - \delta$), the regret is bounded by

$$\begin{aligned}
 \text{Reg}(T) &\leq 2l_T v_T \sum_{t=1}^T \frac{1}{\sqrt{B(t)}} + 2l_T \sqrt{T} \underbrace{\sqrt{2d \log\left(1 + \frac{TL^2}{d\lambda}\right)}}_{\stackrel{\text{def}}{=} v_T} \\
 &= 2l_T v_T \left(\sum_{t=1}^T \frac{1}{\sqrt{B(t)}} + \sqrt{T}\right),
 \end{aligned}$$

for all $T \geq 1$. □

D.1. CBM-OFUL: Budget Constraint is Satisfied

Lemma 12 (CBM-OFUL: Budget Constraint is Satisfied). *Conditioned on the good event, for any $T \geq 1$ the budget constraint is not violated, $B^q(T) \leq B(T)$.*

Proof. Before supplying the proof, notice that the function $f(x) = x/\log(1 + x\alpha)$ is a strictly increasing function in $\alpha > 0, x \geq 0$ (since $f'(x) > 0$ for $x > 0$). This implies that the ratio

$$\frac{B(T)}{v_{B(T)}^2} = \frac{B(T)}{4d \log(1 + B(T)L^2/\lambda)} = \frac{f(B(T))}{4d}$$

is a non-decreasing function in T since $B(T)$ is non-decreasing in T . Also, recall that $CI_t(x_t) = 2l_t \min\{\|x_t\|_{V_t^{-1}}, 1\}$. Using these facts we get the following relations for any $T \geq 1$.

$$\begin{aligned} B^q(T) &= \sum_{t=1}^T 1\{q_t = 1\} \\ &\leq \sum_{t=1}^T 1\{q_t = 1\} \frac{CI_t(x_t)}{2l_{t-1}v_{B(t)}\sqrt{1/B(t)}} && \text{(Update rule)} \\ &= \sum_{t=1}^T 1\{q_t = 1\} \frac{2l_{t-1} \min\{\|x_t\|_{V_{t-1}^{-1}}, 1\}}{2l_{t-1}v_{B(t)}\sqrt{1/B(t)}} && \text{(Definition of } CI_t(x_t)\text{)} \\ &\leq \sqrt{B(T)/v_{B(T)}} \sum_{t=1}^T 1\{q_t = 1\} \min\{\|x_t\|_{V_{t-1}^{-1}}, 1\} && (B(T)/v_{B(T)}^2 \text{ is non-decreasing in } T) \\ &\leq \sqrt{B(T)/v_{B(T)}} \sqrt{\sum_{t=1}^T 1\{q_t = 1\}} \sqrt{\sum_{t=1}^T 1\{q_t = 1\} \min\{\|x_t\|_{V_{t-1}^{-1}}^2, 1\}} && \text{(Cauchy-Schwartz Inequality)} \\ &\leq \sqrt{B(T)/v_{B(T)}} \sqrt{B^q(T)} \sqrt{2d \log\left(1 + \frac{B^q(T)L^2}{d}\right)} && \text{(Lemma 30)} \\ &= \sqrt{\frac{B(T)}{\log\left(1 + \frac{B(T)L^2}{d\lambda}\right)}} \sqrt{B^q(T) \log\left(1 + \frac{B^q(T)L^2}{d}\right)}. \end{aligned}$$

Rearranging we get

$$\frac{B^q(T)}{\log(1 + B^q(T)L^2/\lambda)} \leq \frac{B(T)}{\log(1 + B(T)L^2/\lambda)}. \quad (12)$$

Since $f(x) = x/\log(1 + x\alpha)$ is strictly increasing function for any $\alpha > 0, x \geq 0$, (12) implies that $B^q(T) \leq B(T)$. \square

E. Confidence Budget Matching for Reinforcement Learning: CBM-UCBVI

We start by defining the feedback model of RL with budget constraints and by introducing some notations. At the beginning of each round t , the learner acts with a non-stationary policy $\pi : \mathcal{S} \times [H] \rightarrow \mathcal{A}$ for h time steps. The learner observes the trajectory of the t^{th} episode, $\{(s_{t,h}, a_{t,h})\}_{h=1}^H$, and an adaptive adversary reveals to the learner the budget $B(t)$. Then, the learner is allowed to query for reward feedback on the states observed at the t^{th} episode. That is, the learner can ask for a noisy version of the reward $R_{t,h}(s, a) = r_h(s, a) + \eta_t$ where $\mathbb{E}[R_{t,h}(s, a)] = r_h(s, a)$ and $R_{t,h}(s, a) \in [0, 1]$ a.s. as long as (i) (s, a) was observed at the h^{th} time step at the t^{th} episode, i.e., $s_{t,h} = s, a_{t,h} = a$, and, (ii) there is a spare budget. In this section, we assume that querying rewards incur a unit costs and we denote $q_{t,h} = 1$ as the event reward feedback is queried in the t^{th} episode at the h^{th} time step (which implies there is an available budget for this event to occur). Furthermore, we define the set

$$\mathcal{L}_R \stackrel{\text{def}}{=} \{(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H] : r_h(s, a) > 0\}, \quad (13)$$

that is, the set of state, action and time step tuples such that the reward is not zero¹.

The CBM-UCBVI algorithm (Algorithm 5) combines the CBM principle into the UCBVI algorithm (Azar et al., 2017). Similarly to UCBVI, CBM-UCBVI solves an optimistic MDP, $\mathcal{M}_t = (\mathcal{S}, \mathcal{A}, \bar{r}_{t-1} + b_t^r + b_t^p, \bar{P}_{t-1}, H)$, at the beginning of each round, where $\bar{r}_{t-1}, \bar{P}_{t-1}$ are the empirical reward and transition model and $b_t^r + b_t^p$ are bonus terms (all defined below). CBM-UCBVI interacts with the environment with the optimal policy of \mathcal{M}_t and samples a trajectory of state-action pairs $\{(s_{t,h}, a_{t,h})\}_{h=1}^H$. Then, instead of receiving the reward feedback on all the state-action pairs within the trajectory, CBM-UCBVI utilizes the CBM principle to decide in which state-action pairs reward feedback is queried. Specifically, it queries for reward feedback at $(s_{t,h}, a_{t,h})$ if

$$2b_{t,h}^r(s_{t,h}, a_{t,h}) \triangleq CI_{t,h}^r(s_{t,h}, a_{t,h}) \geq 2L_{t,\delta} \left(\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + \frac{SAH \log(1 + B(t))}{B(t)} \right).$$

where $L_{t,\delta} \stackrel{\text{def}}{=} \log\left(\frac{12S^2AHt^2(t+1)}{\delta}\right)$. Indeed, Lemma 16 establishes that this query rule does not violate the budget constraint. Thus, if reward feedback is queried, there is an available budget.

Let $n_{t,h}(s, a) = \sum_{k=1}^t 1\{s_{k,h} = s, a_{k,h} = a\}$ be the number of times (s, a) was sampled at the h^{th} time step until the end of the t^{th} episode, and, $n_{t,h}^q(s, a) = \sum_{k=1}^t 1\{s_{k,h} = s, a_{k,h} = a, q_{k,h} = 1\}$ be the number of times the learner queried for reward feedback in (s, a) at the h^{th} time step until the end of the t^{th} episode. We denote by \bar{r}_t, \bar{P}_t the empirical reward and empirical transition model, that is

$$\begin{aligned} \bar{P}_{t,h}(s'|s, a) &= \frac{1}{n_{t,h}(s, a) \vee 1} \sum_{k=1}^t 1\{s_{k,h+1} = s', s_{k,h} = s, a_{k,h} = a\}, \\ \bar{r}_{t,h} &= \frac{1}{n_{t-1,h}^q(s, a) \vee 1} \sum_{k=1}^t 1\{s_{k,h} = s, a_{k,h} = a, q_{k,h} = 1\} R_{t,h}(s, a). \end{aligned}$$

Observe that, unlike in classic RL, $n_{t,h}(s, a) \neq n_{t,h}^q(s, a)$; the number of times *reward feedback* was queried in (s, a) for the h^{th} time step is not equal to the number of times (s, a) was *visited* at the h^{th} time step. Lastly, the bonus terms b_t^r, b_t^p used by CBM-UCBVI are given as follows.

$$\begin{aligned} b_{t,h}^r(s, a) &\stackrel{\text{def}}{=} \sqrt{\frac{2\widehat{\text{Var}}_{R,t-1,h}(s, a)L_{t,\delta}}{n_{t-1,h}^q(s, a) \vee 1}} + \frac{5L_{t,\delta}}{n_{t-1,h}^q(s, a) \vee 1}, \\ b_{t,h}^p(s, a) &\stackrel{\text{def}}{=} \sqrt{\frac{2H^2L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{5HL_{t,\delta}}{n_{t-1,h}(s, a)}, \end{aligned}$$

¹A bit more generally, our results holds would we define the set of (s, a, h) with deterministic reward, $\{(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H] : \text{Var}(R_h(s, a)) = 0\}$. However, since sparse reward is a more natural measure, we chose to work with the set that defined in (13).

Algorithm 5 CBM-UCBVI

Require: $\delta \in (0, 1)$
for $t = 1, 2, \dots$ **do**

$$\mathcal{M}_t = (\mathcal{S}, \mathcal{A}, \bar{r}_{t-1} + b_t^r + b_t^p, \bar{P}_{t-1}, H)$$

 Get optimistic Q-functions $\{\bar{Q}_{t,h}(s, a)\}_{(s,a) \in \mathcal{S} \times \mathcal{A}, h \in [H]}$ via truncated value iteration on \mathcal{M}_t (Algorithm 6)

 Act with $\pi_{t,h} = \arg \max_a \bar{Q}_{t,h}(s, a)$ and observe a trajectory $\{(s_{t,h}, a_{t,h})\}_{h=1}^H$

 Observe current budget $B(t)$

 Ask for feedback on $(s_{t,h}, a_{t,h}) \in \{(s_{t,h}, a_{t,h})\}_{h=1}^H$ if

$$CI_{t,h}^R(s_{t,h}, a_{t,h}) \geq L_{t,\delta} \left(6\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + 4SAH \frac{\log(1 + B(t)) + 1}{B(t)} \right).$$

end for
Algorithm 6 Truncated Value Iteration

Require: An MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, r, P, H)$.

Initialize: $V_{H+1}(s) = 0$ for all $s \in \mathcal{S}$.

for $h = H, H-1, \dots, 1$ **do**

 For all $(s, a) \in \mathcal{S} \times \mathcal{A}$, $Q_h(s, a) = r_h(s, a) + \mathbb{E}_{P_h}[V_{h+1}(s') | s_h = s, a_h = a]$

 For all $s \in \mathcal{S}$, $V_h(s) = \min\{\max_{a \in \mathcal{A}} Q_h(s, a), H - h\}$
end for
Return: $\{Q(s, a)\}_{(s,a) \in \mathcal{S} \times \mathcal{A}, h \in [H]}$

 where $L_{t,\delta} \stackrel{\text{def}}{=} \log\left(\frac{12S^2AHt^2(t+1)}{\delta}\right)$ and

$$\widehat{\text{Var}}_{R,t-1,h}(s, a) = \begin{cases} \frac{\sum_{k,k'=1}^{n_{t-1,h}^q(s,a)} (R_{k',h}(s,a) - R_{k,h}(s,a))^2}{2(n_{t-1,h}^q(s,a)(n_{t-1,h}^q(s,a)-1) \vee 1)} & n_{t-1,h}^q(s, a) \geq 2 \\ 0 & o.w. \end{cases}$$

is the (unbiased) empirical estimate of the reward's variance.

Notations for the Proof Let $f : \mathcal{S} \rightarrow \mathbb{R}$ and $P(s'|s, a)$ be a transition model. We use the following notation

$$\mathbb{E}_{P_h(\cdot|s,a)}[f(s')] = \sum_{s'} P_h(s'|s, a) f(s').$$

 Furthermore, we define $F_{t,h-1}$ as the σ -algebra generated by all the event until the h^{th} time step within the t^{th} episode; namely, $s_{t,h}$ and $a_{t,h}$ are $F_{t,h-1}$ -measurable, (but not $R_{t,h}$, which we assume that is generated by the end of the episode). Thus, this definition implies that $F_{t,0} = F_{t-1}$ (where F_{t-1} is defined in Section 2). Observe that this definition also implies that

$$\mathbb{E}_{P_h(\cdot|s_{t,h}, a_{t,h})}[f(s')] = \mathbb{E}[f(s_{t,h+1}) | F_{t,h-1}], \quad (14)$$

 for f which is an $F_{t,h-1}$ measurable function (i.e., that depends on the history until time step h at the t^{th} episode).

We are now ready to state the central result of this section which gives a performance guarantee on the regret of CBM-UCBVI.

Theorem 4 (CBM-UCBVI). *For any adversarially adaptive sequence $\{B(t), s_{t,1}\}_{t \geq 1}$ of budget and initial states the regret of CBM-ULCBVI is upper bounded by*

$$\text{Reg}(T) \leq 36L_{T,\delta} \left(\sqrt{SAH^4T} + \sum_{t=1}^T \left(\sqrt{\frac{|\mathcal{L}_R|H^2}{B(t)}} + SAH^2 \frac{\log(1 + B(t)) + 1}{B(t)} \right) \right) + 306H^3S^2AL_{T,\delta}^2,$$

 for any $T \geq 1$ with probability greater than $1 - \delta$.

To establish to proof of this result we prove some preliminary results. We first define the set of good events, which hold with high probability (Section E.1 and Section E.3), and establish the optimism of CBM-UCBVI (Section E.2). Furthermore, in Section E.4, we establish that CBM-UCBVI does not violate the budget constraint. Later, this important property allows us to bound the reward bonus when reward feedback is not queried (see Lemma 18). Given these tools, we prove Theorem 4 (Section E.5) by relying on a key recursion lemma that bounds the on-policy errors at time step h by the on-policy errors of time step $h + 1$.

E.1. The First Good Event - Concentration Events

$$\begin{aligned}
 E^r(t) &= \left\{ \forall s \in \mathcal{S}, a \in \mathcal{A} : |\bar{r}_{t-1,h}(s, a) - r_h(s, a)| \leq \sqrt{\frac{2\widehat{\text{Var}}_{R,t-1,h}(s, a) \log \frac{12SAHt^2(t+1)}{\delta}}{n_{t-1,h}^q(s, a) \vee 1}} + \frac{5 \log \frac{12SAHt^2(t+1)}{\delta}}{n_{t-1,h}^q(s, a) \vee 1} \right\} \\
 E^p(t) &= \left\{ \forall s, s' \in \mathcal{S}, a \in \mathcal{A} : |P_h(s'|s, a) - \bar{P}_{t,h}(s'|s, a)| \leq \sqrt{\frac{2P(s'|s, a) \log \frac{12S^2AHt^2(t+1)}{\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{2 \log \frac{12S^2AHt^2(t+1)}{\delta}}{n_{t-1,h}(s, a) \vee 1} \right\} \\
 E^{pv}(t) &= \left\{ \forall s, a, h : \left| (\bar{P}_{t,h}(\cdot | s, a) - P_h(\cdot | s, a))^T V_{h+1}^* \right| \leq \sqrt{\frac{2H^2 \log \frac{12SAHt^2(t+1)}{\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{5H \log \frac{12SAHt^2(t+1)}{\delta}}{n_{t-1,h}(s, a)} \right\}
 \end{aligned}$$

Notice that the bonus $b_{t,h}^{pv}(s, a)$ depends on the number of times we visited the state s and took action a , denoted by n , whereas the bonus $b_{t,h}^r(s, a)$ depends on the number of times we queried a reward, denoted by n^q . Proving this set of events hold jointly is standard, based upon the empirical Bernstein concentration bound (Maurer & Pontil, 2009).

Lemma 13 (The First Good Event). *Let $\mathbb{G}_1 = \cap_{t \geq 1} E^r(t) \cap_{t \geq 1} E^p(t) \cap_{t \geq 1} E^{pv}(t)$ be the good event. It holds that $\Pr(\mathbb{G}_1) \geq 1 - \delta/2$.*

Proof. We prove that each part of the good event holds with a probability of at least $1 - \delta/6$.

The event $\cap_{t \geq 1} E^r(t)$ holds with high probability. Fix an episode $t \geq 1$ and $s, a, h \in \mathcal{S} \times \mathcal{A} \times [H]$. There are at most t reward samples from (s, a) at time step h at episode t . Taking a union bound over these possible values and scaling $\delta \rightarrow \delta/t$ we get that

$$\Pr \left(|\bar{r}_{t-1,h}(s, a) - r_h(s, a)| \leq \sqrt{\frac{2\widehat{\text{Var}}_{R,t-1,h}(s, a) \log \frac{2t}{\delta}}{n_{t-1,h}^q(s, a) \vee 1}} + \frac{14 \log \frac{2t}{\delta}}{3n_{t-1,h}^q(s, a) \vee 1} \right) \geq 1 - \delta,$$

by (Maurer & Pontil, 2009, Theorem 4). Note that we used the relation $1/(n-1) \leq 2/n$ for $n \geq 2$ and that the bound trivially holds for $n < 2$ since the reward is in $[0, 1]$. Furthermore, taking a union bound on all $s, a, h \in \mathcal{S} \times \mathcal{A} \times [H]$ and setting $\delta \rightarrow \delta/6SAHt(t+1)$ results in $\Pr(\cap_{t \geq 1} E^r(t)) \geq 1 - \delta/6$ by using $\sum_{t=1}^{\infty} \frac{\delta}{t(t+1)} = \delta$.

The event $\cap_{t \geq 1} E^p(t)$ holds with high probability. Fix an episode $t \geq 1$ and $s, a, s', h \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$. At the t^{th} episode the tuple can be sampled for at most t times. Taking a union bound on all these possible values and applying Bennet's inequality (e.g., Maurer & Pontil, 2009, Theorem 3), we get

$$\Pr \left(|P_h(s'|s, a) - \bar{P}_{t,h}(s'|s, a)| \leq \sqrt{\frac{2P(s'|s, a) \log \frac{2t}{\delta}}{n_{t-1}(s, a) \vee 1}} + \frac{4 \log \frac{2t}{\delta}}{3n_{t-1}^p(s, a) \vee 1} \right) \geq 1 - \delta,$$

since the variance of a Bernoulli random variable is $P(s'|s, a)(1 - P(s'|s, a)) \leq P(s'|s, a)$. Taking a union bound over all $s, a, s', h \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$ and scaling $\delta \rightarrow \delta/6S^2AHt(t+1)$ and repeating the same reasoning as before concludes the high probability bound for $\cap_{t \geq 1} E^p(t)$.

The event $\cap_{t \geq 1} E^{pv}(t)$ holds with high probability. Repeating the same arguments as for the event $\cap E^r(t)$ while noticing that $\bar{P}_{t,h}(\cdot | s, a)^T V_{h+1}^* \in [0, H]$ (and, thus, also the empirical variance, for all $n \geq 2$) concludes the proof.

Combining the results. Taking a union bound concludes the proof. \square

E.2. Optimism

We can prove that the value is optimistic using standard techniques.

Lemma 14 (Value Function is Optimistic). *Conditioning on the first good event \mathbb{G}_1 the value function of CBM-UCBVI is optimism for all $s \in \mathcal{S}, h \in [H], t \geq 1$, i.e.,*

$$\forall s \in \mathcal{S}, h \in [H], t \geq 1 : V_h^*(s) \leq \bar{V}_{t,h}(s).$$

Proof. We prove this result via induction.

Base case, $h = H$ and for all $s \in \mathcal{S}$. Let $a^*(s) \in \arg \max_{a \in \mathcal{A}} r_H(s, a)$. For any $s \in \mathcal{S}$ it holds that

$$V_H^*(s) - \bar{V}_{t,H}(s) = r_H(s, a^*(s)) - \min \left\{ \max_a \{ \bar{r}_{t,H}(s, a) + b_{t,H}^r(s, a) \}, 1 \right\}. \quad (15)$$

Assume that $\max_a \bar{r}_{t,H}(s, a) + b_{t,H}^r(s, a) < 1$ then,

$$\begin{aligned} (15) &\leq r_H(s, a^*(s)) - \bar{r}_{t,H}(s, a^*(s)) - b_{t,H}^r(s, a^*(s)) \\ &\leq b_{t,H}^r(s, a^*(s)) - b_{t,H}^r(s, a^*(s)) \leq 0. \end{aligned} \quad (\text{Event } \cap_t E^r(t) \text{ holds})$$

If $\max_a \{ \bar{r}_{t,H}(s, a) + b_{t,H}^r(s, a) \} \geq 1$ then trivially (15) ≤ 0 since $r_H(s, a) \in [0, 1]$ for all $a \in \mathcal{A}$. Overall, we conclude that $V_H^*(s) \leq \bar{V}_{t,H}(s)$ for all $s \in \mathcal{S}$ for $h = H$.

Induction step, for $h \in [H]$ $s \in \mathcal{S}$ assuming it holds for all $h' \geq h + 1$. Let $a^*(s) \in \arg \max_{a \in \mathcal{A}} Q_h^*(s, a)$ The following relations hold.

$$V_h^*(s) - \bar{V}_{t,h}(s) = Q_h^*(s, a^*(s)) - \min \left\{ \max_a \bar{Q}_{t,h}(s, a), H - h \right\} \quad (16)$$

Assume that $\max_a \bar{Q}_{t,h}^\pi(s, a) < H - h$, then,

$$\begin{aligned} (16) &\leq Q_h^*(s, a^*(s)) - \bar{Q}_{t,h}(s, a^*(s)) \\ &= r_h(s, a^*(s)) + P_h(\cdot | s, a^*(s))^T V_{h+1}^* \\ &\quad - \left(\bar{r}_{t-1,h}(s, a^*(s)) + b_{t,h}^r(s, a^*(s)) + b_{t,h}^p(s, a^*(s)) + \bar{P}_h(\cdot | s, a^*(s))^T \bar{V}_{t,h+1} \right) \\ &= r_h(s, a^*(s)) - \bar{r}_{t-1,h}(s, a^*(s)) - b_{t,h}^r(s, a^*(s)) + (P_h - \bar{P}_{t-1,h})(\cdot | s, a^*(s))^T V_{h+1}^* - b_{t,h}^p(s, a^*(s)) \\ &\quad + \underbrace{\mathbb{E}_{\bar{P}_{t-1,h}(\cdot | s, a^*(s))} [V_{h+1}^*(s') - \bar{V}_{t,h+1}(s')]}_{\leq 0 \text{ Induction hypothesis}} \\ &\leq b_{t,h}^r(s, a^*(s)) - b_{t,h}^r(s, a^*(s)) - b_{t,h}^p(s, a^*(s)) + b_{t,h}^p(s, a^*(s)) \quad (\cap_t E^{pv}(t) \cup \cap_t E^r(t) \text{ holds}) \\ &= 0. \end{aligned}$$

If $\max_a \bar{Q}_{t,h}^\pi(s, a) \geq H - h$ then trivially (16) ≤ 0 since $Q_h^*(s, a) \leq H - h$ since it is an expectation over $H - h$ terms, each bounded by $[0, 1]$. Thus, $V_h^*(s) \leq \bar{V}_{t,h}(s)$ for all $s \in \mathcal{S}$ for $h' \geq h$ which proves the induction step and concludes the proof. \square

E.3. The Second Good Event - Optimism Bound

We now prove a high probability bound which holds conditioned on the first good event \mathbb{G}_1 .

Lemma 15 (The Second Good Event). *Let \mathbb{G}_1 be the event defined in Lemma 13. Let $\{Y_{t,h}\}_{t \geq 1}$ the random variables defined as*

$$Y_{t,h} \stackrel{\text{def}}{=} \bar{V}_{t,h+1}(s_{t,h+1}) - V_{h+1}^{\pi_t}(s_{t,h+1})$$

The second good event is defined as $\mathbb{G}_2 = E^O$ where

$$E^O = \left\{ \forall h \in [H-1], T \geq 1 : \sum_{t=1}^T \mathbb{E}[Y_{t,h} | F_{t,h-1}] \leq \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T Y_{t,h} + 18H^2 \log \frac{2HT(T+1)}{\delta} \right\}$$

Then, the good event $\mathbb{G} = \mathbb{G}_1 \cap \mathbb{G}_2$ holds with probability greater than $1 - \delta$.

Proof. Fix $h \in [H], T$. We start by defining the random variable $W_t = 1\{\bar{V}_{t,h+1}(s) - V_h^{\pi_t}(s) \geq 0, \forall h \in [H], s \in \mathcal{S}\}$ and the random process $Y_t = Y_{t,h} = \bar{V}_{t,h+1}(s_{t,h+1}) - V_{h+1}^{\pi_t}(s_{t,h+1})$ w.r.t. the filtration $\{F_{t+1,h-1}\}_{t \geq 1}$ (observe that $Y_{t,h}$ is $F_{t+1,h-1}$ measurable). Also notice that W_t is $F_{t,h-1}$ measurable, as both π_t and $\bar{V}_{t,h}$ are F_{t-1} -measurable. Finally, define $\tilde{Y}_t = W_t Y_t$. Importantly, $\tilde{Y}_t \in [0, H]$ almost surely, by definition of W_t and since $\bar{V}_{t,h+1}(s), V_h^{\pi_t}(s) \in [0, H]$ by the update rule. Thus, using Lemma 27 with $C = H \geq 1$, we get

$$\sum_{t=1}^T \mathbb{E}[\tilde{Y}_{t,h} | F_{t,h-1}] \leq \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T \tilde{Y}_{t,h} + 18H^2 \log \frac{1}{\delta},$$

with probability greater than $1 - \delta$, and since W_t is $F_{t,h-1}$ -measurable, we can write

$$\sum_{t=1}^T W_t \mathbb{E}[Y_{t,h} | F_{t,h-1}] \leq \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T W_t Y_{t,h} + 18H^2 \log \frac{1}{\delta}. \quad (17)$$

Importantly, notice that under \mathbb{G}_1 , it holds that $W_t \equiv 1$ (by Lemma 14). Therefore, applying the union bound and setting $\delta = \delta/2HT(T+1)$ we get

$$\begin{aligned} & \Pr(\overline{E^O} \cap \mathbb{G}_1) \\ & \leq \sum_{h=1}^H \sum_{T=1}^{\infty} \Pr\left(\left\{ \sum_{t=1}^T \mathbb{E}[Y_{t,h} | F_{t,h-1}] \geq \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T Y_{t,h} + 18H^2 \log \frac{2HT(T+1)}{\delta} \right\} \cap \mathbb{G}_1\right) \quad (\text{Union bound}) \\ & = \sum_{h=1}^H \sum_{T=1}^{\infty} \Pr\left(\left\{ \sum_{t=1}^T W_t \mathbb{E}[Y_{t,h} | F_{t,h-1}] \geq \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T W_t Y_{t,h} + 18H^2 \log \frac{2HT(T+1)}{\delta} \right\} \cap \mathbb{G}_1\right) \\ & \quad (W_t \equiv 1 \text{ under } \mathbb{G}_1) \\ & \leq \sum_{h=1}^H \sum_{T=1}^{\infty} \Pr\left(\sum_{t=1}^T W_t \mathbb{E}[Y_{t,h} | F_{t,h-1}] \geq \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T W_t Y_{t,h} + 18H^2 \log \frac{2HT(T+1)}{\delta}\right) \\ & \leq \sum_{h=1}^H \sum_{T=1}^{\infty} \frac{\delta}{2HT(T+1)} = \delta/2, \quad (\text{By (17)}) \end{aligned}$$

Finally, we have

$$\Pr(\overline{\mathbb{G}}) \leq \Pr(\overline{\mathbb{G}_2} \cap \mathbb{G}_1) + \Pr(\overline{\mathbb{G}_1}) \leq \frac{\delta}{2} + \frac{\delta}{2} = \delta$$

□

E.4. CBM-UCBVI: Budget Constraint is Satisfied

Lemma 16 (CBM-UCBVI: Budget Constraint is Satisfied). *For any $T \geq 0$ the budget constraint is not violated, that is $B^q(T) \leq B(T)$ almost surely.*

The number of times feedback was queried after T episodes $B^q(T)$ is by definition $B^q(T) = \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\}$. We prove the lemma by extending the techniques used for CBM-UCB and CBM-OFUL (see Lemma 11 and Lemma 12).

Proof. It holds that

$$\begin{aligned}
 B^q(T) &= \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} \\
 &\leq \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} \frac{CI_{t,h}^R(s_{t,h}, a_{t,h})}{L_{t,\delta} \left(6\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + 4SAH \frac{\log(1+B(t))+1}{B(t)} \right)} \\
 &\leq \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} \frac{\sqrt{\frac{2\widehat{\text{Var}}_{R,t-1,h}(s_{t,h}, a_{t,h})}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1}} + \frac{2}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1}}{6\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + 4SAH \frac{\log(1+B(t))+1}{B(t)}} \quad \left(\frac{\sqrt{L_{t,\delta}}}{L_{t,\delta}}, \frac{L_{t,\delta}}{L_{t,\delta}} \leq 1 \right) \\
 &\leq \left(6\sqrt{\frac{|\mathcal{L}_R|}{B(T)}} + 4SAH \frac{\log(1+B(T))+1}{B(T)} \right)^{-1} \\
 &\quad \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} \left(\sqrt{\frac{2\widehat{\text{Var}}_{R,t-1,h}(s_{t,h}, a_{t,h})}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1}} + \frac{2}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1} \right), \quad (18)
 \end{aligned}$$

where the last relation holds since the budget is non-decreasing in T and both $\frac{1}{\sqrt{x}}$, $\frac{\log(1+x)+1}{x}$ are monotonically decreasing functions for $x > 0$. We now upper bound the sum in the last relation. Recall that up to time t , a state action pair (s, a) at time step h was queried $n_{T,h}^q(s, a)$ times. Also, notice that $\widehat{\text{Var}}_{R,t-1}(s, a) = 0$ if $(s, a, h) \notin \mathcal{L}_R$ and $\widehat{\text{Var}}_{R,t-1}(s, a) \leq 1$ if $(s, a, h) \in \mathcal{L}_R$. Then, we can write

$$\begin{aligned}
 &\sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} \left(\sqrt{\frac{2\widehat{\text{Var}}_{R,t-1,h}(s_{t,h}, a_{t,h})}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1}} + \frac{2}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1} \right) \\
 &\leq \sum_{s,a,h \in \mathcal{L}_R} n_{T,h}^q(s,a) \sum_{i=0}^{n_{T,h}^q(s,a)} \sqrt{\frac{2}{i \vee 1}} + \sum_{s,a,h} n_{T,h}^q(s,a) \sum_{i=0}^{n_{T,h}^q(s,a)} \frac{2}{i \vee 1} \quad \text{(Bounding } \widehat{\text{Var}}_{R,t-1}(s, a) \text{)} \\
 &\leq 6 \sum_{s,a,h \in \mathcal{L}_R} \sqrt{n_{T,h}^q(s,a)} + 2 \sum_{s,a,h} \left(2 + \log(n_{T,h}^q(s,a)) \vee 1 \right) \\
 &\stackrel{(a)}{\leq} 6\sqrt{|\mathcal{L}_R| B^q(T)} + 4SAH \log(1+B^q(T)) + 4SAH, \quad (19)
 \end{aligned}$$

where (a) follows by Jensen's inequality and $\sum_{s,a,h \in \mathcal{L}_R} n_{T,h}^q(s,a) \leq \sum_{s,a,h} n_{T,h}^q(s,a) = B^q(T)$. Plugging this back into (18) we get

$$\begin{aligned}
 B^q(T) &\leq (18) \leq \left(6\sqrt{\frac{|\mathcal{L}_R|}{B(T)}} + 4SAH \frac{\log(1+B(T))+1}{B(T)} \right)^{-1} \left(6\sqrt{|\mathcal{L}_R| B^q(T)} + 4SAH \log(1+B^q(T)) + 4SAH \right) \\
 &\iff 6\sqrt{\frac{|\mathcal{L}_R|}{B(T)}} + 4SAH \frac{\log(1+B(T))+1}{B(T)} \leq 6\sqrt{\frac{|\mathcal{L}_R|}{B^q(T)}} + 4SAH \frac{\log(1+B^q(T))+1}{B^q(T)}. \quad (20)
 \end{aligned}$$

Remember that if $f(x)$ is strictly monotonically decreasing function then $f(x_1) \leq f(x_2) \iff x_1 \geq x_2$. Furthermore, see that $f_{\alpha,\beta}(x) = \frac{\alpha}{\sqrt{x}} + \beta \frac{\log(1+x)+1}{x}$ is a strictly monotonically decreasing function for $x > 0, \alpha \geq 0, \beta > 0$ (since both $\frac{1}{\sqrt{x}}$

and $\frac{\log(1+x)+1}{x}$ are strictly decreasing). We can then write (20), equivalently as

$$f_{\alpha,\beta}(B(T)) \leq f_{\alpha,\beta}(B^q(T)),$$

for $\alpha = 6|\mathcal{L}_R|$, $\beta = 4SAH$, and since $f_{\alpha,\beta}$ is strictly monotonically decreasing it implies that $B^q(T) \leq B(T)$. \square

E.5. Proof of Theorem 4

Before establishing the proof Theorem 4 we establish the following key lemma that bounds the on-policy errors at time step h by the on-policy errors at time step $h+1$ and additional additive terms. Given this result, the analysis follows with relative ease. We are now ready to establish Theorem 4.

Lemma 17 (CBM-UCBVI, Key Recursion Bound). *Conditioned on the good event \mathbb{G} , the following bound holds for all $h \in [H]$.*

$$\begin{aligned} & \sum_{t=1}^T \bar{V}_{t,h}(s_{t,h}) - V_h^{\pi_t}(s_{t,h}) \\ & \leq 27H^2 \log\left(\frac{2HT(T+1)}{\delta}\right) + 2 \sum_{t=1}^T b_{t,h}^r(s_{t,h}, a_{t,h}) + \sum_{t=1}^T 3H \sqrt{\frac{L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + \frac{13H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \\ & \quad + \left(1 + \frac{1}{2H}\right)^2 \sum_{t=1}^T \bar{V}_{t,h+1}(s_{t,h+1}) - V_{h+1}^{\pi_t}(s_{t,h+1}). \end{aligned}$$

Proof. Let $\Delta P_{t-1,h}(\cdot|s, a) = (\bar{P}_{t-1,h} - P_h)(\cdot|s, a)$. We bound each of the terms in the sum as follows.

$$\begin{aligned} & \bar{V}_{t,h}(s_{t,h}) - V_h^{\pi_t}(s_{t,h}) \\ & = \bar{r}_{t-1,h}(s_1, a_{t,h}) + b_{t,h}^r(s_{t,h}, a_{t,h}) - r_h(s_{t,h}, a_{t,h}) + b_{t,h}^p(s_{t,h}, a_{t,h}) \\ & \quad + \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s_{t,h}, a_{t,h})}[\bar{V}_{t,h+1}(s_{h+1})] - \mathbb{E}_{P_h(\cdot|s_{t,h}, a_{t,h})}[V_{h+1}^{\pi_t}(s_{h+1})] \\ & \leq 2b_{t,h}^r(s_{t,h}, a_{t,h}) + b_{t,h}^p(s_{t,h}, a_{t,h}) \\ & \quad + \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s_{t,h}, a_{t,h})}[\bar{V}_{t,h+1}(s_{h+1})] - \mathbb{E}_{P_h(\cdot|s_{t,h}, a_{t,h})}[V_{h+1}^{\pi_t}(s_{h+1})] \quad (\mathbb{G}_1 \text{ holds}) \\ & = 2b_{t,h}^r(s_{t,h}, a_{t,h}) + b_{t,h}^p(s_{t,h}, a_{t,h}) + \mathbb{E}_{P_h(\cdot|s_{t,h}, a_{t,h})}[\bar{V}_{t,h+1}(s_{h+1}) - V_{h+1}^{\pi_t}(s_{h+1})] \\ & \quad + \Delta P_{t-1,h}(\cdot|s_{t,h}, a_{t,h})^T \bar{V}_{t-1,h+1} \\ & = 2b_{t,h}^r(s_{t,h}, a_{t,h}) + b_{t,h}^p(s_{t,h}, a_{t,h}) + \mathbb{E}_{P_h(\cdot|s_{t,h}, a_{t,h})}[\bar{V}_{t,h+1}(s_{h+1}) - V_{h+1}^{\pi_t}(s_{h+1})] \\ & \quad + \Delta P_{t-1,h}(\cdot|s_{t,h}, a_{t,h})^T (\bar{V}_{t,h+1} - V_{h+1}^*) + \underbrace{\Delta P_{t-1,h}(\cdot|s_{t,h}, a_{t,h})^T V_{h+1}^*}_{\leq b_{t,h}^p(s_{t,h}, a_{t,h})} \\ & \leq 2b_{t,h}^r(s_{t,h}, a_{t,h}) + 2b_{t,h}^p(s_{t,h}, a_{t,h}) + \mathbb{E}_{P_h(\cdot|s_{t,h}, a_{t,h})}[\bar{V}_{t,h+1}(s_{h+1}) - V_{h+1}^{\pi_t}(s_{h+1})] \\ & \quad + \Delta P_{t-1,h}(\cdot|s_{t,h}, a_{t,h})^T (\bar{V}_{t,h+1} - V_{h+1}^*) \quad (\mathbb{G}_1 \text{ holds}) \\ & \stackrel{(a)}{\leq} 2b_{t,h}^r(s_{t,h}, a_{t,h}) + 3H \sqrt{\frac{L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + \frac{13H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \\ & \quad + \left(1 + \frac{1}{2H}\right) \mathbb{E}_{P_h(\cdot|s_{t,h}, a_{t,h})}[\bar{V}_{t,h+1}(s_{h+1}) - V_{h+1}^{\pi_t}(s_{h+1})], \quad (21) \end{aligned}$$

where (a) by substituting $b_{t,h}^p(s, a)$ while bounding its last term by $\frac{5H^2 S L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}$ and by applying Lemma 28. Specifically, we set $\alpha = 2H$, $C_1 = C_2 = 2L_{t,\delta}$ and, thus, $HS(C_2 + \alpha C_2/4) \leq 3H^2 S L_{t,\delta}$ and observe that the conditions of Lemma 28 hold since the event $\cap_t E^p(t)$ holds and by the optimism of Lemma 14 under the good event.

Taking the sum over the latter inequality we conclude the proof since

$$\begin{aligned}
 & \sum_{t=1}^T \bar{V}_{t,h}(s_{t,h}) - V_h^{\pi_t}(s_{t,h}) \\
 & \leq 2 \sum_{t=1}^T b_{t,h}^r(s_{t,h}, a_{t,h}) + \sum_{t=1}^T 3H \sqrt{\frac{L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + \frac{13H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \\
 & \quad + \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T \mathbb{E}_{P_h(\cdot|s_{t,h}, a_{t,h})} [\bar{V}_{t,h+1}(s_{h+1}) - V_{h+1}^{\pi_t}(s_{h+1})] \\
 & = 2 \sum_{t=1}^T b_{t,h}^r(s_{t,h}, a_{t,h}) + \sum_{t=1}^T 3H \sqrt{\frac{L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + \frac{13H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \\
 & \quad + \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T \mathbb{E}[\bar{V}_{t,h+1}(s_{h+1}) - V_{h+1}^{\pi_t}(s_{h+1}) | F_{t,h-1}] \\
 & \leq 27H^2 \log\left(\frac{2HT(T+1)}{\delta}\right) + 2 \sum_{t=1}^T b_{t,h}^r(s_{t,h}, a_{t,h}) + \sum_{t=1}^T 3H \sqrt{\frac{L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + \frac{13H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \\
 & \quad + \left(1 + \frac{1}{2H}\right)^2 \sum_{t=1}^T \bar{V}_{t,h+1}(s_{t,h+1}) - V_{h+1}^{\pi_t}(s_{t,h+1}), \quad (\text{The second good event holds, } \mathbb{G}_2 = E^O)
 \end{aligned}$$

where in the last relation we also bounded $18(1 + \frac{1}{2H}) \leq 27$. \square

We are now ready to prove Theorem 4.

Proof. We start by conditioning on the good event \mathbb{G} , which holds with probability greater than $1 - \delta$. Conditioned on the good event, the value is optimistic (Lemma 14). This fact, together with the key recursion lemma (Lemma 17) yields the following.

$$\begin{aligned}
 & \sum_{t=1}^T V_1^*(s_{t,1}) - V_1^{\pi_t}(s_{t,1}) \leq \sum_{t=1}^T \bar{V}_{t,1}(s_{t,1}) - V_1^{\pi_t}(s_{t,1}) \\
 & \leq 27H^2 \log\left(\frac{2HT(T+1)}{\delta}\right) + \sum_{t=1}^T 2b_{t,1}^r(s_{t,1}, a_{t,1}) + 3H \sqrt{\frac{L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + \frac{13H^2 S L_{t,\delta}}{n_{t-1,1}(s_{t,1}, a_{t,1}) \vee 1} \\
 & \quad + \left(1 + \frac{1}{2H}\right)^2 \sum_{t=1}^T \bar{V}_{t,2}(s_{t,h+1}) - V_2^{\pi_t}(s_{t,2}). \quad (\text{Lemma 17})
 \end{aligned}$$

Iterating on this relation over $h \in \{2, \dots, H\}$ and using $(1 + \frac{1}{2H})^{2h} \leq e \leq 3$ for $h \leq H$ and $\bar{V}_{t,H+1}(s) = V_{H+}^{\pi_t}(s) = 0$, we conclude that

$$\begin{aligned}
 & \sum_{t=1}^T V_1^*(s_{t,1}) - V_1^{\pi_t}(s_{t,1}) \leq 81H^3 \log\left(\frac{2HT(T+1)}{\delta}\right) + 6 \sum_{t=1}^T \sum_{h=1}^H b_{t,h}^r(s_{t,h}, a_{t,h}) \\
 & \quad + \sum_{t=1}^T \sum_{h=1}^H 9H \sqrt{\frac{L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + \frac{39H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}. \quad (22)
 \end{aligned}$$

The first sum in (22) is bounded in Lemma 18 as

$$\begin{aligned}
 & 6 \sum_{t=1}^T \sum_{h=1}^H b_{t,h}^r(s_{t,h}, a_{t,h}) \\
 & \leq 3L_{t,\delta} \left(6\sqrt{|\mathcal{L}_R|HT} + 10SAH \log(HT) + 23SAH + \sum_{t=1}^T \left(6\sqrt{\frac{|\mathcal{L}_R|H^2}{B(t)}} + 4SAH \frac{\log(1+B(t))+1}{B(t)} \right) \right).
 \end{aligned}$$

The second sum in (22) is bounded via the following standard analysis as follows.

$$\begin{aligned}
 \sum_{t=1}^T \sum_{h=1}^H \frac{1}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} &= \sum_{s,a,h} \sum_{i=0}^{n_{T,h}(s,a)} \frac{1}{i \vee 1} && \text{(Rewriting the summation)} \\
 &\leq \sum_{s,a,h} (2 + \log(n_{T,h}(s,a) \vee 1)) \\
 &\leq SAH(2 + \log(TH)) && (23)
 \end{aligned}$$

where the last relation holds by Jensen's inequality while using $\sum_{s,a,h} n_{T,h}(s,a) = TH$. Similarly, we bound

$$\begin{aligned}
 \sum_{t=1}^T \sum_{h=1}^H \frac{1}{\sqrt{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} &= \sum_{s,a,h} \sum_{i=0}^{n_{T,h}(s,a)} \frac{1}{\sqrt{i \vee 1}} && \text{(Rewriting the summation)} \\
 &\leq \sum_{s,a,h} \left(1 + 2\sqrt{n_{T,h}(s,a)}\right) \\
 &\leq SAH + 2\sqrt{SAH^2 T} && (24)
 \end{aligned}$$

Thus, the second sum in (22) is bounded by

$$\begin{aligned}
 &\sum_{t=1}^T \sum_{h=1}^H 9H \sqrt{\frac{L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + \sum_{t=1}^T \sum_{h=1}^H \frac{39H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \\
 &\leq 9H \sqrt{L_{T,\delta}} \sum_{t=1}^T \sum_{h=1}^H \sqrt{\frac{1}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} + 39H^2 S L_{T,\delta} \sum_{t=1}^T \sum_{h=1}^H \frac{1}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \quad (L_{t,\delta} \text{ is increasing in } t) \\
 &\leq 18\sqrt{SAH^4 T L_{T,\delta}} + 39H^3 S^2 A L_{T,\delta} \log(TH) + 87H^3 S^2 A L_{T,\delta}. && \text{(By (23), (24) and since } L_{T,\delta} \geq 1)
 \end{aligned}$$

Plugging the bounds on the first and second sums in (22) concludes the proof. \square

Lemma 18 (Bound on Cumulative Reward Bonus). *The following bound holds almost surely*

$$\begin{aligned}
 \sum_{t=1}^T \sum_{h=1}^H C I_{t,h}^r(s_{t,h}, a_{t,h}) &= 2 \sum_{t=1}^T \sum_{h=1}^H b_{t,h}^r(s_{t,h}, a_{t,h}) \\
 &\leq L_{t,\delta} \left(6\sqrt{|\mathcal{L}_R| HT} + 10SAH \log(HT) + 23SAH + \sum_{t=1}^T \left(6\sqrt{\frac{|\mathcal{L}_R| H^2}{B(t)}} + 4SAH^2 \frac{\log(1+B(t))+1}{B(t)} \right) \right).
 \end{aligned}$$

We bound this sum based on the analysis technique that was utilized for CBM-UCB and UCB-OFUL.

Proof. Let $q_{t,h} = 1$ be the event reward feedback was queried at the h^{th} time step of the t^{th} episode and $q_{t,h} = 0$ be its complement event. The following relations hold.

$$\begin{aligned}
 \sum_{t=1}^T \sum_{h=1}^H C I_{t,h}^r(s_{t,h}, a_{t,h}) &= 2 \sum_{t=1}^T \sum_{h=1}^H b_{t,h}^r(s_{t,h}, a_{t,h}) \\
 &= \underbrace{\sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} C I_{t,h}^r(s_{t,h}, a_{t,h})}_{(i)} + \underbrace{\sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 0\} C I_{t,h}^r(s_{t,h}, a_{t,h})}_{(ii)} && (25)
 \end{aligned}$$

Bound on term (i), episodes in which feedback is queried. The following relations hold.

$$\begin{aligned}
 (i) &= 2 \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} b_{t,h}^r(s_{t,h}, a_{t,h}) \\
 &\leq 2 \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} \sqrt{\frac{2\widehat{\text{Var}}_{r,t-1,h}(s_{t,h}, a_{t,h})L_{t,\delta}}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1}} + 2 \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 1\} \frac{5L_{t,\delta}}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1} \\
 &\stackrel{(a)}{\leq} 3\sqrt{L_{T,\delta}} \sum_{s,a,h \in \mathcal{L}_R} \sum_{i=0}^{n_{T,h}^q(s,a)} \sqrt{\frac{1}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1}} + 10L_{T,\delta} \sum_{s,a,h} \sum_{i=0}^{n_{T,h}(s,a)} \frac{1}{n_{t-1,h}^q(s_{t,h}, a_{t,h}) \vee 1} \\
 &\leq 6\sqrt{|\mathcal{L}_R|B^q(T)L_{T,\delta}} + 10L_{T,\delta}SAH \log(B^q(T)) + 23SAHL_{T,\delta} \quad (\text{As (23),(24), \& } \sum_{s,a,h} n_{T,h}^q(s,a) = B^q(T)) \\
 &\leq 6\sqrt{|\mathcal{L}_R|THL_{T,\delta}} + 10L_{T,\delta}SAH \log(TH) + 23SAHL_{T,\delta}, \quad (B^q(T) \leq TH)
 \end{aligned}$$

Relation (a) holds since $\widehat{\text{Var}}_{r,t-1,h}(s_{t,h}, a_{t,h}) = 0$ for all rewards with deterministic reward (and zero reward is deterministic); for non-deterministic rewards, $\widehat{\text{Var}}_{r,t-1,h}(s_{t,h}, a_{t,h}) \leq 1$.

Bound on term (ii), episodes in which feedback is not queried. Due to Lemma 16 it holds that $B^q(T) \leq B(T)$, that is, the budget constraint is never violated. For this reason, if reward is not queried, it does not satisfy the query rule of CBM-UCBVI. Using this, we get the following relations.

$$\begin{aligned}
 (ii) &= \sum_{t=1}^T \sum_{h=1}^H 1\{q_{t,h} = 0\} CI_{t,h}^r(s_{t,h}, a_{t,h}) \\
 &\leq \sum_{t=1}^T \sum_{h=1}^H \left(L_{t,\delta} \left(6\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + 4SAH \frac{\log(1+B(t))+1}{B(t)} \right) \right) \\
 &\leq L_{T,\delta} \sum_{t=1}^T \sum_{h=1}^H \left(6\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + 4SAH \frac{\log(1+B(t))+1}{B(t)} \right).
 \end{aligned}$$

Combining the bounds. Combining the bounds on (i) and (ii) we conclude the proof,

$$(25) \leq L_{t,\delta} \left(6\sqrt{|\mathcal{L}_R|HT} + 10SAH \log(HT) + 23SAH + \sum_{t=1}^T \left(6\sqrt{\frac{|\mathcal{L}_R|H^2}{B(t)}} + 4SAH^2 \frac{\log(1+B(t))+1}{B(t)} \right) \right).$$

□

F. Confidence Budget Matching for Reinforcement Learning – CBM-ULCVI

Algorithm 7 CBM-ULCVI

Require: $\delta \in (0, 1)$

for $t = 1, 2, \dots$ **do**

$\left\{ \bar{Q}_{t,h}(s, a), \underline{Q}_{t,h}(s, a), \pi_{t,h}(a|s) \right\}_{(s,a) \in \mathcal{S} \times \mathcal{A}, h \in [H]}$ via Optimistic-Pessimistic VI (Algorithm 8)

Act with $\pi_{t,h}(a|s)$ and observe a trajectory $\{(s_{t,h}, a_{t,h})\}_{h=1}^H$

Observe current budget $B(t)$

Ask for feedback on $(s_{t,h}, a_{t,h})$ if

$$CI_{t,h}^R(s_{t,h}, a_{t,h}) \geq L_{t,\delta} \left(6\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + 4SAH \frac{\log(1 + B(t)) + 1}{B(t)} \right).$$

end for

In the previous section, we analyzed the performance of CBM-UCBVI algorithm, which incorporates the CBM principle into the Hoeffding-based UCBVI-CH (Azar et al., 2017). However, the performance of CBM-UCBVI is worse by a factor of \sqrt{H} relatively to the minimax performance even when the budget is $B(t) = Ht$. Then, the minimax regret bound is $\sqrt{SAH^3T}$ (e.g., Jin & Luo, 2019). Building on ideas from (Azar et al., 2017; Dann et al., 2019) we analyze the CBM Upper Lower Confidence Interval VI (CBM-ULCVI), which uses tighter CI to shave an additional \sqrt{H} from the asymptotic performance.

The idea is essentially the one used in (Azar et al., 2017; Jin et al., 2018; Zanette & Brunskill, 2019; Dann et al., 2019; Efroni et al., 2019). Instead of bounding $\text{Var}_{P(\cdot|s,a)} V_{h+1}^* \leq H^2$ and constructing the bonus as $b_{t,h}^p(s, a) \sim \frac{H}{\sqrt{n_{t-1,h}(s,a)}}$ (see Section E), we build the following bonus to compensate on errors in transition model:

$$b_{t,h}^p(s, a) \sim \frac{\sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{h+1}^k)}}{\sqrt{n_{t-1,h}(s, a)}} + \frac{1}{n_{t-1,h}(s, a)}. \quad (26)$$

This allows to get a performance guarantee that depends on the sum of variances along the trajectories $\sum_t \sum_h \text{Var}_{P_h(\cdot|s_t, a_t, h)}$. These can be bounded by $\sim H^2T$ via the law of total variance (Azar et al., 2017). Ultimately, this shaves a \sqrt{H} factor in the final bound: a ‘naïve’ bound, without applying the law of total variance, would result in H^3T . Our actual bonus is similar to the one used in (Dann et al., 2019), i.e., we use proper lower and upper value functions in addition to the bonus (26). However, our analysis is more ‘direct’ relatively to (Dann et al., 2019); we bound the regret by on-policy errors, and not expected on-policy errors as in (Dann et al., 2019). This becomes crucial due to the usage in the CBM scheme in our algorithm. Specifically, the CBM allows us to bound only

$$CI_{t,h}^R(s_{t,h}, a_{t,h}) 1\{q_{t,h} = 0\} \leq \tilde{\mathcal{O}} \left(\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + \frac{SAH}{B(t)} \right),$$

only on encountered state-action pairs, i.e., not under expectation.

Formally, we work with the following bonuses:

$$\begin{aligned} b_{t,h}^r(s, a) &= \sqrt{\frac{2\widehat{\text{Var}}_{R,t-1,h}(s, a)L_{t,\delta}}{n_{t-1,h}^q(s, a) \vee 1}} + \frac{5L_{t,\delta}}{n_{t-1,h}^q(s, a) \vee 1}, \\ b_{t,h}^p(s, a; \bar{V}_{t,h+1}, \underline{V}_{t,h+1}) &= \sqrt{\frac{2\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1})L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{44H^2SL_{t,\delta}}{n_{t-1,h}(s, a) \vee 1} \\ &\quad + \frac{1}{16H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - \underline{V}_{t,h+1}(s')] \end{aligned}$$

Algorithm 8 Optimistic-Pessimistic Value Iteration

Require: For all $s, a, h, n_{t-1,h}(s, a), n_{t-1,h}^q(s, a), \bar{P}_{t-1,h}(s), \bar{r}_{t-1,h}(s, a)$
Initialize: $\bar{V}_{H+1}(s) = \underline{V}_{H+1}(s) = 0$ for all $s \in \mathcal{S}$.
for $h = H, H-1, \dots, 1$ **do**
 for $s \in \mathcal{S}$ **do**
 $\forall a \in \mathcal{A}, \bar{Q}_{t,h}(s, a) = \bar{r}_{t-1,h}(s, a) + b_{t,h}^r(s, a) + b_{t,h}^p(s, a; \bar{V}_{t,h+1}, \underline{V}_{t,h+1}) + \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)}[\bar{V}_{t,h+1}(s')]$
 $\pi_{t,h}(s) \in \arg \max_a \bar{Q}_{t,h}(s, a)$
 $\forall a \in \mathcal{A}, \underline{Q}_{t,h}(s, a) = \bar{r}_{t-1,h}(s, a) - b_{t,h}^r(s, a) - b_{t,h}^p(s, a; \bar{V}_{t,h+1}, \underline{V}_{t,h+1}) + \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)}[\underline{V}_{t,h+1}(s')]$
 // Truncate values outside the interval $[0, H-h+1]$
 $\bar{V}_{t,h}(s) = \min\{\max_{a \in \mathcal{A}} \bar{Q}_{t,h}(s, a), H-h+1\}, \underline{V}_{t,h}(s) = \max(\underline{Q}_{t,h}(s, \pi_{t,h}(s)), 0)$
 end for
end for
Return: $\left\{ \bar{Q}_t(s, a), \underline{Q}_t(s, a), \pi_{t,h}(a|s) \right\}_{(s,a) \in \mathcal{S} \times \mathcal{A}, h \in [H]}$

where $L_{t,\delta} = \log \frac{16S^2 AH t^2 (t+1)}{\delta}$. As with CBM-UCBVI, we define $CI_{t,h}^R(s_{t,h}, a_{t,h}) = 2b_{t,h}^r(s, a)$ and similarly ask for feedback only if

$$CI_{t,h}^R(s_{t,h}, a_{t,h}) \geq L_{t,\delta} \left(6\sqrt{\frac{|\mathcal{L}_R|}{B(t)}} + 4SAH \frac{\log(1+B(t))+1}{B(t)} \right),$$

where \mathcal{L}_R is the set of rewarding state-actions, defined in Appendix E. This leads to Algorithm 7, which has the following regret guarantees:

Theorem 5 (CBM-ULCVDI). For any adversarially adaptive sequence $\{B(t), s_{t,1}\}_{t \geq 1}$ of budget and initial states the regret of CBM-ULCVDI is upper bounded by

$$\text{Reg}(T) \leq 45L_{T,\delta}^2 \left(\sqrt{SAH^3 T} + \sum_{t=1}^T \left(\sqrt{\frac{|\mathcal{L}_R| H^2}{B(t)}} + SAH^2 \frac{\log(1+B(t))+1}{B(t)} \right) \right) + 2904H^3 S^2 AL_{T,\delta}^2.$$

for any $T \geq 1$ with probability greater than $1 - \delta$.

F.1. The First Good Event - Concentration

Define the following events:

$$\begin{aligned}
 E^r(t) &= \left\{ \forall s \in \mathcal{S}, a \in \mathcal{A} : |\bar{r}_{t-1,h}(s, a) - r_h(s, a)| \leq \sqrt{\frac{2\widehat{\text{Var}}_{R,t-1,h}(s, a) \log \frac{16SAHt^2(t+1)}{\delta}}{n_{t-1,h}^q(s, a) \vee 1}} + \frac{5 \log \frac{16SAHt^2(t+1)}{\delta}}{n_{t-1,h}^q(s, a) \vee 1} \right\} \\
 E^p(t) &= \left\{ \forall s, s' \in \mathcal{S}, a \in \mathcal{A} : |P_h(s'|s, a) - \bar{P}_{t-1,h}(s'|s, a)| \leq \sqrt{\frac{2P(s'|s, a) \log \frac{16S^2 AH t^2 (t+1)}{\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{2 \log \frac{16S^2 AH t^2 (t+1)}{\delta}}{n_{t-1,h}(s, a) \vee 1} \right\} \\
 E^{pv1}(t) &= \left\{ \forall s, a, h : \left| (\bar{P}_{t-1,h}(\cdot|s, a) - P_h(\cdot|s, a))^T V_{h+1}^* \right| \leq \sqrt{\frac{2\text{Var}_{P_h(\cdot|s, a)}(V_{t+1}^*) \log \frac{16SAHt^2(t+1)}{\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{5H \log \frac{16SAHt^2(t+1)}{\delta}}{n_{t-1,h}(s, a) \vee 1} \right\} \\
 E^{pv2}(t) &= \left\{ \forall s, a, h : \left| \sqrt{\text{Var}_{P_h(\cdot|s, a)}(V_{h+1}^*)} - \sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s, a)}(V_{h+1}^*)} \right| \leq \sqrt{\frac{12H^2 \log \frac{16SAHt^2(t+1)}{\delta}}{n_{t-1,h}(s, a) \vee 1}} \right\}
 \end{aligned}$$

For brevity, we denote

$$b_{t,h}^{pv1}(s, a) = \sqrt{\frac{2\text{Var}_{P_h(\cdot|s, a)}(V_{t+1}^*) L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{5HL_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}.$$

Lemma 19 (The First Good Event). *Let $\mathbb{G}_1 = \cap_{t \geq 1} E^r(t) \cap_{t \geq 1} E^p(t) \cap_{t \geq 1} E^{pv1}(t) \cap_{t \geq 1} E^{pv2}(t)$ be the good event. It holds that $\Pr(\mathbb{G}_1) \geq 1 - \delta/2$.*

Proof. Each one of the events $\cap_t E^r(t), \cap_t E^p(t), \cap_t E^{pv1}(t)$ holds with probability greater than $1 - \delta/8$ via similar analysis as in Lemma 13. We now turn to establish the high probability guarantee for the rest of events.

The event $\cap_t E^{pv2}(t)$ holds with high probability. The event holds with high probability due to (Maurer & Pontil, 2009, Theorem 10). To see this, fix s, a, h and assume it was sampled for a fixed number of times $n_{t-1,h}(s, a) = n > 1$. Let $V_{s,a,h,n} = \frac{1}{2n(n-1)} \sum_{i,j} (V_{h+1}^*(s_i) - V_{h+1}^*(s_j))^2$ where $s_i \sim P_h(\cdot|s, a)$ and are i.i.d.. See that this is an unbiased estimator for the variance, i.e., $\mathbb{E}[V_{s,a,h,n}] = \text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)$. Applying (Maurer & Pontil, 2009, Theorem 10), we get that

$$\begin{aligned} \Pr\left(\left|\sqrt{V_{s,a,h,n}} - \sqrt{\mathbb{E}[V_{s,a,h,n}]}\right| \leq \sqrt{\frac{2H^2 \log(2/\delta)}{n-1}}\right) \\ = \Pr\left(\left|\sqrt{V_{s,a,h,n}} - \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)}\right| \leq \sqrt{\frac{2H^2 \log(2/\delta)}{n-1}}\right) \leq \delta. \end{aligned} \quad (27)$$

On the other hand,

$$\begin{aligned} V_{s,a,h,n} &= \frac{1}{2n(n-1)} \sum_{i,j=1}^n (V_{h+1}^*(s_i) - V_{h+1}^*(s_j))^2 = \sum_{s'} \frac{N(s')}{n} V_{h+1}^*(s')^2 - \left(\sum_{s'} \frac{N(s')}{n} V_{h+1}^*(s')\right)^2 \\ &\quad + \frac{1}{n(n-1)} \sum_{i=1}^n V_{h+1}^*(s_i)^2 - \frac{1}{n-1} \left(\frac{1}{n} \sum_{i=1}^n V_{h+1}^*(s_i)\right)^2, \end{aligned}$$

(by elementary algebra)

where $N(s')$ is the number of times s' was sampled. By definition, $\frac{N(s')}{n} \equiv \bar{P}_{n,h}(s'|s, a)$, that is, $\frac{N(s')}{n}$ is the estimated transition model. Thus, the $\text{Var}_{\bar{P}_{n,h}(\cdot|s,a)} = \sum_{s'} \frac{N(s')}{n} V_{h+1}^*(s')^2 - \left(\sum_{s'} \frac{N(s')}{n} V_{h+1}^*(s_i)\right)^2$. This, together with $|V_h^*(s)| \leq H$, implies that $\left|V_{s,a,h,n} - \text{Var}_{\bar{P}_{n,h}(\cdot|s,a)}(V_{h+1}^*)\right| \leq \frac{H^2}{n-1}$, and, thus

$$\left|\sqrt{V_{s,a,h,n}} - \sqrt{\text{Var}_{\bar{P}_{n,h}(\cdot|s,a)}(V_{h+1}^*)}\right| \leq \sqrt{\left|V_{s,a,h,n} - \text{Var}_{\bar{P}_{n,h}(\cdot|s,a)}(V_{h+1}^*)\right|} \leq \sqrt{\frac{H^2}{n-1}}, \quad (28)$$

by $|\sqrt{a} - \sqrt{b}| \leq \sqrt{|a-b|}$. Combining (27) and (28) yields

$$\Pr\left(\left|\sqrt{\text{Var}_{\bar{P}_{n,h}(\cdot|s,a)}(V_{h+1}^*)} - \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)}\right| \leq \sqrt{\frac{12H^2 \log(2/\delta)}{n \vee 1}}\right) \leq \delta,$$

since $\frac{2}{n \vee 1} \geq \frac{1}{n-1}$ for $n \geq 2$ and for $n \leq 1$ the bound trivially holds. Utilizing the same techniques as in Lemma 13 to generalize this result to hold for any (random) value of $n = n_{t-1,h}(s, a)$ and any s, a, h establishes the high probability guarantee.

Combining the results. Taking a union bound concludes the proof. \square

F.2. Optimism and Pessimism of Upper and Lower Value Functions

We can prove that the value is optimistic using standard techniques.

Lemma 20 (Upper Value Function is Optimistic, Lower Value Function is Pessimistic). *Conditioned on the first good event \mathbb{G}_1 the value function of CBM-UCBVI is optimism for all $s \in \mathcal{S}, h \in [H]$, i.e.,*

$$\forall s \in \mathcal{S}, h \in [H] : V_{t,h}(s) \leq V_h^{\pi_t}(s) \leq V_h^*(s) \leq \bar{V}_{t,h}(s).$$

Proof. Since $V_h^\pi(s) \leq V_h^*(s)$ for all $\pi, s \in \mathcal{S}, h \in [H]$, we only need to prove the leftmost and rightmost inequalities of the claim. We prove this result via induction.

Base case, the claim holds for $h = H$.

Rightmost inequality, optimism. Let $a^*(s) \in \arg \max_{a \in \mathcal{A}} r_H(s, a)$. For any $s \in \mathcal{S}$ it holds that

$$V_H^*(s) - \bar{V}_{t,H}(s) = r_H(s, a^*(s)) - \min \left\{ \max_a \{ \bar{r}_{t,H}(s, a) + b_{t,H}^r(s, a) \}, 1 \right\}. \quad (29)$$

Assume that $\max_a \bar{r}_{t,H}(s, a) + b_{t,H}^r(s, a) < 1$ (otherwise, the inequality holds since $r(s, a) \leq 1$). Then,

$$\begin{aligned} (29) &\leq r_H(s, a^*(s)) - \bar{r}_{t,H}(s, a^*(s)) - b_{t,H}^r(s, a^*(s)) \\ &\leq b_{t,H}^r(s, a^*(s)) - b_{t,H}^r(s, a^*(s)) \leq 0. \end{aligned} \quad (\text{event } \cap_t E^r(t) \text{ holds})$$

Thus, $V_H^*(s) \leq \bar{V}_{t,H}(s)$ for all $s \in \mathcal{S}$ for $h = H$.

Leftmost inequality, pessimism. For any $s \in \mathcal{S}$ it holds that

$$V_H^{\pi_t}(s) - \bar{V}_{t,H}(s) = r_H(s, \pi_{t,h}(s)) - \max(\bar{r}_{t,H}(s, \pi_{t,h}(s)) - b_{t,H}^r(s, \pi_{t,h}(s)), 0). \quad (30)$$

Assume that $\bar{r}_{t,H}(s, \pi_{t,h}(s)) - b_{t,H}^r(s, \pi_{t,h}(s)) > 0$ (otherwise, the inequality holds since $r(s, a) \geq 0$). Then,

$$(30) \geq -b_{t,H}^r(s, \pi_{t,h}(s)) + b_{t,H}^r(s, \pi_{t,h}(s)) \geq 0. \quad (\text{event } \cap_t E^r(t) \text{ holds})$$

Thus, $V_H^*(s) \geq \bar{V}_{t,H}(s)$ for all $s \in \mathcal{S}$ for $h = H$.

Induction step, prove for $h \in [H]$ assuming the claim holds for all $h+1 \leq h' \leq H$.

Rightmost inequality, optimism. Let $a^*(s) \in \arg \max_{a \in \mathcal{A}} Q_h^*(s, a)$. The following relations hold.

$$V_h^*(s) - \bar{V}_{t,h}(s) = Q_h^*(s, a^*(s)) - \min \left\{ \max_a \bar{Q}_{t,h}(s, a), H - h + 1 \right\} \quad (31)$$

Assume that $\max_a \bar{Q}_{t,h}^*(s, a) < H - h + 1$ (otherwise, the inequality is satisfied since $Q_h^*(s, a) \leq H - h + 1$). Then,

$$\begin{aligned} (31) &\leq Q_h^*(s, a^*(s)) - \bar{Q}_{t,h}(s, a^*(s)) \\ &= r_h(s, a^*(s)) - \bar{r}_{t-1,h}(s, a^*(s)) - b_{t,h}^r(s, a^*(s)) - b_{t,h}^p(s, a^*(s)) \\ &\quad + (P_h - \bar{P}_{t-1,h})(\cdot | s, a^*(s)) V_{h+1}^* + \underbrace{\mathbb{E}_{\bar{P}_{t-1,h}(\cdot | s, a^*(s))} [V_{h+1}^*(s') - \bar{V}_{t,h+1}(s')]}_{\leq 0 \text{ Induction hypothesis}} \\ &\leq b_{t,h}^r(s, a^*(s)) - b_{t,h}^r(s, a^*(s)) - b_{t,h}^p(s, a^*(s)) + b_{t,h}^{pv1}(s, a^*(s)) \quad (\text{events } \cap_t E^{pv1}(t) \ \& \ \cap_t E^r(t) \text{ holds}) \\ &= b_{t,h}^{pv1}(s, a^*(s)) - b_{t,h}^p(s, a^*(s)) \end{aligned} \quad (32)$$

We now analyze this term.

$$\begin{aligned} (32) &= b_{t,h}^{pv1}(s, a^*(s)) - b_{t,h}^p(s, a^*(s)) \\ &\stackrel{(a)}{\leq} \sqrt{\frac{2\text{Var}_{P(\cdot | s, a^*(s))}(V_{h+1}^*)L_{t,\delta}}{n_{t-1,h}(s, a^*(s)) \vee 1}} + \frac{5HL_{t,\delta}}{n_{t-1,h}(s, a^*(s)) \vee 1} \\ &\quad - \sqrt{\frac{2\text{Var}_{\bar{P}_{t-1,h}(\cdot | s, a^*(s))}(\bar{V}_{t,h+1})L_{t,\delta}}{n_{t-1,h}(s, a^*(s)) \vee 1}} - \frac{18H^2L_{t,\delta}}{n_{t-1,h}(s, a^*(s)) \vee 1} - \frac{1}{16H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot | s, a)} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] \\ &= \sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{P_h(\cdot | s, a^*(s))}(V_{h+1}^*)} - \sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot | s, a^*(s))}(\bar{V}_{t,h+1})}}{\sqrt{n_{t-1,h}(s, a^*(s)) \vee 1}} \\ &\quad - \frac{1}{16H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot | s, a)} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] - \frac{13H^2L_{t,\delta}}{n_{t-1,h}(s, a^*(s)) \vee 1} \\ &\stackrel{(b)}{\leq} \frac{1}{16H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot | s, a)} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] + \frac{13H^2L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1} \\ &\quad - \frac{1}{16H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot | s, a)} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] - \frac{13H^2L_{t,\delta}}{n_{t-1,h}(s, a^*(s)) \vee 1} = 0 \end{aligned}$$

where (a) holds by plugging the bonus and bounding $-44H^2SL_{t,\delta} \leq -18H^2L_{t,\delta}$, and by the induction hypothesis ($V_{t-1,h+1}(s) \leq V_{h+1}^*(s), \forall s \in \mathcal{S}$). (b) holds by Lemma 26 while setting $\alpha = 16H$ and bounding $(5H + H\alpha/2) \leq 13H^2$. Combining all the above we conclude the proof of the rightmost inequality since

$$V_h^*(s) - \bar{V}_{t,h}(s) \leq (31) \leq (32) \leq 0.$$

Leftmost inequality, pessimism. The following relations hold.

$$V_h^{\pi_t}(s) - \underline{V}_{t,h}(s) = Q_h^{\pi_t}(s, \pi_{t,h}(s)) - \max(\bar{Q}_{t,h}(s, \pi_{t,h}(s)), 0). \quad (33)$$

Assume that $\bar{Q}_{t,h}(s, \pi_{t,h}(s)) > 0$ (otherwise, the claim holds since $Q_h^{\pi_t}(s, \pi_{t,h}(s)) \geq 0$). Then,

$$\begin{aligned} (33) &= Q_h^{\pi_t}(s, \pi_{t,h}(s)) - \underline{Q}_{t,h}(s, \pi_{t,h}(s)) \\ &= r_h(s, \pi_{t,h}(s)) - \bar{r}_{t-1,h}(s, \pi_{t,h}(s)) + b_{t,h}^r(s, \pi_{t,h}(s)) + b_{t,h}^p(s, \pi_{t,h}(s)) \\ &\quad + (P_h - \bar{P}_{t-1,h})(\cdot|s, \pi_{t,h}(s))V_{h+1}^{\pi_t} + \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)} \underbrace{[V_{h+1}^{\pi_t}(s') - \underline{V}_{t,h+1}(s')]}_{\geq 0 \text{ Induction hypothesis}} |s_h = s, a_h = \pi_{t,h}(s) \\ &\geq -b_{t,h}^r(s, \pi_{t,h}(s)) + b_{t,h}^r(s, \pi_{t,h}(s)) + b_{t,h}^p(s, \pi_{t,h}(s)) + (P_h - \bar{P}_{t-1,h})(\cdot|s, \pi_{t,h}(s))V_{h+1}^{\pi_t} \quad (\cap_t E^r(t) \text{ holds}) \\ &= b_{t,h}^p(s, \pi_{t,h}(s)) + (P_h - \bar{P}_{t-1,h})(\cdot|s, \pi_{t,h}(s))V_{h+1}^{\pi_t}. \quad (34) \end{aligned}$$

We now focus on the last term. Observe that

$$\begin{aligned} (P_h - \bar{P}_{t-1,h})(\cdot|s, \pi_{t,h}(s))V_{h+1}^{\pi_t} &= (P_h - \bar{P}_{t-1,h})(\cdot|s, \pi_{t,h}(s))V_{h+1}^* + (P_h - \bar{P}_{t-1,h})(\cdot|s, \pi_{t,h}(s))(V_{h+1}^{\pi_t} - V_{h+1}^*) \\ &\geq -b_{t-1,h}^{pv1}(s, \pi_{t,h}(s)) + (P_h - \bar{P}_{t-1,h})(\cdot|s, \pi_{t,h}(s))(V_{h+1}^{\pi_t} - V_{h+1}^*) \quad (\cap_t E^{pv1}(t) \text{ holds}) \\ &\stackrel{(a)}{\geq} -b_{t-1,h}^{pv1}(s, \pi_{t,h}(s)) - \frac{18H^2SL_{t,\delta}}{n_{t-1,h}(s, \pi_{t,h}(s)) \vee 1} - \frac{1}{32H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s, \pi_{t,h}(s))} [(V_{h+1}^{\pi_t} - V_{h+1}^*)(s')] \\ &\geq -b_{t-1,h}^{pv1}(s, \pi_{t,h}(s)) - \frac{18H^2SL_{t,\delta}}{n_{t-1,h}(s, \pi_{t,h}(s)) \vee 1} - \frac{1}{32H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s, \pi_{t,h}(s))} [(\bar{V}_{t,h+1} - \underline{V}_{t-1,h+1})(s')] \\ &\hspace{15em} (\text{Induction hypothesis}) \\ &\geq -\sqrt{\frac{2\text{Var}_{P_h(\cdot|s, \pi_{t,h}(s))}(V_{h+1}^*)L_{t,\delta}}{n_{t-1,h}(s, \pi_{t,h}(s)) \vee 1}} - \frac{23H^2SL_{t,\delta}}{n_{t-1,h}(s, \pi_{t,h}(s)) \vee 1} - \frac{1}{32H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s, \pi_{t,h}(s))} [(\bar{V}_{t,h+1} - \underline{V}_{t-1,h+1})(s')], \\ &\hspace{15em} (\text{Plugging } b_{t-1,h}^{pv1} \text{ and elementary bounds}) \end{aligned}$$

where (a) holds by applying Lemma 28 while setting $\alpha = 32H$, $C_1 = 2L_{t,\delta}$, $C_2 = 2L_{t,\delta}$ and bounding $HS(C_2 + \alpha C_1/4) \leq 18H^2SL_{t,\delta}$ (assumption holds since $\cap_t E^p(t)$ holds). Plugging this back into (34) and plugging the explicit form of the bonus $b_{t,h}^p(s, a)$ we get

$$\begin{aligned} (34) &= \sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s, \pi_{t,h}(s))}(\bar{V}_{t,h+1})} - \sqrt{\text{Var}_{P_h(\cdot|s, \pi_{t,h}(s))}(V_{h+1}^*)}}{\sqrt{n_{t-1,h}(s, \pi_{t,h}(s)) \vee 1}} \\ &\quad + \frac{21H^2SL_{t,\delta}}{n_{t-1,h}(s, \pi_{t,h}(s)) \vee 1} + \frac{1}{32H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s, \pi_{t,h}(s))} [\bar{V}_{t,h+1}(s') - \underline{V}_{t-1,h+1}(s')] \\ &\stackrel{(a)}{\geq} -\frac{1}{32H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s, \pi_{t,h}(s))} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] - \frac{21H^2L_{t,\delta}}{n_{t-1,h}(s, \pi_{t,h}(s))} \\ &\quad + \frac{1}{32H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s, \pi_{t,h}(s))} [\bar{V}_{t,h+1}(s') - \underline{V}_{t-1,h+1}(s')] + \frac{21H^2SL_{t,\delta}}{n_{t-1,h}(s, \pi_{t,h}(s))} \geq 0 \end{aligned}$$

where (a) holds by Lemma 26 while setting $\alpha = 32H$ and bounding $(5H + H\alpha/2)L_{t,\delta} \leq 21H^2L_{t,\delta}$. Combining all the above we concludes the proof as

$$V_h^{\pi_t}(s) - \underline{V}_{t,h}(s) \geq (33) \geq (34) \geq 0.$$

□

F.3. The Good Event

We now prove an additional high probability bound which holds alongside first good event \mathbb{G}_1 .

Lemma 21 (The Good Event). *Let \mathbb{G}_1 be the event defined in Lemma 19. Let $\{Y_{1,t,h}, Y_{2,t,h}\}_{t \geq 1}$ the following random variables.*

$$\begin{aligned} Y_{1,t,h} &\stackrel{\text{def}}{=} \bar{V}_{t,h+1}(s_{t,h+1}) - V_{t,h+1}(s_{t,h+1}) \\ Y_{2,t,h} &= \text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t}). \end{aligned}$$

The second good event is the intersection of two events $\mathbb{G}_2 = E^{OP} \cap E^{\text{Var}}$ defined as follows.

$$\begin{aligned} E^{OP} &= \left\{ \forall h \in [H], T \geq 1 : \sum_{t=1}^T \mathbb{E}[Y_{1,t,h}|F_{t,h-1}] \leq \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T Y_{1,t,h} + 18H^2 \log \frac{4HT(T+1)}{\delta} \right\} \\ E^{\text{Var}} &= \left\{ T \geq 1 : \sum_{t=1}^T \sum_{h=1}^H Y_{2,t,h} \leq 2 \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}[Y_{2,t,h}|F_{t-1}] + 4H^3 \log \frac{4HT(T+1)}{\delta} \right\}, \end{aligned}$$

Then, the good event $\mathbb{G} = \mathbb{G}_1 \cap \mathbb{G}_2$ holds with probability greater than $1 - \delta$.

Proof. Event E^{OP} . The proof follows the lines of Lemma 15 while noting that $0 \leq \bar{V}_{t,h+1}(s_{t,h+1}) - V_{t,h+1}(s_{t,h+1}) \leq H$ (replacing $\bar{V}_{t,h+1}(s_{t,h+1}) - V_h^{\pi_t}(s_{t,h})$ as in Lemma 15). This holds conditioned on the first good event due to the optimism-pessimism lemma (Lemma 20). Notice that we replaced $\delta \rightarrow \delta/2$, and therefore, the proof results with $\Pr(\overline{E^{OP}} \cap \mathbb{G}_1) \leq \frac{\delta}{4}$.

Event E^{Var} . Fix $h \in [H]$. Set $Y_t = Y_{2,t,h} = \text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})$, and the filtration as $\{F_t\}_{t \geq 0}$ (observe that Y_t is F_t measurable). Furthermore, see that $0 \leq Y_{2,t,h} \leq H$ a.s. . Applying the second statement of Lemma 27 we get that

$$\sum_{t=1}^T Y_{2,t,h} \leq 2 \sum_{t=1}^T \mathbb{E}[Y_{2,t,h}|F_{t-1}] + 4H^2 \log \frac{1}{\delta}.$$

By taking union bound, as in the proof of the first statement of the lemma and on all $h \in [H]$ and summing over $h \in [H]$, we get that with probability greater than $1 - \delta/4$ for all $T \geq 1$ it holds that

$$\sum_{t=1}^T \sum_{h=1}^H Y_{2,t,h} \leq 2 \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}[Y_{2,t,h}|F_{t-1}] + 4H^3 \log \frac{4HT(T+1)}{\delta}.$$

Combining all the above We bound the probability of $\overline{\mathbb{G}}$ as follows:

$$\Pr(\overline{\mathbb{G}}) \leq \Pr(\overline{\mathbb{G}_1}) + \Pr(\overline{E^{OP}} \cap \mathbb{G}_1) + \Pr(\overline{E^{\text{Var}}}) \leq \frac{\delta}{2} + \frac{\delta}{4} + \frac{\delta}{4} = \delta.$$

□

F.4. CBM-ULCBVI: Budget Constraint is Satisfied

Similarly to all algorithms that follow the CBM paradigm, CBM-ULCBVI does not violate the budget constraint. That is, it queries reward feedback only if there is an available budget. The proof of this result is identical to the one of Lemma 16 (i.e., for the CBM-UCBVI algorithm), since the reward bonus and query rule for rewards of the two algorithms are the same, except to the log-factors. We restate the result here for convenience.

Lemma 22 (CBM-ULCBVI: Budget Constraint is Satisfied). *For any $T \geq 1$ the budget constraint is not violated, that is $B^q(T) \leq B(T)$ almost surely.*

E.5. Proof of Theorem 5

As in the proof of CBM-UCBVI, before establishing the proof of Theorem 5 we establish the following key lemma that bounds the on-policy errors at time step h by the on-policy errors at time step $h + 1$ and additional additive terms. Given this result, the analysis follows with relative ease.

Lemma 23 (CBM-ULCBVI, Key Recursion Bound). *Conditioned on the good event \mathbb{G} , the following bound holds for all $h \in [H]$.*

$$\begin{aligned} \sum_{t=1}^T \bar{V}_{t,h}(s_{t,h}) - \underline{V}_{t,h}(s_{t,h}) &\leq 27H^2 \log\left(\frac{4HT(T+1)}{\delta}\right) + 2 \sum_{t=1}^T b_{t,h}^r(s_{t,h}, a_{t,h}) + \sum_{t=1}^T \frac{224H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \\ &\quad + 2\sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t,h}(s_{t,h}, a_{t,h})} \vee 1} + \left(1 + \frac{1}{2H}\right)^2 \sum_{t=1}^T \bar{V}_{t,h+1}(s_{t,h+1}) - \underline{V}_{t,h+1}(s_{t,h+1}). \end{aligned}$$

Proof. We bound each of the terms in the sum as follows.

$$\begin{aligned} &\bar{V}_{t,h}(s_{t,h}) - \underline{V}_{t,h}(s_{t,h}) \\ &= 2b_{t,h}^r(s_{t,h}, a_{t,h}) + 2b_{t-1}^p(s_{t,h}, a_{t,h}) + \left(1 + \frac{1}{16H}\right) \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s_{t,h}, a_{t,h})}[\bar{V}_{t,h+1}(s_{h+1}) - \underline{V}_{t,h+1}(s_{h+1})] \\ &= 2b_{t,h}^r(s_{t,h}, a_{t,h}) + 2b_{t-1}^p(s_{t,h}, a_{t,h}) + \left(1 + \frac{1}{16H}\right) \mathbb{E}[\bar{V}_{t,h+1}(s_{h+1}) - \underline{V}_{t,h+1}(s_{h+1}) | F_{t,h-1}] \\ &\quad + \left(1 + \frac{1}{16H}\right) (\bar{P}_{t-1,h} - P_h)(\cdot|s, a)^T (\bar{V}_{t,h+1} - \underline{V}_{t,h+1}) \\ &\leq 2b_{t,h}^r(s_{t,h}, a_{t,h}) + 2b_{t-1}^p(s_{t,h}, a_{t,h}) + \frac{6SH^2 L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1} + \left(1 + \frac{1}{4H}\right) \mathbb{E}[\bar{V}_{t,h+1}(s_{h+1}) - \underline{V}_{t,h+1}(s_{h+1}) | F_{t,h-1}], \end{aligned} \tag{35}$$

where the last relation holds by Lemma 28 while setting $\alpha = 8H$, $C_1 = C_2 = 2L_{t,\delta}$ and bounding $HS(C_2 + \alpha SC_1/4) \leq 6SH^2 L_{t,\delta}$ (the assumption of the lemma holds since the event $\cap_t E^p(t)$ holds). By Lemma 24 it holds that

$$\begin{aligned} \sum_{t=1}^T b_{t,h}^p(s_{t,h}, a_{t,h}) &\leq \sum_{t=1}^T \frac{109H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} + \sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t,h}(s_{t,h}, a_{t,h})} \vee 1} \\ &\quad + \frac{1}{8H} \sum_{t=1}^T \mathbb{E}[\bar{V}_{t,h+1}(s_{t,h+1}) - \underline{V}_{t,h+1}(s_{t,h+1}) | F_{t,h-1}]. \end{aligned}$$

Plugging this into (35) and rearranging the terms we get

$$\begin{aligned} \sum_{t=1}^T \bar{V}_{t,h}(s_{t,h}) - \underline{V}_{t,h}(s_{t,h}) &\leq 2 \sum_{t=1}^T b_{t,h}^r(s_{t,h}, a_{t,h}) + \sum_{t=1}^T \frac{224H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} + 2\sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t,h}(s_{t,h}, a_{t,h})} \vee 1} \\ &\quad + \left(1 + \frac{1}{2H}\right) \sum_{t=1}^T \mathbb{E}[\bar{V}_{t,h+1}(s_{t,h+1}) - \underline{V}_{t,h+1}(s_{t,h+1}) | F_{t,h-1}] \\ &\leq 27H^2 \log\left(\frac{4HT(T+1)}{\delta}\right) + 2 \sum_{t=1}^T b_{t,h}^r(s_{t,h}, a_{t,h}) + \sum_{t=1}^T \frac{224H^2 S L_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} \\ &\quad + 2\sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t,h}(s_{t,h}, a_{t,h})} \vee 1} + \left(1 + \frac{1}{2H}\right)^2 \sum_{t=1}^T \bar{V}_{t,h+1}(s_{t,h+1}) - \underline{V}_{t,h+1}(s_{t,h+1}), \end{aligned}$$

(event E^{OP} holds (the second good event))

where in the last relation we also bounded $18(1 + \frac{1}{2H}) \leq 27$. \square

We are now ready to establish Theorem 5.

Proof. Start by conditioning on the good event which holds with probability greater than $1 - \delta$. Applying the optimism-pessimism of the upper and lower value function we get

$$\sum_{t=1}^T V_1^*(s_{t,1}) - V_1^{\pi_t}(s_{t,1}) \leq \sum_{t=1}^T \bar{V}_{t,1}(s_{t,1}) - \underline{V}_{t,1}(s_{t,1}). \quad (36)$$

Iteratively applying Lemma 23 and bound the exponential growth by $(1 + \frac{1}{2H})^{2h} \leq e \leq 3$ for any $h \in \{0, \dots, H\}$, The following upper bound on the cumulative regret is obtained.

$$\begin{aligned} (36) &\leq 81H^3 \log\left(\frac{HT(T+1)}{\delta}\right) + 6 \sum_{t=1}^T \sum_{h=1}^H b_{t,h}^r(s_{t,h}, a_{t,h}) + \sum_{t=1}^T \sum_{h=1}^H \frac{672H^2 SL_{t,\delta}}{n_{t-1,h}^p(s_{t,h}, a_{t,h}) \vee 1} \\ &\quad + 9 \sum_{t=1}^T \sum_{h=1}^H \frac{\sqrt{L_{t,\delta} \text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s_{t,h}, a_{t,h})}}. \end{aligned} \quad (37)$$

The first sum in (37) is bounded in Lemma 18 (the reward bonus is exactly the same as for CBM-UCBVI which implies the same upper bound on the sum of bonuses hold) by

$$\begin{aligned} &6 \sum_{t=1}^T \sum_{h=1}^H b_{t,h}^r(s_{t,h}, a_{t,h}) \\ &\leq 3L_{t,\delta} \left(6\sqrt{|\mathcal{L}_R|HT} + 10SAH \log(HT) + 23SAH + \sum_{t=1}^T \left(6\sqrt{\frac{|\mathcal{L}_R|H^2}{B(t)}} + 4SAH^2 \frac{\log(1+B(t))+1}{B(t)} \right) \right). \end{aligned}$$

The second sum is bounded via standard analysis as follows:

$$\begin{aligned} \sum_{t=1}^T \sum_{h=1}^H \frac{672H^2 SL_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} &\leq 672H^2 SL_{T,\delta} \sum_{t=1}^T \sum_{h=1}^H \frac{1}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} && (L_{t,\delta} \text{ increasing in } t) \\ &= 672H^2 SL_{T,\delta} \sum_{s,a,h} \sum_{i=0}^{n_{T,h}(s,a)} \frac{1}{i \vee 1} && (\text{Reorganizing summation}) \\ &\leq 672H^2 SL_{T,\delta} \sum_{s,a,h} (2 + \log(n_{T,h}(s,a) \vee 1)) \\ &\leq 672H^3 S^2 AL_{T,\delta} (2 + \log(TH)) && (\text{Jensen's inequality and } \sum_{s,a,h} n_{T,h}(s,a) = HT) \\ &\leq 2688H^3 S^2 AL_{T,\delta} \log(TH + 1) \end{aligned}$$

The third sum in (37) is bounded in Lemma 25 by

$$\begin{aligned} 9 \sum_{t=1}^T \sum_{h=1}^H \frac{\sqrt{L_{t,\delta} \text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s_{t,h}, a_{t,h})}} &\leq 9\sqrt{L_{T,\delta}} \sum_{t=1}^T \sum_{h=1}^H \frac{\sqrt{\text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s_{t,h}, a_{t,h})}} && (L_{t,\delta} \text{ increasing in } t) \\ &\leq 27\sqrt{SAH^3 T \log(TH + 1)} L_{T,\delta} + 36H^2 \sqrt{SA \log(TH + 1) \log\left(\frac{4HT(T+1)}{\delta}\right)} L_{T,\delta} \end{aligned}$$

Combining the above bounds with proper simplification yields a bound on (37) and concludes the proof,

$$(37) \leq 45L_{T,\delta}^2 \left(\sqrt{SAH^3 T} + \sum_{t=1}^T \left(\sqrt{\frac{|\mathcal{L}_R|H^2}{B(t)}} + SAH^2 \frac{\log(1+B(t))+1}{B(t)} \right) \right) + 2904H^3 S^2 AL_{T,\delta}^2.$$

□

E.6. Results that Hold Conditioned on the Good Event

Lemma 24 (Bound on the Cumulative Transition Model Bonus). *Conditioning on the good event \mathbb{G}_1 the following bound holds for all $h \in [H]$.*

$$\begin{aligned} \sum_{t=1}^T b_t^p(s_{t,h}, a_{t,h}) &\leq \sum_{t=1}^T \frac{109H^2 SL_{t,\delta}}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1} + \sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} \\ &\quad + \frac{1}{8H} \sum_{t=1}^T \mathbb{E}[\bar{V}_{t,h+1}(s') - V_{t,h+1}(s') | F_{t,h-1}]. \end{aligned}$$

Proof. First, observe that

$$\begin{aligned} &\mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - V_{t,h+1}(s')] \\ &= \mathbb{E}_{P_h(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - V_{t,h+1}(s')] + (\bar{P}_{t-1,h} - P_h)(\cdot|s,a)^T [\bar{V}_{t,h+1}(s') - V_{t,h+1}(s')] \\ &\leq \frac{9}{8} \mathbb{E}_{P_h(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - V_{t,h+1}(s')] + \frac{6H^2 SL_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}, \end{aligned} \quad (38)$$

by applying Lemma 28 with $\alpha = 8H$, $C_1 = C_2 = 2L_{t,\delta}$ and $HS(C_2 + \alpha C_1/4) \leq 6H^2 S$ (applicable since $\cap_t EP(t)$ holds).

The bonus $b_{t,h}^p(s, a)$ can be upper bounded as follows.

$$\begin{aligned} b_{t,h}^p(s, a) &\leq \sqrt{2} \sqrt{\frac{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1}) L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{1}{16H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - V_{t,h+1}(s')] + \frac{44H^2 SL_{t,\delta}}{n_{t-1,h}(s, a) \vee 1} \\ &\leq \sqrt{2} \sqrt{\frac{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1}) L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}} + \frac{9}{128H} \mathbb{E}_{P_h(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - V_{t,h+1}(s')] + \frac{50H^2 SL_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}, \end{aligned} \quad (39)$$

by (38). We bound the first term to establish the lemma.

Bound on the first term of (39) It holds that

$$\begin{aligned} &\sqrt{2L_{t,\delta}} \sqrt{\frac{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1})}{n_{t-1,h}(s, a) \vee 1}} \\ &= \underbrace{\sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1})} - \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)}}{\sqrt{n_{t-1,h}(s, a) \vee 1}}}_{(i)} + \underbrace{\sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)} - \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s, a) \vee 1}}}_{(ii)} \\ &\quad + \frac{\sqrt{2L_{t,\delta}} \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s, a) \vee 1}}. \end{aligned}$$

Term (i) is bounded by Lemma 26 (by setting $\alpha = 32H$ and $(5H + H\alpha/2) \leq 21H^2$),

$$\sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1})} - \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)}}{\sqrt{n_{t-1,h}(s, a) \vee 1}} \leq \frac{1}{32H} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] + \frac{21H^2 L_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}.$$

Following the same steps as in (38), we get

$$\mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] \leq \frac{9}{8} \mathbb{E}_{P_h(\cdot|s,a)}[\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] + \frac{6H^2 SL_{t,\delta}}{n_{t-1,h}(s, a) \vee 1},$$

and, thus,

$$(i) \leq \frac{9}{256H} \mathbb{E}_{P_h(\cdot|s,a)}[\bar{V}_{t-1,h+1}(s') - V_{h+1}^*(s')] + \frac{27H^2 SL_{t,\delta}}{n_{t-1,h}(s, a) \vee 1}.$$

Term (ii) is bounded as follows.

$$\begin{aligned}
 (ii) &\leq \sqrt{2L_{t,\delta}} \frac{\sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^* - V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s,a) \vee 1}} && \text{(By Lemma 31)} \\
 &\leq \sqrt{2L_{t,\delta}} \frac{\sqrt{\mathbb{E}_{P_h(\cdot|s,a)}[(V_{h+1}^*(s') - V_{h+1}^{\pi_t}(s'))^2]}}{\sqrt{n_{t-1,h}(s,a) \vee 1}} \\
 &\leq \sqrt{2L_{t,\delta}} \frac{\sqrt{H \mathbb{E}_{P_h(\cdot|s,a)}[(V_{h+1}^*(s') - V_{h+1}^{\pi_t}(s'))]}}{\sqrt{n_{t-1,h}(s,a) \vee 1}} && (0 \leq V_{h+1}^*(s') - V_{h+1}^{\pi_t}(s') \leq H) \\
 &\leq \frac{1}{64H} \mathbb{E}_{P_h(\cdot|s,a)}[(V_{h+1}^*(s') - V_{h+1}^{\pi_t}(s'))] + \frac{32H^2 L_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}. && (ab \leq \frac{1}{\alpha} a^2 + \frac{\alpha}{4} b^2 \text{ for } \alpha = 64H)
 \end{aligned}$$

Thus, applying $V_{h+1} \leq V_{h+1}^{\pi_t} \leq V_{h+1}^* \leq \bar{V}_{h+1}$ (Lemma 20) in the bounds of (i) and (ii) we get

$$b_{t,h}^p(s,a) \leq \frac{1}{8H} \mathbb{E}_{P_h(\cdot|s,a)}[(\bar{V}_{t,h}(s') - V_{t,h}(s'))] + \frac{109H^2 S L_{t,\delta}}{n_{t-1,h}(s,a) \vee 1} + \frac{\sqrt{2L_{t,\delta}} \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s,a) \vee 1}}.$$

and summing over t concludes the proof. \square

Lemma 25 (Bound on Variance Term). *Conditioning on the good event E^{Var} it holds that*

$$\sum_{t=1}^T \sum_{h=1}^H \frac{\sqrt{\text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s_{t,h}, a_{t,h})}} \leq 3\sqrt{SAH^3 T \log(TH+1)} + 4H^2 \sqrt{SA \log(TH+1) \log\left(\frac{4HT(T+1)}{\delta}\right)}.$$

Proof. Applying Cauchy-Schwartz inequality we get

$$\begin{aligned}
 &\sum_{t=1}^T \sum_{h=1}^H \frac{\sqrt{\text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})}}{\sqrt{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} \\
 &\leq \sqrt{\sum_{t=1}^T \sum_{h=1}^H \text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})} \sqrt{\sum_{t=1}^T \sum_{h=1}^H \frac{1}{n_{t-1,h}(s_{t,h}, a_{t,h}) \vee 1}} && \text{(Cauchy-Schwarz inequality)} \\
 &\leq \sqrt{\sum_{t=1}^T \sum_{h=1}^H \text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t})} \sqrt{SAH(2 + \log(TH))} && \text{(By eq. (23))} \\
 &\leq 2 \sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\sum_{h=1}^H \text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t}) | F_{t-1} \right]} + 4H^3 \log\left(\frac{4HT(T+1)}{\delta}\right) \sqrt{SAH \log(TH+1)} && (E^{\text{Var}} \text{ holds}) \\
 &\leq 3 \sqrt{\sum_{t=1}^T \mathbb{E} \left[\sum_{h=1}^H \text{Var}_{P_h(\cdot|s_{t,h}, a_{t,h})}(V_{h+1}^{\pi_t}) | F_{t-1} \right]} \sqrt{SAH \log(TH+1) + 4H^2 \sqrt{SA \log(TH+1) \log\left(\frac{4HT(T+1)}{\delta}\right)}} \\
 &\hspace{15em} (\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}) \\
 &= 3 \sqrt{\sum_{t=1}^T \mathbb{E} \left[\left(V_1^{\pi_t}(s_1) - \sum_{h=1}^H r_h(s_{t,h}, a_{t,h}) \right)^2 | F_{t-1} \right]} \sqrt{SAH \log(TH+1) + 4H^2 \sqrt{SA \log(TH+1) \log\left(\frac{4HT(T+1)}{\delta}\right)}} \\
 &\hspace{10em} \text{(Law of total variance (Azar et al., 2017), see Lemma 32)} \\
 &\leq 3\sqrt{SAH^3 T \log(TH+1)} + 4H^2 \sqrt{SA \log(TH+1) \log\left(\frac{4HT(T+1)}{\delta}\right)}. && (V_1^{\pi_t}(s) \in [0, H], r_h(s, a) \in [0, 1])
 \end{aligned}$$

\square

Lemma 26 (Variance Difference is Upper Bounded by Value Difference). *Assume that the value at time step $h + 1$ is optimistic, $\bar{V}_{t,h+1}(s) \geq V_{h+1}^*(s)$ for all $s \in \mathcal{S}$. Conditioning on the event $\cap_t E^{pv2}(t)$ it holds for all $s, a \in \mathcal{S} \times \mathcal{A}, \in [H]$ that*

$$\sqrt{2L_{t,\delta}} \left| \frac{\sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1})} - \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)}}{\sqrt{n_{t-1,h}(s,a)}} \right| \leq \frac{1}{\alpha} \mathbb{E}_{\bar{P}_{t-1,h}(\cdot|s,a)} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] + \frac{(5H + H\alpha/2)L_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}.$$

Proof. Conditioning on the first good event the following relations hold.

$$\begin{aligned} & \left| \sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1})} - \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)} \right| \\ & \leq \left| \sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1})} - \sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(V_{h+1}^*)} \right| + \sqrt{\frac{12H^2L_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}} \quad (\cap_t E^{pv2}(t) \text{ holds}) \\ & \leq \sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(V_{h+1}^* - \bar{V}_{t,h+1})} + \sqrt{\frac{12H^2L_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}} \quad (\text{Lemma 31}) \\ & \leq \sqrt{\mathbb{E}_{\bar{P}_{t-1,h}} [(V_{h+1}^*(s') - \bar{V}_{t,h+1}(s'))^2]} + \sqrt{\frac{12H^2L_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}} \\ & \leq \sqrt{H\mathbb{E}_{\bar{P}_{t-1,h}} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')]} + \sqrt{\frac{12H^2L_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}}, \end{aligned}$$

where the last relation holds since $V_{h+1}^*(s'), \bar{V}_{t,h+1}(s') \in [0, H]$ (the first, by model assumption, and the second, by the update rule) and since $V_{h+1}^*(s') \leq \bar{V}_{t,h+1}(s')$ by the assumption the value is optimistic. Thus,

$$\begin{aligned} & \sqrt{2L_{t,\delta}} \left| \frac{\sqrt{\text{Var}_{\bar{P}_{t-1,h}(\cdot|s,a)}(\bar{V}_{t,h+1})} - \sqrt{\text{Var}_{P_h(\cdot|s,a)}(V_{h+1}^*)}}{\sqrt{n_{t-1,h}(s,a)}} \right| \\ & \leq \sqrt{\mathbb{E}_{\bar{P}_{t-1,h}} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')]} \sqrt{\frac{2HL_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}} + \frac{\sqrt{24}HL_{t,\delta}}{n_{t-1,h}(s,a) \vee 1} \\ & \leq \frac{1}{\alpha} \mathbb{E}_{\bar{P}_{t-1,h}} [\bar{V}_{t,h+1}(s') - V_{h+1}^*(s')] + \frac{(5H + H\alpha/2)L_{t,\delta}}{n_{t-1,h}(s,a) \vee 1}. \quad (\text{Youngs inequality, } ab \leq \frac{1}{\alpha}a^2 + \frac{\alpha}{4}b^2) \end{aligned}$$

□

G. Useful Results

Lemma 27 (Consequences of Freedman's Inequality for Bounded and Positive Sequence of Random Variables). *Let $\{Y_t\}_{t \geq 1}$ be a real valued sequence of random variables adapted to a filtration $\{F_t\}_{t \geq 0}$. Assume that for all $t \geq 1$ it holds that $0 \leq Y_t \leq C$ a.s., and $T \in \mathbb{N}$. Then each of the following inequalities hold with probability greater than $1 - \delta$.*

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}[Y_t | F_{t-1}] &\leq \left(1 + \frac{1}{2C}\right) \sum_{t=1}^T Y_t + 2(2C+1)^2 \log \frac{1}{\delta}, \\ \sum_{t=1}^T Y_t &\leq 2 \sum_{t=1}^T \mathbb{E}[Y_t | F_{t-1}] + 4C \log \frac{1}{\delta}. \end{aligned}$$

Proof. First statement. Let $X_t \stackrel{\text{def}}{=} \mathbb{E}[Y_t | F_{t-1}] - Y_t$. Observe that X_t is a martingale difference sequence w.r.t. to $\{F_t\}_{t \geq 0}$, and that $|X_t| \leq C$ a.s.. Furthermore, observe that

$$\begin{aligned} \mathbb{E}[X_t^2 | F_{t-1}] &\leq 2\mathbb{E}[(\mathbb{E}[Y_t | F_{t-1}])^2 | F_{t-1}] + 2\mathbb{E}[Y_t^2 | F_{t-1}] && ((a-b)^2 \leq 2a^2 + 2b^2) \\ &\leq 4\mathbb{E}[Y_t^2 | F_{t-1}] && (\text{Jensen's inequality}) \\ &\leq 4C\mathbb{E}[Y_t | F_{t-1}], && (40) \end{aligned}$$

where the last relation holds since $Y_t \geq 0$ and $Y_t \leq C$. Applying Freedman's inequality we get that

$$\begin{aligned} \sum_{t=1}^T (\mathbb{E}[Y_t | F_{t-1}] - Y_t) &= \sum_{t=1}^T X_t \leq \eta \sum_{t=1}^T \mathbb{E}[X_t^2 | F_{t-1}] + \frac{\log(1/\delta)}{\eta} && (\text{Lemma 29}) \\ &\leq \eta \sum_{t=1}^T 4C\mathbb{E}[Y_t | F_{t-1}] + \frac{\log(1/\delta)}{\eta}. && (\text{Eq. (40)}) \end{aligned}$$

Choosing $\eta = \frac{1}{4C(2C+1)} \in (0, \frac{1}{C})$ and rearranging we get

$$\begin{aligned} \sum_{t=1}^T \left(1 - \frac{1}{2C+1}\right) \mathbb{E}[Y_t | F_{t-1}] &= \sum_{t=1}^T \left(\frac{2C}{2C+1}\right) \mathbb{E}[Y_t | F_{t-1}] \\ &\leq \sum_{t=1}^T Y_t + 4C(2C+1) \log(1/\delta). \end{aligned}$$

Thus,

$$\sum_{t=1}^T \mathbb{E}[Y_t | F_{t-1}] \leq \sum_{t=1}^T \left(1 + \frac{1}{2C}\right) Y_t + 2(2C+1)^2 \log(1/\delta).$$

Second statement. Let $X_t \stackrel{\text{def}}{=} Y_t - \mathbb{E}[Y_t | F_{t-1}]$. Similarly to before, it holds that X_t is a martingale difference sequence w.r.t. to the filtration F_{t-1} , and that $|X_{t,h}| \leq C$ a.s. . Similarly to (40) which is sign invariant we get $\mathbb{E}[X_t^2 | F_{t-1}] \leq 4C\mathbb{E}[Y_t | F_{t-1}]$. Applying Freedman's inequality we get

$$\begin{aligned} \sum_{t=1}^T Y_t - \sum_{t=1}^T \mathbb{E}[Y_t | F_{t-1}] &= \sum_{t=1}^T X_t \leq \eta \sum_{t=1}^T \mathbb{E}[X_t^2 | F_{t-1}] + \frac{\log(1/\delta)}{\eta} && (\text{Lemma 29}) \\ &\leq \eta \sum_{t=1}^T 4C\mathbb{E}[Y_t | F_{t-1}] + \frac{\log(1/\delta)}{\eta}. \end{aligned}$$

Setting $\eta = 1/4C$ and rearranging concludes the proof of the second statement. \square

Lemma 28 (Transition Difference to Next State Expectation). *Let $Y \in \mathbb{R}^{\mathcal{S}}$ be a vector such that $0 \leq Y(s) \leq H$ for all $s \in \mathcal{S}$. Let P_1 and P_2 be two transition models and $n \in \mathbb{R}_+^{\mathcal{S}^A}$. Let $\Delta P_h(\cdot|s, a) \in \mathbb{R}^{\mathcal{S}}$ and $\Delta P_h(s'|s, a) \stackrel{\text{def}}{=} P_{1,h}(s'|s, a) - P_{2,h}(s'|s, a)$. Assume that*

$$\left\{ \forall s, a, s' \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}, h \in [H] : |\Delta P_h(s'|s, a)| \leq \sqrt{\frac{C_1 L_{t,\delta} P_{1,h}(s'|s, a)}{n(s, a) \vee 1}} + \frac{C_2 L_{t,\delta}}{n(s, a) \vee 1} \right\},$$

for some $C_1, C_2 > 0$, then, for any $\alpha > 0$,

$$|\Delta P_h(\cdot|s, a)Y| \leq \frac{1}{\alpha} \mathbb{E}_{P_1(\cdot|s, a)}[Y(s')] + \frac{HSL_{t,\delta}(C_2 + \alpha C_1/4)}{n(s, a) \vee 1},$$

Proof. The following relations hold.

$$\begin{aligned} |\Delta P_h(\cdot|s, a)Y| &\leq \sum_{s'} |\Delta P_h(s'|s, a)| \cdot |Y(s')| \\ &\leq \sum_{s'} \left(\sqrt{\frac{C_1 P_{1,h}(s'|s, a)}{n(s, a) \vee 1}} Y(s') + \frac{HC_2}{n(s, a) \vee 1} \right) \quad (\text{By assumption of the lemma \& } Y(s) \in [0, H]) \\ &= \sum_{s'} \sqrt{\frac{C_1 P_{1,h}(s'|s, a) Y^2(s')}{n(s, a) \vee 1}} + \frac{HSC_2}{n(s, a) \vee 1} \\ &\leq \sum_{s'} \sqrt{P_{1,h}(s'|s, a) Y(s')} \sqrt{\frac{HC_1}{n(s, a) \vee 1}} + \frac{HSC_2}{n(s, a) \vee 1} \quad (0 \leq Y(s') \leq H) \\ &\leq \sqrt{\sum_{s'} P_{1,h}(s'|s, a) Y(s')} \sqrt{\frac{SHC_1}{n(s, a) \vee 1}} + \frac{HSC_2}{n(s, a) \vee 1} \quad (\text{Cauchy-Schwarz inequality}) \\ &\stackrel{(*)}{\leq} \frac{1}{2\alpha} \sum_{s'} P_{1,h}(s'|s, a) Y(s') + \frac{\alpha SHC_1}{2n(s, a) \vee 1} + \frac{HSC_2}{n(s, a) \vee 1}, \end{aligned}$$

where (*) is by Young's inequality, namely, $ab \leq \frac{a^2}{2\alpha} + \frac{\alpha}{2}b^2$ for any $\alpha > 0$. Re-scaling $2\alpha = \alpha'$ we conclude the proof. \square

H. Useful Existing Results

Lemma 29 (Freedman's Inequality, [Beygelzimer et al. 2011](#), Theorem 1). *Let $\{X_t\}_{t \geq 1}$ be a real valued martingale difference sequence adapted to a filtration $\{F_t\}_{t \geq 0}$. If $|X_t| \leq R$ a.s. then for any $\eta \in (0, 1/R]$, $T \in \mathbb{N}$ it holds with probability greater than $1 - \delta$,*

$$\sum_{t=1}^T X_t \leq \eta \sum_{t=1}^T \mathbb{E}[X_t^2 | F_{t-1}] + \frac{\log(1/\delta)}{\eta}.$$

Theorem 6 ([Abbasi-Yadkori et al. 2011](#), Theorem 1). *Let $\{F_t\}_{t=0}^\infty$ be a filtration. Let $\{\eta_t\}_{t=0}^\infty$ be a real-valued stochastic process such that η_t is F_t -measurable and η_t is conditionally σ -sub-Gaussian for $\sigma \geq 0$. Let $\{x_t\}_{t=0}^\infty$ be an \mathbb{R}^d -valued stochastic process s.t. X_t is F_{t-1} -measurable and $\|x_t\| \leq L$. Assume that V is a $d \times d$ positive-definite matrix. For any $t \geq 0$, define $V_t = V + \sum_{s=1}^t x_s x_s^T$, and $S_t = \sum_{s=1}^t \eta_s X_s$. Then, for any $\delta > 0$, with a probability of at least $1 - \delta$, for all $t \geq 0$,*

$$\|S_t\|_{V_t^{-1}}^2 \leq 2\sigma^2 \log\left(\frac{\det(V_t)^{1/2} \det(V)^{-1/2}}{\delta}\right).$$

The following results is an adaptation of [Abbasi-Yadkori et al. 2011](#), Theorem 2. It establishes a concentration guarantee for the least square estimator with a skipping process.

Theorem 7 (Concentration of Least-Square with Skipping Process). *Let $\hat{\theta}_t = (X_t^T X_t + \lambda I_d)^{-1} X_t^T Y_t$, where X_t is the matrix whose rows are $1\{q_1 = 1\}x_1^T, \dots, 1\{q_t = 1\}x_t^T$, $Y_t = (1\{q_1 = 1\}y_1, \dots, 1\{q_t = 1\}y_t)^T$, $y_t = \langle x_t, \theta^* \rangle + \eta_t$, and $\{q_t\}_{t \geq 1}$ is a sequence of binary F_{t-1} measurable events. Also assume that $\|\theta^*\| \leq D$. Then, for any $\delta > 0$ with a probability of at least $1 - \delta$ for all, $t \geq 0$, θ^* lies in the set*

$$C_t \stackrel{\text{def}}{=} \left\{ \theta \in \mathbb{R}^d : \|\hat{\theta}_t - \theta\|_{V_t} \leq \sigma \sqrt{d \log\left(\frac{1 + tL^2/\lambda}{\delta}\right)} + \lambda^{1/2} D \right\}.$$

Proof. Let $\eta = (1\{q_1 = 1\}\eta_1, \dots, 1\{q_t = 1\}\eta_t)^T$ and denote $X = X_t, Y = Y_t$. Observe that

$$\begin{aligned} \hat{\theta}_t &= (X^T X + \lambda I_d)^{-1} X^T (X \theta^* + \eta) \\ &= (X^T X + \lambda I_d)^{-1} X^T \eta + \theta^* - \lambda (X^T X + \lambda I_d)^{-1} \theta^*. \end{aligned}$$

Thus, for any $x \in \mathbb{R}^d$

$$\langle x, \hat{\theta}_t - \theta^* \rangle = \langle x, X^T \eta \rangle_{V_t^{-1}} - \lambda \langle x, \theta^* \rangle_{V_t^{-1}},$$

where $V_t = X^T X + \lambda I$. Since V_t is positive definite and symmetric the inner product is well define. By Cauchy-Schwartz inequality, we get

$$\begin{aligned} |\langle x, \hat{\theta}_t - \theta^* \rangle| &\leq \|x\|_{V_t^{-1}} \left(\|X^T \eta\|_{V_t^{-1}} + \lambda \|\theta^*\|_{V_t^{-1}} \right) \\ &\leq \|x\|_{V_t^{-1}} \left(\|X^T \eta\|_{V_t^{-1}} + \lambda^{1/2} \|\theta^*\| \right). \end{aligned} \tag{41}$$

We now apply Theorem 6 to bound $\|X^T \eta\|_{V_t^{-1}}$. Let

$$X^T \eta = \sum_{s=1}^t x_s 1\{q_t = 1\} \eta_t = \sum_{s=1}^t x_s 1\{q_t = 1\} \eta_t = \sum_{s=1}^t x'_s \eta_t \stackrel{\text{def}}{=} S_t,$$

where we defined $x'_t = x_t 1\{q_t = 1\}$. See that x'_t is F_{t-1} measurable since both x_t and $1\{q_t = 1\}$ are F_{t-1} measurable. Thus, we can apply Theorem 6 and get the following bound with probability greater than $1 - \delta$ for all $t \geq 0$

$$\|X^T \eta\|_{V_t^{-1}} \leq \sqrt{2\sigma^2 \log\left(\frac{\det(V_t)^{1/2} \det(V)^{-1/2}}{\delta}\right)}.$$

Conditioning on the event this bound holds and setting $x = V_t(\hat{\theta}_t - \theta^*)$ in (41) and using $\|\theta^*\| \leq D$ we get

$$\|\hat{\theta}_t - \theta^*\|_{V_t}^2 \leq \|V_t(\hat{\theta}_t - \theta^*)\|_{V_t^{-1}} \left(\sqrt{2\sigma^2 \log \left(\frac{\det(V_t)^{1/2} \det(V)^{-1/2}}{\delta} \right)} + \lambda^{1/2} D \right).$$

Using $\|V_t(\hat{\theta}_t - \theta^*)\|_{V_t^{-1}} = \|\hat{\theta}_t - \theta^*\|_{V_t}$ and rearranging leads to

$$\|\hat{\theta}_t - \theta^*\|_{V_t} \leq \sigma \sqrt{2 \log \left(\frac{\det(V_t)^{1/2} \det(V)^{-1/2}}{\delta} \right)} + \lambda^{1/2} D. \quad (42)$$

To obtain the final form, we us bound the term $\log \left(\det(V_t)^{1/2} \det(V)^{-1/2} \right)$. The trace of V_t is bounded by $\text{trace}(V) + tL^2$ since $\|x'_t\| \leq \|x_t\| \leq L$. Hence, $\det(V_t) = \prod_{i=1}^d \lambda_i \leq \left(\frac{\text{trace}(V) + tL^2}{d} \right)^d$ by the GM-AM inequality and $\sum_i \lambda_i = \text{trace}(V_t)$. Therefore,

$$\log \left(\det(V_t) \det(V)^{-1} \right) = \log \det(V_t) - d \log \lambda \leq d \log \left(\frac{\lambda + tL^2}{d} \right) - d \log \lambda = d \log \left(\frac{1 + tL^2/\lambda}{d} \right).$$

Plugging this back into (42) concludes the proof. \square

Lemma 30 (Elliptical Potential Lemma, (Abbasi-Yadkori et al., 2011), Lemma 11). *Let $\{x_t\}_{t=1}^\infty$ be a sequence in \mathbb{R}^d and $V_t = V + \sum_{i=1}^t x_i x_i^T$. Assume $\|x_t\| \leq L$ for all t . Then,*

$$\sum_{i=1}^t \min \left(\|x_i\|_{V_{i-1}}^2, 1 \right) \leq 2 \log \left(\frac{\det(V_t)}{\det(V)} \right) \leq 2d \log \left(\frac{\text{trace}(V) + tL^2}{d} \right) - 2 \log \det(V).$$

Furthermore, if $\lambda_{\min}(V) \geq \max(1, L^2)$ then

$$\sum_{i=1}^t \|x_i\|_{V_{i-1}}^2 \leq 2 \log \frac{\det(V_t)}{\det(V)} \leq 2d \log \frac{\text{trace}(V) + tL^2}{d}.$$

The next result appears in (Zanette & Brunskill, 2019), proposition 2, lines 48-51, and is a variation of a result utilized by (Azar et al., 2017).

Lemma 31 (Standard Deviation Difference (Zanette & Brunskill, 2019)). *Let $V_1, V_2 : \mathcal{S} \rightarrow \mathbb{R}$ be fixed mappings. Let $P(s)$ be a probability measure over the state space. Then, $\sqrt{\text{Var}(V_1)} - \sqrt{\text{Var}(V_2)} \leq \sqrt{\text{Var}(V_1 - V_2)}$.*

Lemma 32 (Law of Total Variance, e.g., (Azar et al., 2017; Zanette & Brunskill, 2019)). *For any π the following holds.*

$$\mathbb{E} \left[\sum_{h=1}^H \text{Var}_{P_h(\cdot | s_h, a_h)}(V_{h+1}^\pi | \pi) \right] = \mathbb{E} \left[\left(\sum_{h=1}^H r(s_h, a_h) - V_1^\pi(s_1) \right)^2 | \pi \right].$$

I. Useful Identities

Lemma 33. Let $C = \{\theta : \|\bar{\theta} - \theta\|_A \leq \alpha\}$ for some $\bar{\theta} \in \mathbb{R}^d$, $\alpha \in \mathbb{R}$ and $A \in \mathbb{R}^{d \times d}$ a PD matrix. Then, for any $x \in \mathbb{R}^d$

$$\max_{\theta \in C} \langle x, \theta - \bar{\theta} \rangle = \max_{\theta \in C} \langle x, \bar{\theta} - \theta \rangle = \alpha \|x\|_{A^{-1}}.$$

Proof. The following relations hold.

$$\begin{aligned} \max_{\theta \in C} \langle x, \theta - \bar{\theta} \rangle &= \max_{\theta \in C} \langle A^{-1/2}x, A^{1/2}(\theta - \bar{\theta}) \rangle \\ &\leq \|x\|_{A^{-1}} \max_{\theta \in C} \|\theta - \bar{\theta}\|_A && \text{(Cauchy-Schwartz Inequality)} \\ &\leq \alpha \|x\|_{A^{-1}}. && \text{(Definition of the set } C) \end{aligned}$$

Setting $\theta - \bar{\theta} = \frac{\alpha}{\|x\|^2} A^{-1/2}x$ the above inequalities hold with equality (see that $\|\theta - \bar{\theta}\|_A = \alpha$ which implies that $\theta - \bar{\theta} \in C$). Thus,

$$\max_{\theta \in C} \langle x, \theta - \bar{\theta} \rangle = \alpha \|x\|_{A^{-1}}.$$

Lastly, since $\|x\|_{A^{-1}} = \|-x\|_{A^{-1}}$ we conclude that

$$\max_{\theta \in C} \langle x, \bar{\theta} - \theta \rangle = \max_{\theta \in C} \langle -x, \theta - \bar{\theta} \rangle = \alpha \|-x\|_{A^{-1}} = \alpha \|x\|_{A^{-1}}.$$

□