

# Minimax Regret for Stochastic Shortest Path with Adversarial Costs and Known Transition

**Liyu Chen**

*University of Southern California*

LIYUC@USC.EDU

**Haipeng Luo**

*University of Southern California*

HAIPENGL@USC.EDU

**Chen-Yu Wei**

*University of Southern California*

CHENYU.WEI@USC.EDU

**Editors:** Mikhail Belkin and Samory Kpotufe

## Abstract

We study the stochastic shortest path problem with adversarial costs and known transition, and show that the minimax regret is  $\tilde{O}(\sqrt{DT_*K})$  and  $\tilde{O}(\sqrt{DT_*SAK})$  for the full-information setting and the bandit feedback setting respectively, where  $D$  is the diameter,  $T_*$  is the expected hitting time of the optimal policy,  $S$  is the number of states,  $A$  is the number of actions, and  $K$  is the number of episodes. Our results significantly improve upon the recent work of (Rosenberg and Mansour, 2020) which only considers the full-information setting and achieves suboptimal regret. Our work is also the first to consider bandit feedback with adversarial costs.

Our algorithms are built on top of the Online Mirror Descent framework with a variety of new techniques that might be of independent interest, including an improved multi-scale expert algorithm, a reduction from general stochastic shortest path to a special loop-free case, a skewed occupancy measure space, and a novel correction term added to the cost estimators. Interestingly, the last two elements reduce the variance of the learner via positive bias and the variance of the optimal policy via negative bias respectively, and having them simultaneously is critical for obtaining the optimal high-probability bound in the bandit feedback setting.

## 1. Introduction

We study the stochastic shortest path (SSP) problem, where a learner tries to reach a goal state in a Markov Decision Process (MDP) with minimum total cost. The problem proceeds in  $K$  episodes. In each episode, the learner starts from a fixed state, sequentially selects an available action, incurs a cost, and transits to the next state sampled from a fixed transition function. The episode ends when the learner reaches a fixed goal state. The performance of the learner is measured by her regret, which is the difference between her total cost over the  $K$  episodes and that of the best fixed policy.

The special case of SSP where the learner is guaranteed to reach the goal state within a fixed number of steps is extensively studied in recent years. It is often known as episodic finite-horizon reinforcement learning or equivalently loop-free SSP. The general case, however, is much less understood. Recently, Tarbouriech et al. (2020) and Cohen et al. (2020) study the case where the costs are fixed or generated stochastically and develop algorithms with sub-linear regret. Another recent work by Rosenberg and Mansour (2020) considers adversarial costs that are chosen arbitrarily but revealed at the end of each episode (the so-called full-information setting). When the transition

Table 1: Summary of our minimax optimal results and comparisons with prior work. Here,  $D, S, A$  are the diameter, number of states, and number of actions of the MDP,  $c_{\min}$  is the minimum cost,  $T_{\star} \leq D/c_{\min}$  is the expected hitting time of the optimal policy, and  $K$  is the number of episodes. Logarithmic terms are omitted. All algorithms can be implemented efficiently. [Algorithm 2](#) is completely parameter-free, while others require the knowledge of  $T_{\star}$ .

	Minimax Regret ( <b>this work</b> )	( <a href="#">Rosenberg and Mansour, 2020</a> )
Full information	$\Theta(\sqrt{DT_{\star}K})$ <a href="#">Algorithm 2</a> (expected bound) <a href="#">Algorithm 3</a> (high probability bound) <a href="#">Theorem 3</a> (lower bound)	$\tilde{O}\left(\frac{D}{c_{\min}}\sqrt{K}\right)$ or $\tilde{O}\left(\sqrt{DT_{\star}K^{\frac{3}{4}}}\right)$
Bandit feedback	$\Theta(\sqrt{DT_{\star}SAK})$ <a href="#">Algorithm 4</a> (expected bound) <a href="#">Algorithm 5</a> (high probability bound) <a href="#">Theorem 10</a> (lower bound)	N/A

function is known, their algorithm achieves  $\tilde{O}\left(\frac{D}{c_{\min}}\sqrt{K}\right)$  regret where  $D$  is the diameter of the MDP and  $c_{\min} \in (0, 1]$  is a global lower bound of the cost for any state-action pair. When  $c_{\min} = 0$ , they provide a different algorithm with regret  $\tilde{O}(\sqrt{DT_{\star}K^{3/4}})$  where  $T_{\star}$  is the expected time for the optimal policy to reach the goal state. They also further study the case with unknown transition.

In this work, we significantly improve the state-of-the-art for the general SSP problem with adversarial costs and known transition, by developing matching upper and lower bounds for both the full-information setting and the bandit feedback setting. More specifically, our results are (see also [Table 1](#) for a summary):

- In the full-information setting, we show that the minimax regret is of order  $\Theta(\sqrt{DT_{\star}K})$  (ignoring logarithmic terms), with no dependence on  $1/c_{\min}$  (it can be shown that  $T_{\star} \leq D/c_{\min}$ ). We develop two algorithms, one with optimal expected regret ([Algorithm 2](#)) and another with optimal high probability regret ([Algorithm 3](#)). Note that, as pointed out by [Rosenberg and Mansour \(2020\)](#), achieving high probability bounds for SSP is significantly more challenging even in the full-information setting, since the learner is often not guaranteed to reach the goal within a fixed number of steps with high probability. We complement our algorithms and upper bounds with a matching lower bound in [Theorem 3](#).
- Next, we further consider the more challenging bandit feedback setting where the learner only observes the cost for the visited state-action pairs, which has not been studied before in the adversarial cost case to the best of our knowledge. We show that the minimax regret is of order  $\Theta(\sqrt{DT_{\star}SAK})$  (ignoring logarithmic terms) where  $S$  is the number of states and  $A$  is the number of actions. We again developed two algorithms, one with optimal expected regret ([Algorithm 4](#)) and another more complex one with optimal high probability regret ([Algorithm 5](#)). A matching lower bound is shown in [Theorem 10](#).

**Techniques** Similarly to (Rosenberg and Mansour, 2020), our algorithms are all based on the standard Online Mirror Descent (OMD) framework. However, a variety of new techniques are developed on top of OMD to achieve our results. For example, to obtain the optimal expected regret in the full-information setting without knowing  $T_*$  ahead of time, we reduce the problem to the *multi-scale expert* problem studied in (Bubeck et al., 2017; Foster et al., 2017; Cutkosky and Orabona, 2018) and develop a new algorithm with an improved guarantee necessary to achieve our results, which might be of independent interest.<sup>1</sup>

Our other algorithms all require a reduction from a general SSP instance to its loop-free version (Definition 5) as well as executing OMD over a *skewed occupancy measure space*, both of which are novel as far as we know. The skewed occupancy measure can be viewed as adding *positive* bias to the costs, as a way to reduce the variance of the learner. These algorithms require setting some parameters in terms of  $T_*$  to achieve the optimal regret though (see discussions after Theorem 8).

In addition, the two algorithms in the bandit feedback setting require the usage of the log-barrier regularizer, an increasing learning rate schedule similar to (Lee et al., 2020a), and injecting another *negative* bias term into the cost estimator to reduce the variance of the optimal policy. We find the necessity of both positive and negative bias in the bandit setting intriguing.

**Related work** Earlier research studies SSP as a control problem and focuses on finding the optimal policy efficiently with all the parameters known; see for example (Bertsekas and Tsitsiklis, 1991; Bertsekas and Yu, 2013). Learning with low regret in SSP was first studied in (Tarbouriech et al., 2020), which considers fixed or stochastic costs and proposes algorithms with sub-linear regret that depends on  $1/c_{\min}$ . Cohen et al. (2020) remove the  $1/c_{\min}$  dependence and propose an algorithm with almost optimal regret. Note that their bounds do not depend on the parameter  $T_*$ ; see our discussions after Theorem 3 on why  $T_*$  shows up in the adversarial cost case.

To the best of our knowledge, (Rosenberg and Mansour, 2020) is the only existing work that studies SSP with adversarial costs. They only study the full-information setting, with either known or unknown transition, while we consider both the full-information setting and the bandit feedback setting, but only with known transition. We note that our loop-free reduction is readily applied to the unknown transition case, but it only leads to some suboptimal bounds (details omitted).

As mentioned, the special case of SSP with a fixed horizon is extensively studied in recent years, for both stochastic costs (see e.g., (Azar et al., 2017; Jin et al., 2018; Zanette and Brunskill, 2019; Shani et al., 2020)) and adversarial costs (see e.g., (Neu et al., 2012; Zimin and Neu, 2013; Rosenberg and Mansour, 2019; Jin et al., 2020)). The latter also heavily relies on the OMD framework, but the occupancy measure space that OMD operates over is much simpler compared to general SSP. Note that, although one of our key algorithmic ideas is to reduce general SSP to this special case, it does not mean that one can directly apply these existing algorithms after the reduction, as it only leads to suboptimal bounds. Instead, one must further utilize different properties of the original SSP instance to achieve the minimax regret, as we will discuss in detail.

## 2. Preliminaries

A stochastic shortest path (SSP) instance is defined by an MDP  $M = (\mathcal{S}, s_0, g, \mathcal{A}, P)$  and a sequence of  $K$  cost functions  $\{c_k\}_{k=1}^K$ . Here,  $\mathcal{S}$  is a finite state space,  $s_0 \in \mathcal{S}$  is the initial state,

1. See also concurrent work (Chen et al., 2021) by the same authors for in-depth discussions and significant extensions of this idea.

$g \notin \mathcal{S}$  is the goal state, and  $\mathcal{A} = \{\mathcal{A}_s\}_{s \in \mathcal{S}}$  is a finite action space where  $\mathcal{A}_s$  is the available action set at state  $s$ . We denote by  $\Gamma = \{(s, a) : s \in \mathcal{S}, a \in \mathcal{A}_s\}$  the set of valid state-action pairs, and by  $S = |\mathcal{S}|$  and  $A = (\sum_{s \in \mathcal{S}} |\mathcal{A}_s|) / S$  the total number of states and the average number of available actions respectively. The transition function  $P : \Gamma \times \mathcal{S} \cup \{g\} \rightarrow [0, 1]$  is such that  $P(s' | s, a)$  is the probability of transiting to  $s'$  after taking action  $a \in \mathcal{A}_s$  at state  $s$ , and it satisfies  $\sum_{s' \in \mathcal{S} \cup \{g\}} P(s' | s, a) = 1$  for each  $(s, a) \in \Gamma$ . Finally, the cost function  $c_k : \Gamma \rightarrow [0, 1]$  specifies the cost for each state-action pair during episode  $k$ .

The learning protocol is as follows. The learner interacts with a known MDP  $M$  through  $K$  episodes. In each episode  $k = 1, \dots, K$ , the environment adaptively decides the cost function  $c_k$ , which can depend on the learner's algorithm and the randomness before episode  $k$ . Simultaneously, starting from the initial state  $s_0 \in \mathcal{S}$ , the learner sequentially selects an action and transits to the next state according to the transition function, until reaching the goal state  $g$ . More formally, in each step  $i$  of the episode, the learner observes its current state  $s_k^i$  (with  $s_k^1 = s_0$  always). If  $s_k^i \neq g$ , the learner selects an action  $a_k^i \in \mathcal{A}_{s_k^i}$  and moves to the next state  $s_k^{i+1}$  sampled from  $P(\cdot | s_k^i, a_k^i)$ . The episode ends when the current state is the goal state, and we denote by  $I_k$  the number of steps in this episode such that  $s_k^{I_k+1} = g$ .

We consider two different types of feedback on the cost functions for the learner after the goal state is reached. In the *full-information* setting, the entire cost function  $c_k$  is revealed to the learner, while in the *bandit feedback* setting, only the costs for the visited state-action pairs, that is,  $c_k(s_k^i, a_k^i)$  for  $i = 1, \dots, I_k$ , are revealed to the learner.

**Proper policies and related concepts** Before discussing the goal of the learner, we introduce several necessary concepts. A stationary policy is a mapping  $\pi$  such that  $\pi(a|s)$  specifies the probability of taking action  $a \in \mathcal{A}_s$  in state  $s$ . It is deterministic if  $\pi(\cdot|s)$  concentrates on one single action (denoted by  $\pi(s)$ ) for all  $s$ . It is *proper* if executing it in the MDP starting from any state ensures that the goal state is reached within a finite number of steps with probability 1 (otherwise it is called improper). The set of all deterministic and proper policies is denoted by  $\Pi_{\text{proper}}$ . Following (Rosenberg and Mansour, 2020), we make the basic assumption  $\Pi_{\text{proper}} \neq \emptyset$ .

Let  $T^\pi(s)$  denote the expected hitting time it takes for  $\pi$  to reach  $g$  starting from state  $s$ . If  $\pi$  is proper, then  $T^\pi(s) < \infty$  for any state  $s$ . The *fast policy*  $\pi^f$  is the (deterministic) policy that achieves the minimum expected hitting time starting from any state, and the diameter of the MDP is defined as  $D = \max_{s \in \mathcal{S}} \min_{\pi \in \Pi_{\text{proper}}} T^\pi(s) = \max_{s \in \mathcal{S}} T^{\pi^f}(s)$ . Note that both  $\pi^f$  and  $D$  can be computed ahead of time since we consider the known transition setting.

Given a cost function  $c$  and a proper policy  $\pi$ , we define the *cost-to-go function*  $J^\pi : \mathcal{S} \rightarrow [0, \infty)$  such that  $J^\pi(s) = \mathbb{E} \left[ \sum_{i=1}^I c(s^i, a^i) \middle| P, \pi, s^1 = s \right]$ , where the expectation is over the randomness of the action  $a^i$  drawn from  $\pi(\cdot|s^i)$ , the state  $s^{i+1}$  drawn from  $P(\cdot|s^i, a^i)$ , and the number of steps  $I$  before reaching  $g$ . We use  $J_k^\pi$  to denote the cost-to-go function with respect to the cost  $c_k$ .

**Learning objective** The learner's goal is to minimize her *regret*, defined as the difference between her total cost and the total expected cost of the best deterministic proper policy in hindsight:  $R_K = \sum_{k=1}^K \sum_{i=1}^{I_k} c_k(s_k^i, a_k^i) - \sum_{k=1}^K J_k^{\pi^*}(s_0)$ , where  $\pi^* \in \operatorname{argmin}_{\pi \in \Pi_{\text{proper}}} \sum_{k=1}^K J_k^\pi(s_0)$ . By the Markov property, it is clear that  $\pi^*$  is in fact also the optimal policy starting from any other state, that is,  $\pi^* \in \operatorname{argmin}_{\pi \in \Pi_{\text{proper}}} \sum_{k=1}^K J_k^\pi(s)$  for any  $s \in \mathcal{S}$ . Two quantities related to  $\pi^*$  play an important role in our analysis: its expected hitting time starting from the initial state  $T_\star = T^{\pi^*}(s_0)$  and its largest expected hitting time starting from any state  $T_{\max} = \max_s T^{\pi^*}(s)$ . Let

$c_{\min} = \min_k \min_{(s,a)} c_k(s, a)$  be the minimum cost, and  $s_{\max} \in \mathcal{S}$  be such that  $T_{\max} = T^{\pi^*}(s_{\max})$ . We have  $T_{\max} c_{\min} \leq J_k^{\pi^*}(s_{\max})$  and  $J_k^{\pi^f}(s_{\max}) \leq D$  by definition. Together with the fact  $\sum_{k=1}^K J_k^{\pi^*}(s_{\max}) \leq \sum_{k=1}^K J_k^{\pi^f}(s_{\max})$ , this implies  $T_{\star} \leq T_{\max} \leq \frac{D}{c_{\min}}$  if  $c_{\min} > 0$  (which is one of the reasons why  $c_{\min}$  shows up in existing results).

**Occupancy measure** For a fixed MDP, a proper policy  $\pi$  induces an occupancy measure  $q_{\pi} \in \mathbb{R}_{\geq 0}^{\Gamma}$  such that  $q_{\pi}(s, a)$  is the expected number of visits to  $(s, a)$  when executing  $\pi$ , that is:  $q_{\pi}(s, a) = \mathbb{E} \left[ \sum_{i=1}^I \mathbb{I}\{s^i = s, a^i = a\} \mid P, \pi, s^1 = s_0 \right]$ . Similarly,  $q_{\pi}(s) = \sum_{a \in \mathcal{A}_s} q_{\pi}(s, a)$  is the expected number of visits to  $s$  when executing  $\pi$ . Clearly, we have  $J_k^{\pi}(s_0) = \sum_{(s,a) \in \Gamma} q_{\pi}(s, a) c_k(s, a) = \langle q_{\pi}, c_k \rangle$ , and if the learner executes a stationary proper policy  $\pi_k$  in episode  $k$ , then the expected regret can be written as

$$\mathbb{E}[R_K] = \mathbb{E} \left[ \sum_{k=1}^K J_k^{\pi_k}(s_0) - J_k^{\pi^*}(s_0) \right] = \mathbb{E} \left[ \sum_{k=1}^K \langle q_{\pi_k} - q_{\pi^*}, c_k \rangle \right], \quad (1)$$

converting the problem into a form of online linear optimization and making Online Mirror Descent a natural solution to the problem. Note that, given a function  $q : \Gamma \rightarrow [0, \infty)$ , if it corresponds to an occupancy measure, then the corresponding policy  $\pi_q$  can clearly be obtained by  $\pi_q(a|s) \propto q(s, a)$ . Also note that  $T^{\pi}(s_0) = \sum_{(s,a)} q_{\pi}(s, a) = \sum_{s \in \mathcal{S}} q_{\pi}(s)$ .

**Other notations** We let  $N_k(s, a)$  denote the (random) number of visits of the learner to  $(s, a)$  during episode  $k$ , so that the regret can be re-written as  $R_K = \sum_{k=1}^K \langle N_k - q_{\pi^*}, c_k \rangle$ . Throughout the paper, we use the notation  $\langle f, g \rangle$  as a shorthand for  $\sum_{s \in \mathcal{S}} f(s)g(s)$ ,  $\sum_{(s,a)} f(s, a)g(s, a)$ , or  $\sum_{h=1}^H \sum_{(s,a)} f(s, a, h)g(s, a, h)$  when  $f$  and  $g$  are functions in  $\mathbb{R}^{\mathcal{S}}$ ,  $\mathbb{R}^{\Gamma}$ , or  $\mathbb{R}^{\Gamma \times [H]}$  (for some  $H$ ) respectively. Let  $\mathcal{F}_k$  denote the  $\sigma$ -algebra of events up to the beginning of episode  $k$ , and  $\mathbb{E}_k$  be a shorthand of  $\mathbb{E}[\cdot | \mathcal{F}_k]$ . For a convex function  $\psi$ , the Bregman divergence between  $u$  and  $v$  is defined as:  $D_{\psi}(u, v) = \psi(u) - \psi(v) - \langle \nabla \psi(v), u - v \rangle$ . For an integer  $n$ ,  $[n]$  denotes the set  $\{1, \dots, n\}$ .

### 3. Minimax Regret for the Full-information Setting

In this section, we consider the simpler full-information setting where the learner observes  $c_k$  in the end of episode  $k$ . Somewhat surprisingly, even in this case, ensuring optimal regret is rather challenging. We first propose an algorithm with expected regret  $\tilde{O}(\sqrt{DT_{\star}K})$  and a matching lower bound in Section 3.1. Notably, our algorithm is parameter-free and does not need to know  $T_{\star}$  ahead of time.<sup>2</sup> Next, in Section 3.2, by converting the problem into another loop-free SSP instance and using a skewed occupancy measure space, we develop an algorithm that achieves the same regret bound with high probability, although this requires the knowledge of  $T_{\star}$ .

#### 3.1. Optimal expected regret

To introduce our algorithm, we first briefly review the SSP-O-REPS algorithm of Rosenberg and Mansour (2020), which only achieves regret  $\tilde{O}(\frac{D}{c_{\min}} \sqrt{K})$ . The idea is to run the standard Online Mirror Descent (OMD) algorithm over an appropriate occupancy measure space. Specifically, they

2. The knowledge of  $K$  is also unnecessary due to the standard doubling trick.

**Algorithm 1** SSP-O-REPS**Input:** upper bound on expected hitting time  $T$ .**Define:** regularizer  $\psi(q) = \frac{1}{\eta} \sum_{(s,a)} q(s,a) \ln q(s,a)$  and  $\eta = \min \left\{ \frac{1}{2}, \sqrt{\frac{T \ln(SAT)}{DK}} \right\}$ .**Initialization:**  $q_1 = \operatorname{argmin}_{q \in \Delta(T)} \psi(q)$  where  $\Delta(T)$  is defined in Eq. (2).**for**  $k = 1, \dots, K$  **do**    Execute  $\pi_{q_k}$ , receive  $c_k$ , and update  $q_{k+1} = \operatorname{argmin}_{q \in \Delta(T)} \langle q, c_k \rangle + D_\psi(q, q_k)$ .**Algorithm 2** Adaptive SSP-O-REPS with Multi-scale Experts**Define:**  $j_0 = \lceil \log_2 T^{\pi^f}(s_0) \rceil - 1$ ,  $b(j) = 2^{j_0+j}$ ,  $\eta_j = \frac{1}{\sqrt{b(j)K \max\{D, 16\}}}$ ,  $N = \lceil \log_2 K \rceil - j_0$ .**Define:**  $\Omega = \{p \in \mathbb{R}_{\geq 0}^N : \sum_{j=1}^N p(j) = 1\}$  and  $\psi(p) = \sum_{j=1}^N \frac{1}{\eta_j} p(j) \ln p(j)$ .**Initialize:**  $p_1 \in \Omega$  such that  $p_1(j) = \frac{\eta_j}{N\eta_1}$ ,  $\forall j \in [N]$ .**Initialize:**  $N$  instances of [Algorithm 1](#), where the  $j$ -th instance uses parameter  $T = b(j)$ .**for**  $k = 1, \dots, K$  **do**

- 1   For each  $j \in [N]$ , obtain occupancy measure  $q_k^j$  from SSP-O-REPS instance  $j$ .
- 2   Sample  $j_k \sim p_k$ , execute the policy induced by  $q_k^{j_k}$ , receive  $c_k$ , and feed  $c_k$  to all instances.
- 3   Compute  $\ell_k$  and  $a_k$ :  $\ell_k(j) = \langle q_k^j, c_k \rangle$ ,  $a_k(j) = 4\eta_j \ell_k^2(j)$ ,  $\forall j \in [N]$ .
- 4   Update  $p_{k+1} = \operatorname{argmin}_{p \in \Omega} \langle p, \ell_k + a_k \rangle + D_\psi(p, p_k)$ .

define the occupancy measure space parameterized by size  $T > 0$  as:

$$\Delta(T) = \left\{ q \in \mathbb{R}_{\geq 0}^\Gamma : \sum_{(s,a) \in \Gamma} q(s,a) \leq T, \right. \\ \left. \sum_{a \in \mathcal{A}_s} q(s,a) - \sum_{(s',a') \in \Gamma} P(s|s', a') q(s', a') = \mathbb{I}\{s = s_0\}, \forall s \in \mathcal{S} \right\}. \quad (2)$$

It is shown that every  $q \in \Delta(T)$  is a valid occupancy measure induced by the policy  $\pi_q$  (recall  $\pi_q(a|s) \propto q(s,a)$ ). Therefore, as long as  $T$  is large enough such that  $q_{\pi^*} \in \Delta(T)$ , based on [Eq. \(1\)](#), the problem is essentially translated to an instance of online linear optimization and can be solved by maintaining a sequence of occupancy measures  $q_1, \dots, q_K$  updated according to OMD:  $q_{k+1} = \operatorname{argmin}_{q \in \Delta(T)} \langle q, c_k \rangle + D_\psi(q, q_k)$ , where  $\psi$  is a regularizer with the default choice being the negative entropy  $\psi(q) = \frac{1}{\eta} \sum_{(s,a)} q(s,a) \ln q(s,a)$  for some learning rate  $\eta > 0$ . See [Algorithm 1](#) for the pseudocode and ([Rosenberg and Mansour, 2020](#)) for the details of implementing it efficiently.

[Rosenberg and Mansour \(2020\)](#) show that as long as  $q_{\pi^*} \in \Delta(T)$ , [Algorithm 1](#) ensures  $\mathbb{E}[R_K] = \tilde{O}(T\sqrt{K})$ . To ensure  $q_{\pi^*} \in \Delta(T)$ , they set  $T = \frac{D}{c_{\min}}$  because  $\sum_{(s,a) \in \Gamma} q_{\pi^*}(s,a) = T_* \leq \frac{D}{c_{\min}}$  as discussed in [Section 2](#). This leads to their final regret bound  $\tilde{O}(\frac{D}{c_{\min}} \sqrt{K})$ .

We improve their approach using the following two ideas. First, we show a more careful analysis for the same [Algorithm 1](#) and use the fact that the total expected cost of  $\pi^*$  is bounded by  $DK$  instead of  $TK$  to obtain the following stronger guarantee.

**Lemma 1** *If  $T$  is such that  $q_{\pi^*} \in \Delta(T)$ , then [Algorithm 1](#) guarantees:  $\mathbb{E}[R_K] = \tilde{O}(\sqrt{DTK})$ .*

By using the same  $T = \frac{D}{c_{\min}}$ , this already leads to a better bound  $\tilde{O}(D\sqrt{K/c_{\min}})$ . If  $T_\star$  was known, setting  $T = T_\star$  would also immediately give the claimed bound  $\tilde{O}(\sqrt{DT_\star K})$  (since  $q_{\pi^\star} \in \Delta(T_\star)$ ), which is optimal as we show later.

The second new idea of our approach is thus to deal with unknown  $T_\star$  by learning it on the fly via another online learning meta-algorithm (Algorithm 2). Specifically, we maintain roughly  $\ln K$  instances of Algorithm 1, where the  $j$ -th instance sets the parameter  $T$  as  $b(j)$  which is roughly  $2^j$ , so that there always exists an instance  $j^\star$  with  $b(j^\star)$  very close to the unknown  $T_\star$ . The meta-algorithm treats each instance as an expert, and in each episode, samples one of these experts and follows its policy (Line 2). If the regret of this meta-algorithm to instance  $j^\star$  is no larger than  $\tilde{O}(\sqrt{DT_\star K})$ , then the overall regret to  $\pi^\star$  would clearly also be  $\tilde{O}(\sqrt{DT_\star K})$ .

While seemingly this appears to be a classic expert problem and might be solved by the standard Hedge algorithm (Freund and Schapire, 1997), the key challenge is that the loss for each expert  $j$ , denoted by  $\ell_k(j) = \langle q_k^j, c_k \rangle$  (for episode  $k$ ), has a different scale. Indeed, we have  $\ell_k(j) \leq \sum_{(s,a)} q_k^j(s,a) \leq b(j)$ . Standard algorithms such as Hedge have a regret bound that depends on a uniform upper bound of all losses as large as  $b(N) \approx K$  in our case, leading to a vacuous bound. More advanced “multi-scale” algorithms (Bubeck et al., 2017; Foster et al., 2017; Cutkosky and Orabona, 2018) mitigate the issue and ensure regret  $\tilde{O}(b(j^\star)\sqrt{K})$  comparing to expert  $j^\star$ , which still leads to  $\tilde{O}(T_\star\sqrt{K})$  regret overhead and ruins the final bound.

To address this challenge, we propose a new multi-scale expert algorithm with regret bound  $\tilde{O}\left(\sqrt{b(j^\star)\mathbb{E}[\sum_{k=1}^K \ell_k(j^\star)]}\right)$ , which is always no worse than previous works since  $\sum_{k=1}^K \ell_k(j^\star) \leq b(j^\star)K$ . The algorithm is similar to that of (Bubeck et al., 2017) which is OMD over the  $(N-1)$ -dimensional simplex with a weighted negative entropy regularizer  $\psi(p) = \sum_{j=1}^N \frac{1}{\eta_j} p(j) \ln p(j)$ . Here, each expert uses a different learning rate  $\eta_j$  that depends on the corresponding scale  $b(j)$ . The key difference of our algorithm is that we also add a correction term  $a_k(j) = 4\eta_j \ell_k^2(j)$  to the loss  $\ell_k(j)$  (Line 4), an idea used in previous works such as (Steinhardt and Liang, 2014; Wei and Luo, 2018) to obtain a bound in terms of the loss of the benchmark  $\sum_{k=1}^K \ell_k(j^\star)$ . Another important tweak is to set the initial distribution for expert  $j \neq 1$  to be  $\frac{\eta_j}{N\eta_1}$ . We note that this new and improved multi-scale expert algorithm might be of independent interest.

To see why this improved bound  $\tilde{O}\left(\sqrt{b(j^\star)\mathbb{E}[\sum_{k=1}^K \ell_k(j^\star)]}\right)$  helps, note that Lemma 1 implies:  $\mathbb{E}[\sum_{k=1}^K \ell_k(j^\star)] \leq \mathbb{E}[\sum_{k=1}^K J_k^{\pi^\star}(s_0)] + \tilde{O}(\sqrt{DT_\star K}) \leq DK + \tilde{O}(\sqrt{DT_\star K})$ . Thus, the overhead of the meta-algorithm is of order  $\tilde{O}(\sqrt{DT_\star K})$  as desired. We summarize the final guarantee below.

**Theorem 2** *Algorithm 2 enjoys the following expected regret bound:  $\mathbb{E}[R_K] = \tilde{O}(\sqrt{DT_\star K})$ .*

**Lower bound** Our regret bound stated in Theorem 2 not only improves that of (Rosenberg and Mansour, 2020), but is also optimal up to logarithmic terms as shown in the following lower bound.

**Theorem 3** *For any  $D, T_\star, K$  with  $K \geq T_\star \geq D + 1$ , there exists an SSP instance such that its diameter is  $D + 2$ , the optimal policy has hitting time  $T_\star + 1$ , and the expected regret of any learner after  $K$  episodes is at least  $\Omega(\sqrt{DT_\star K})$  under the full-information and known transition setting.*

Similarly to most lower bound proofs, our proof also constructs an environment with stochastic costs and with a slightly better state hidden among other equally good states, and argues that the expected regret of any learner with respect to the randomness of the environment has to be

$\Omega(\sqrt{DT_*K})$ . At first glance, this appears to be a contradiction to existing results for SSP with stochastic costs (Tarbouriech et al., 2020; Cohen et al., 2020), where the optimal regret is independent of  $T_*$ . However, the catch is that “stochastic costs” has a different meaning in these works. Specifically, it refers to a setting where the cost for each state-action pair is drawn independently from a fixed distribution *every time* it is visited, and is revealed to the learner immediately. On the other hand, “stochastic costs” in our lower bound proof refers to a setting where at the beginning of each episode  $k$ ,  $c_k$  is sampled once from a fixed distribution and then *fixed* throughout the episode. Moreover, it is revealed only after the episode ends. It can be shown that our setting is harder due to the larger variance of costs, explaining our larger lower bound and the seemingly contradiction.

### 3.2. Optimal high-probability regret

To obtain a high-probability regret bound, one needs to control the deviation between the actual total cost of the learner  $\sum_{k=1}^K \langle N_k, c_k \rangle$  and its expectation  $\sum_{k=1}^K \langle q_k, c_k \rangle$ . While for most online learning problems with full information, similar deviation can be easily controlled by the Azuma’s inequality, this is not true for SSP as pointed out in (Rosenberg and Mansour, 2020), due to the lack of an almost sure upper bound on the random variable  $\langle N_k, c_k \rangle$ . Rosenberg and Mansour (2020, Lemma E.1) point out that with high probability  $\sum_{(s,a)} N_{\pi^*}(s, a)$  is bounded by  $T_{\max}$ , and thus it is natural to enforce the same for  $N_k$ . However, this at best leads to a bound of order  $\tilde{O}(\sqrt{DT_{\max}K})$ . To achieve the optimal regret, we start with a closer look at the variance of the actual cost of any policy, showing that it is in fact related to the corresponding cost-to-go function.

**Lemma 4** *Consider executing a stationary policy  $\pi$  in episode  $k$ . Then  $\mathbb{E}[\langle N_k, c_k \rangle^2] \leq 2 \langle q_\pi, J_k^\pi \rangle$ .*

For the optimal policy  $\pi^*$ , although  $\langle q_{\pi^*}, J_k^{\pi^*} \rangle$  can still be as large as  $T_{\max}$ , one key observation is that the sum of these quantities over  $K$  episodes is at most  $DT_*K$  since  $\sum_{k=1}^K \langle q_{\pi^*}, J_k^{\pi^*} \rangle = \sum_{s \in \mathcal{S}} q_{\pi^*}(s) \sum_{k=1}^K J_k^{\pi^*}(s) \leq DK \sum_{s \in \mathcal{S}} q_{\pi^*}(s) = DT_*K$ , where the inequality is again due to the optimality of  $\pi^*$  and the existence of the fast policy  $\pi^f$ :  $\sum_{k=1}^K J_k^{\pi^*}(s) \leq \sum_{k=1}^K J_k^{\pi^f}(s) \leq DK$ . Given this observation, it is tempting to enforce that the learner’s policies  $\pi_1, \dots, \pi_K$  are also such that  $\sum_{k=1}^K \langle q_{\pi_k}, J_k^{\pi_k} \rangle \leq DT_*K$ , which would be enough to control the deviation between  $\sum_{k=1}^K \langle N_k, c_k \rangle$  and  $\sum_{k=1}^K \langle q_{\pi_k}, c_k \rangle$  by  $\tilde{O}(\sqrt{DT_*K})$  as desired by Freedman’s inequality. However, it is unclear how to enforce this constraint since it depends on all the cost functions unknown ahead of time. In fact, even if the cost functions were known, the constraint is also non-convex due to the complicated dependence of  $J_k^\pi$  on  $q_\pi$ . To address these issues, we propose two novel ideas.

**First idea: a loop-free reduction** Our first idea is to reduce the problem to a *loop-free* MDP so that the variance  $\mathbb{E}[\langle N_k, c_k \rangle^2]$  takes a much simpler form that is linear in both the occupancy measure and the cost function. Moreover, the reduction only introduces a small bias in the regret between the original problem and its loop-free version. The construction of the loop-free MDP is basically to duplicate each state by attaching a time step  $h$  for  $H_1$  steps, and then connect all states to some virtual *fast* state that lasts for another  $H_2$  steps. Formally, we define the following.

**Definition 5** *For an SSP instance  $M = (\mathcal{S}, s_0, g, \mathcal{A}, P)$  with cost functions  $c_{1:K}$ , we define, for horizon parameters  $H_1, H_2 \in \mathbb{N}$ , another loop-free SSP instance  $\tilde{M} = (\tilde{\mathcal{S}}, \tilde{s}_0, g, \tilde{\mathcal{A}}, \tilde{P})$  with cost function  $\tilde{c}_{1:K}$  as follows:*

- $\tilde{\mathcal{S}} = (\mathcal{S} \cup \{s_f\}) \times [H]$  where  $s_f$  is an artificially added “fast” state and  $H = H_1 + H_2$ ;



- $\tilde{s}_0 = (s_0, 1)$  and the goal state  $g$  remains the same;
- $\tilde{\mathcal{A}} = \mathcal{A} \cup \{a_f\}$ , where  $a_f$  is an artificially added action that is only available at  $(s_f, h)$  for  $h \in [H]$  (the available action set at  $(s, h)$  is  $\mathcal{A}_s$  for all  $s \neq s_f$  and  $h \in [H]$ );
- transition from  $(s, h)$  to  $(s', h')$  is only possible when  $h' = h + 1$ : for the first  $H_1$  layers, the transition follows the original MDP in the sense that  $\tilde{P}((s', h + 1)|(s, h), a) = P(s'|s, a)$  and  $\tilde{P}(g|(s, h), a) = P(g|s, a)$  for all  $h < H_1$  and  $(s, a) \in \Gamma$ ; from layer  $H_1$  to layer  $H$ , all states transit to the fast state:  $\tilde{P}((s_f, h + 1)|(s, h), a) = 1$  for all  $H_1 \leq h < H$  and  $(s, a) \in \tilde{\Gamma} \triangleq \Gamma \cup \{(s_f, a_f)\}$ ; finally, the last layer transits to the goal state always:  $\tilde{P}(g|(s, H), a) = 1$  for all  $(s, a) \in \tilde{\Gamma}$ ;
- cost function is such that  $\tilde{c}_k((s, h), a) = c_k(s, a)$  and  $\tilde{c}_k((s_f, h), a_f) = 1$  for all  $(s, a) \in \Gamma$  and  $h \in [H]$ ; for notational convenience, we also write  $\tilde{c}_k((s, h), a)$  as  $c_k(s, a, h)$ .

Note that in this definition, there are some redundant states such as  $(s, h)$  for  $s \in \mathcal{S}$  and  $h > H_1$  or  $(s_f, h)$  for  $h \leq H_1$  since they will never be visited. However, having these redundant states greatly simplifies our presentation. For notations related to the loop-free version, we often use a tilde symbol to distinguish them from the original counterparts (such as  $\tilde{M}$  and  $\tilde{\mathcal{S}}$ ), and for a function  $\tilde{f}((s, h), a)$  that takes a state in  $\tilde{M}$  and an action as inputs, we often simplify it as  $f(s, a, h)$  (such as  $c_k$  and  $q_k$ ). For such a function, we will also use the notation  $\mathbf{h} \circ f \in \mathbb{R}^{\tilde{\Gamma} \times [H]}$  such that  $(\mathbf{h} \circ f)(s, a, h) = h \cdot f(s, a, h)$ . Similarly, for a function  $f \in \mathbb{R}^{\tilde{\Gamma}}$ , we use the same notation  $\mathbf{h} \circ f \in \mathbb{R}^{\tilde{\Gamma} \times [H]}$  such that  $(\mathbf{h} \circ f)(s, a, h) = h \cdot f(s, a)$ .

As mentioned, one key reason of considering such a loop-free MDP is that the variance of the learner's actual cost takes a much simpler form that is linear in both the occupancy measure and the cost function, as shown in the lemma below (which is an analogue of [Lemma 4](#)).

**Lemma 6** Consider executing a stationary policy  $\tilde{\pi}$  in  $\tilde{M}$  in episode  $k$  and let  $\tilde{N}_k(s, a, h) \in \{0, 1\}$  denote the number of visits to state-action pair  $((s, h), a)$ . Then  $\mathbb{E}[(\tilde{N}_k, c_k)^2] \leq 2 \langle q_{\tilde{\pi}}, \mathbf{h} \circ c_k \rangle$ .

Next, we complete the reduction by describing how one can solve the original problem via solving its loop-free version. Given a policy  $\tilde{\pi}$  for  $\tilde{M}$ , we define a *non-stationary* policy  $\sigma(\tilde{\pi})$  for  $M$  as follows: for each step  $h \leq H_1$ , follow  $\tilde{\pi}(\cdot|(s, h))$  when at state  $s$ ; after the first  $H_1$  steps (if not reaching  $g$  yet), execute the fast policy  $\pi^f$  until reaching the goal state  $g$ . When executing  $\sigma(\tilde{\pi})$  in  $M$  for episode  $k$ , we overload the notation  $\tilde{N}_k$  defined in [Lemma 6](#) and let  $\tilde{N}_k(s, a, h)$  be 1 if  $(s, a)$  is visited at time step  $h \leq H_1$ , or 0 otherwise; and  $\tilde{N}_k(s_f, a_f, h)$  be 1 if  $H_1 < h \leq H$  and the goal state  $g$  is not reached within  $H_1$  steps, or 0 otherwise. Clearly,  $\tilde{N}_k(s, a, h)$  indeed follows the same distribution as the number of visits to state-action pair  $((s, h), a)$  when executing  $\tilde{\pi}$  in  $\tilde{M}$ . We also define a deterministic policy  $\tilde{\pi}^*$  for  $\tilde{M}$  that mimics the behavior of  $\pi^*$  in the sense that  $\tilde{\pi}^*(s, h) = \pi^*(s)$  for  $s \in \mathcal{S}$  and  $h \leq H_1$  (for larger  $h$ ,  $s$  has to be  $s_f$  and the only available action is  $a_f$ ). The next lemma shows that, as long as the horizon parameters  $H_1$  and  $H_2$  are set appropriately, this reduction makes sure that the regret between these two problems are similar.

**Lemma 7** Suppose  $H_1 \geq 8T_{\max} \ln K$ ,  $H_2 = \lceil 4D \ln \frac{4K}{\delta} \rceil$  and  $K \geq D$  for some  $\delta \in (0, 1)$ . Let  $\tilde{\pi}_1, \dots, \tilde{\pi}_K$  be policies for  $\tilde{M}$  with occupancy measures  $q_1, \dots, q_K \in [0, 1]^{\tilde{\Gamma} \times [H]}$ . Then the regret

---

**Algorithm 3** SSP-O-REPS with Loop-free Reduction and Skewed Occupancy Measure
 

---

**Input:** Upper bound on expected hitting time  $T$ , horizon parameter  $H_1$ , confidence level  $\delta$

**Parameters:**  $\eta = \min \left\{ \frac{1}{2}, \sqrt{\frac{T}{DK}} \right\}$ ,  $\lambda = \sqrt{\frac{\ln(1/\delta)}{DTK}}$ ,  $H_2 = \lceil 4D \ln \frac{4K}{\delta} \rceil$

**Define:**  $H = H_1 + H_2$ , regularizer  $\psi(\phi) = \frac{1}{\eta} \sum_{h=1}^H \sum_{(s,a) \in \tilde{\Gamma}} \phi(s, a, h) \ln \phi(s, a, h)$

**Define:** decision set  $\Omega = \{\phi = q + \lambda \mathbf{h} \circ q : q \in \tilde{\Delta}(T)\}$  (with  $\tilde{\Delta}(T)$  defined in Eq. (8))

**Initialization:**  $\phi_1 = q_1 + \lambda \mathbf{h} \circ q_1 = \operatorname{argmin}_{\phi \in \Omega} \psi(\phi)$ .

**for**  $k = 1, \dots, K$  **do**

Execute  $\sigma(\tilde{\pi}_k)$  where  $\tilde{\pi}_k$  is such that  $\tilde{\pi}_k(a|(s, h)) \propto q_k(s, a, h)$ , and receive  $c_k$ .

Update  $\phi_{k+1} = q_{k+1} + \lambda \mathbf{h} \circ q_{k+1} = \operatorname{argmin}_{\phi \in \Omega} \langle \phi, c_k \rangle + D_\psi(\phi, \phi_k)$ .

---

of executing  $\sigma(\tilde{\pi}_1), \dots, \sigma(\tilde{\pi}_K)$  in  $M$  satisfies: 1) for any  $\lambda \in (0, 2/H]$ , with probability  $1 - \delta$ ,

$$R_K \leq \sum_{k=1}^K \langle \tilde{N}_k - q_{\tilde{\pi}^*}, c_k \rangle + \tilde{\mathcal{O}}(1) \leq \underbrace{\sum_{k=1}^K \langle q_k - q_{\tilde{\pi}^*}, c_k \rangle}_{\text{REG}} + \lambda \underbrace{\sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle}_{\text{VAR}} + \frac{2 \ln(2/\delta)}{\lambda} + \tilde{\mathcal{O}}(1),$$

and 2)  $\mathbb{E}[R_K] \leq \mathbb{E}[\text{REG}] + \tilde{\mathcal{O}}(1)$ .

Note that the REG term is the expected regret (to  $\tilde{\pi}^*$ ) in  $\tilde{M}$  and can again be controlled by OMD. The VAR term comes from the derivation between the actual cost of the learner in  $\tilde{M}$  and its expectation, according to Freedman's inequality and Lemma 6. At this point, one might wonder whether directly applying an existing algorithm such as (Zimin and Neu, 2013) for loop-free MDPs solves the problem, since Lemma 7 shows that the regret in these two problems are close. Doing so, however, leads to a suboptimal bound of order  $\tilde{\mathcal{O}}(H\sqrt{K}) = \tilde{\mathcal{O}}((T_{\max} + D)\sqrt{K})$ . This is basically the same as trivially bounding VAR by  $H^2K$ . It is thus critical to better control this term using properties of the original problem, which requires the second idea described below.

**Second idea: skewed occupancy measure space** Similarly to earlier discussions, it can be shown that  $\sum_{k=1}^K \langle q_{\tilde{\pi}^*}, \mathbf{h} \circ c_k \rangle = \mathcal{O}(DT_*K)$  (Lemma 15), making it hopeful to bound VAR by the same. However, even though the variance now takes a simpler form, it is still unclear how to directly enforce the algorithm to satisfy  $\text{VAR} = \mathcal{O}(DT_*K)$ . Instead, we take a different route and make sure that the REG term is at most  $\tilde{\mathcal{O}}(\sqrt{DT_*K} + \lambda DT_*K) - \lambda \text{VAR}$ , thus canceling the variance term. To do so, thanks to the simple form of VAR, it suffices to inject a small positive bias into the action space of OMD, making it a *skewed* occupancy measure space:  $\Omega = \{\phi = q + \lambda \mathbf{h} \circ q : q \in \tilde{\Delta}(T_*)\}$  where  $\tilde{\Delta}(T_*)$  is the counterpart of  $\Delta(T_*)$  for  $\tilde{M}$  (see Eq. (8) in Appendix A for the spelled out definition). Indeed, by similar arguments from Section 3.1, operating OMD over this space ensures a bound of order  $\mathcal{O}(\sqrt{DT_*K})$  on the ‘‘skewed regret’’:  $\sum_{k=1}^K \langle (q_k + \lambda \mathbf{h} \circ q_k) - (q_{\tilde{\pi}^*} + \lambda \mathbf{h} \circ q_{\tilde{\pi}^*}), c_k \rangle = \text{REG} + \lambda \text{VAR} - \lambda \sum_{k=1}^K \langle q_{\tilde{\pi}^*}, \mathbf{h} \circ c_k \rangle$ , and we already know that the last term is of order  $\mathcal{O}(\lambda DT_*K)$ . Rearranging thus proves the desired bound on REG, and finally picking the optimal  $\lambda$  to trade off the term  $\frac{2 \ln(2/\delta)}{\lambda}$  leads to the optimal bound. We summarize the final algorithm in Algorithm 3 and its regret guarantee below. (Note that the algorithm can be implemented efficiently since  $\Omega$  is a convex polytope with  $\mathcal{O}(SAH)$  constraints.)

**Theorem 8** *If  $T \geq T_* + 1$ ,  $H_1 \geq 8T_{\max} \ln K$ , and  $K \geq H^2 \ln(\frac{1}{\delta})$ , then with probability at least  $1 - \delta$ , Algorithm 3 ensures  $R_K = \tilde{\mathcal{O}}(\sqrt{DTK \ln(1/\delta)})$ .*

To obtain the optimal bound, we need to set  $T = cT_\star + 1$  for any constant  $c \geq 1$ . Unfortunately, we are unable to extend the idea from [Section 3.1](#) to also learn this parameter on the fly, and we leave it as an important future direction (see [Section 5](#) for more discussions). Note that, however, in the construction of the lower bound ([Theorem 3](#)),  $T_\star$  is indeed known to the learner. Thus, assuming the knowledge of  $T_\star$  *does not make the problem any easier information-theoretically*. As for the parameter  $H_1$ , we can always set it to something large such as  $K^{1/3}$  so that the conditions of the theorem hold for large enough  $K$  (though leading to a larger time complexity of the algorithm).

We also remark that instead of injecting bias to the occupancy measure space, one can obtain the same by injecting a similar positive bias to the cost function. However, we use the former approach because it turns out to be critical for the bandit feedback setting that we consider in the next section.

#### 4. Minimax Regret for the Bandit Feedback Setting

We now consider the more challenging case with bandit feedback, that is, at the end of each episode, the learner only receives the cost of the visited state-action pairs. A standard technique in the adversarial bandit literature is to construct an importance-weighted cost estimator  $\hat{c}_k$  for  $c_k$  and then feed it to OMD, which is even applicable to learning loop-free SSP ([Zimin and Neu, 2013](#); [Jin et al., 2020](#); [Lee et al., 2020a](#)). For general SSP, the natural importance-weighted estimator  $\hat{c}_k$  is:  $\hat{c}_k(s, a) = \frac{N_k(s, a)c_k(s, a)}{q_k(s, a)}$  where  $N_k(s, a)$  is the number of visits to  $(s, a)$  and  $q_k$  is the occupancy measure of the policy executed in episode  $k$ . This is clearly unbiased since  $\mathbb{E}_k[N_k(s, a)] = q_k(s, a)$ .

However, it is well-known that unbiasedness alone is not enough — the variance of the estimator also plays a key role in the OMD analysis even if one only cares about expected regret. For example, if we still use the entropy regularizer as in [Section 3](#), the so-called stability term of OMD is in terms of the weighted variance  $\sum_{(s, a)} q_k(s, a) \mathbb{E}_k[\hat{c}_k^2(s, a)] = \sum_{(s, a)} \frac{\mathbb{E}_k[N_k^2(s, a)]c_k^2(s, a)}{q_k(s, a)}$ . While this term is nicely bounded in the loop-free case (since  $N_k(s, a)$  is binary and thus  $\mathbb{E}_k[N_k^2(s, a)] = q_k(s, a)$  cancels out the denominator), unfortunately it can be prohibitively large in the general case. In light of this, it might be tempting to use our loop-free reduction again and then directly apply an existing algorithm such as ([Zimin and Neu, 2013](#)). However, this again leads to a suboptimal bound with dependence on  $H = \tilde{O}(T_{\max})$ . It turns out that this is significantly more challenging than other bandit problems and requires a combination of various techniques, as described below.

**Log-barrier regularizer** Although the entropy regularizer is a classic choice for OMD to deal with bandit problems, in recent years, a line of research discovers various advantages of using a different regularizer called *log-barrier* (see e.g. ([Foster et al., 2016](#); [Agarwal et al., 2017](#); [Wei and Luo, 2018](#); [Luo et al., 2018](#); [Bubeck et al., 2019](#); [Kotłowski and Neu, 2019](#); [Lee et al., 2020b](#))). In our context, the log-barrier regularizer is  $-\frac{1}{\eta} \sum_{(s, a)} \ln q(s, a)$ , and it indeed leads to a smaller stability term in terms of  $\sum_{(s, a)} q_k^2(s, a) \mathbb{E}_k[\hat{c}_k^2(s, a)] = \sum_{(s, a)} \mathbb{E}_k[N_k^2(s, a)]c_k^2(s, a)$  (note the extra  $q_k(s, a)$  factor compared to the case of entropy). This term is further bounded by  $\mathbb{E}_k[\langle N_k, c_k \rangle^2]$ , which is exactly the variance of the learner’s actual cost considered in [Section 3.2](#)!

**Loop-free reduction and skewed occupancy measure** Based on the observation above, it is natural to apply the same ideas of loop-free reduction and skewed occupancy measure from [Section 3.2](#) to deal with the stability term  $\mathbb{E}_k[\langle N_k, c_k \rangle^2]$ . However, some extra care is needed when using log-barrier in the loop-free instance  $\tilde{M}$ . Indeed, directly using  $\psi(\phi) = -\frac{1}{\eta} \sum_h \sum_{(s, a)} \ln \phi(s, a, h)$  would lead to another term of order  $\tilde{O}(HSA/\eta)$  in the OMD analysis and ruin the bound. Instead,

---

**Algorithm 4** Log-barrier Policy Search for SSP
 

---

**Input:** Upper bound on expected hitting time  $T$  and horizon parameter  $H_1$ .

**Parameters:**  $\eta = \sqrt{\frac{SA}{DTK}}$ ,  $\lambda = 8\eta$ ,  $H_2 = \lceil 4D \ln \frac{4K}{\delta} \rceil$ ,  $H = H_1 + H_2$

**Define:** regularizer  $\psi(\phi) = -\frac{1}{\eta} \sum_{(s,a) \in \tilde{\Gamma}} \ln \phi(s, a)$  where  $\phi(s, a) = \sum_{h=1}^H \phi(s, a, h)$

**Define:** decision set  $\Omega = \{\phi = q + \lambda \mathbf{h} \circ q : q \in \tilde{\Delta}(T)\}$  (with  $\tilde{\Delta}(T)$  defined in Eq. (8))

**Initialization:**  $\phi_1 = q_1 + \lambda \mathbf{h} \circ q_1 = \operatorname{argmin}_{\phi \in \Omega} \psi(\phi)$ .

**for**  $k = 1, \dots, K$  **do**

Execute  $\sigma(\tilde{\pi}_k)$  where  $\tilde{\pi}_k$  is such that  $\tilde{\pi}_k(a|(s, h)) \propto q_k(s, a, h)$ .

Construct cost estimator  $\hat{c}_k \in \mathbb{R}_{\geq 0}^{\tilde{\Gamma}}$  such that  $\hat{c}_k(s, a) = \frac{\tilde{N}_k(s, a) c_k(s, a)}{q_k(s, a)}$  where  $\tilde{N}_k(s, a) = \sum_h \tilde{N}_k(s, a, h)$  and  $q_k(s, a) = \sum_h q_k(s, a, h)$  ( $\tilde{N}_k$  is defined after Lemma 6).

Update  $\phi_{k+1} = q_{k+1} + \lambda \mathbf{h} \circ q_{k+1} = \operatorname{argmin}_{\phi \in \Omega} \sum_{(s,a)} \phi(s, a) \hat{c}_k(s, a) + D_\psi(\phi, \phi_k)$ .

---

taking advantage of the fact that  $c_k(s, a, h)$  is the same for a fixed  $(s, a)$  pair regardless of the value of  $h$ ,<sup>3</sup> we propose to perform OMD with  $\phi(s, a) = \sum_h \phi(s, a, h)$  for all  $(s, a) \in \tilde{\Gamma}$  as the variables, even though the skewed occupancy measure  $\Omega$  is still defined in terms of  $\phi(s, a, h)$  as in Algorithm 3. More specifically, this means that our regularizer is  $\psi(\phi) = -\frac{1}{\eta} \sum_{(s,a)} \ln \phi(s, a)$ , and the cost estimator is  $\hat{c}_k(s, a) = \frac{\tilde{N}_k(s, a) c_k(s, a)}{q_k(s, a)}$  where  $\tilde{N}_k(s, a) = \sum_h \tilde{N}_k(s, a, h)$  and  $q_k(s, a) = \sum_h q_k(s, a, h)$ . This completely avoids the factor  $H$  in the analysis (other than lower order terms).

With the ideas above, we can already show an optimal expected regret bound for an *oblivious* adversary who selects  $c_k$  independent of the learner's randomness. We summarize the algorithm in Algorithm 4 and its guarantee in the following theorem.

**Theorem 9** *If  $T \geq T_\star + 1$ ,  $H_1 \geq 8T_{\max} \ln K$ , and  $K \geq 64SAH^2$ , then Algorithm 4 ensures  $\mathbb{E}[R_K] = \tilde{O}(\sqrt{DT_\star AK})$  for an oblivious adversary.*

Setting  $T = T_\star + 1$  leads to  $\tilde{O}(\sqrt{DT_\star AK})$ , which is optimal in light of the following lower bound theorem (the adversary is indeed oblivious in the lower bound construction).

**Theorem 10** *For any  $D, T_\star, K, S \geq 4$  with  $K \geq ST_\star$  and  $T_\star \geq D + 1$ , there exists an SSP problem instance with  $S$  states and  $A = \mathcal{O}(1)$  actions such that its diameter is  $D + 2$ , the optimal policy has expecting hitting time  $T_\star + 1$ , and the expected regret of any learner after  $K$  episodes is at least  $\Omega(\sqrt{DT_\star AK})$  under the bandit feedback and known transition setting.*

To further obtain a high probability regret bound for general adaptive adversaries (thus also a more general expected regret bound), it is important to analyze the derivation between the optimal policy's estimated total loss  $\sum_k \langle q_{\tilde{\pi}_k}, \hat{c}_k \rangle$  and its expectation  $\sum_k \langle q_{\tilde{\pi}_k}, c_k \rangle$ . Using Freedman's inequality, we need to carefully control the conditional variance  $\mathbb{E}_k[\hat{c}_k^2(s, a)] = \frac{\mathbb{E}_k[\tilde{N}_k^2(s, a) c_k^2(s, a)]}{q_k^2(s, a)}$  for each  $(s, a)$ , which is much more difficult than the aforementioned stability term due to the lack of

---

3. This also explains why injecting the bias to the occupancy space instead of the cost vectors is important here, as mentioned in the end of Section 3, since the latter makes the cost different for different  $h$ .

the extra  $q_k^2(s, a)$  factor. To address this, we first utilize the simpler form of  $\mathbb{E}_k[\tilde{N}_k^2(s, a)]$  in the loop-free setting and bound it by  $\sum_h h q_k(s, a, h)$  (see [Lemma 19](#)). Then, with  $\rho_K(s, a) = \max_k \frac{1}{q_k(s, a)}$  and  $b_k(s, a) = \frac{\sum_h h q_k(s, a, h) c_k(s, a)}{q_k(s, a)}$ , we bound the key term in the derivation  $\sum_k \langle q_{\tilde{\pi}^*}, \hat{c}_k - c_k \rangle$  by

$$\sum_{(s, a)} q_{\tilde{\pi}^*}(s, a) \sqrt{\rho_K(s, a) \sum_{k=1}^K b_k(s, a)} \leq \frac{1}{\eta} \langle q_{\tilde{\pi}^*}, \rho_K \rangle + \eta \sum_{k=1}^K \langle q_{\tilde{\pi}^*}, b_k \rangle,$$

where  $q_{\tilde{\pi}^*}(s, a) = \sum_h q_{\tilde{\pi}^*}(s, a, h)$  and the last step is by AM-GM inequality (see [Lemma 17](#) for details). The last two terms above are then handled by the following two ideas respectively.

**Increasing learning rate** The first term  $\frac{1}{\eta} \langle q_{\tilde{\pi}^*}, \rho_K \rangle$  appears in the work of ([Lee et al., 2020a](#)) already for loop-free MDPs and can be canceled by a negative term introduced by an increasing learning rate schedule. (See the last for loop of [Algorithm 5](#) and [Lemma 16](#).)

**Injecting negative bias to the costs** To handle the second term  $\eta \sum_{k=1}^K \langle q_{\tilde{\pi}^*}, b_k \rangle$ , note again that its counterpart  $\eta \sum_{k=1}^K \langle q_k, b_k \rangle$  is exactly  $\eta \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle$ , a term that can be canceled by the skewed occupancy measure as discussed. Therefore, if we could inject another negative bias term into the cost vectors, that is, replacing  $\hat{c}_k$  with  $\hat{c}_k - \eta b_k$ , then this bias would cancel the term  $\eta \sum_{k=1}^K \langle q_{\tilde{\pi}^*}, b_k \rangle$  while introducing the term  $\eta \sum_{k=1}^K \langle q_k, b_k \rangle$  that could be further canceled by the skewed occupancy measure. However, the issue is that  $b_k$  depends on the unknown true cost  $c_k$ . We address this by using  $\hat{b}_k$  instead which replaces  $c_k$  with  $\hat{c}_k$ , that is,  $\hat{b}_k(s, a) = \frac{\sum_h h q_k(s, a, h) \hat{c}_k(s, a)}{q_k(s, a)}$ .

This leads to yet another derivation term between  $\hat{b}_k$  and  $b_k$  that needs to be controlled in the analysis. Fortunately, this term is of lower order compared to others since it is multiplied by  $\eta$  (see [Lemma 18](#)). Note that at this point we have used both the positive bias from the skewed occupancy measure space and the negative bias from the cost estimators, which we find intriguing.

Combining everything, our final algorithm is summarized in [Algorithm 5](#) (see [Appendix B](#) due to space limit). The following theorem shows that, with the knowledge of  $T_*$  or a suitable upper bound, our algorithm again achieves the optimal regret bound with high probability.

**Theorem 11** *If  $T \geq T_* + 1$ ,  $H_1 \geq 8T_{\max} \ln K$ , and  $K$  is large enough ( $K \gtrsim SAH^2 \ln(\frac{1}{\delta})$ ), then [Algorithm 5](#) ensures  $R_K = \tilde{O}\left(\sqrt{DTSAK \ln(1/\delta)}\right)$  with probability at least  $1 - 6\delta$ .*

## 5. Conclusion

In this paper, we develop matching upper and lower bounds for the stochastic shortest path problem with adversarial costs and unknown transition, significantly improving previous results. Our algorithms are built on top of a variety of techniques that might be of independent interest.

There are two key future directions. The first one is to develop parameter-free and optimal algorithms without the knowledge of  $T_*$ . We only achieve this in the full-information setting for expected regret bounds. Indeed, generalizing our techniques that learn  $T_*$  automatically to obtain a high-probability bound in the full-information setting boils down to getting the same multi-scale expert result with high probability, which is still open unfortunately (see also discussions in ([Chen et al., 2021](#), Section 5)). The difficulty lies in bounding the deviation between the learner's expected loss and the actual loss in terms of the loss of the unknown comparator. On the other hand, it is

also difficult to generalize our technique to obtain an expected bound in the bandit setting (without knowing  $T_*$ ), since this becomes a bandit-of-bandits type of framework and is known to suffer some tuning issues; see for example (Foster et al., 2019, Appendix A.2).

The second future direction is to figure out the minimax regret of the more challenging setting where the transition is unknown. We note that our loop-free reduction is readily to be applied to this case, but due to some technical challenges, it is highly unclear how to avoid having the dependence on  $T_{\max}$  in the regret bounds. A follow-up work by the first two authors (Chen and Luo, 2021) makes some progress in this direction, but the minimax regret remains unknown in this case.

## Acknowledgments

The authors thank Tiancheng Jin for many helpful discussions. This work is supported by NSF Award IIS-1943607 and a Google Faculty Research Award.

## References

- Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E Schapire. Corraling a band of bandit algorithms. In *Conference on Learning Theory*, 2017.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, pages 263–272, 2017.
- Dimitri P Bertsekas and John N Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
- Dimitri P Bertsekas and Huizhen Yu. Stochastic shortest path problems under weak conditions. *Lab. for Information and Decision Systems Report LIDS-P-2909, MIT*, 2013.
- Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *International Conference on Artificial Intelligence and Statistics*, 2011.
- Sébastien Bubeck, Nikhil R Devanur, Zhiyi Huang, and Rad Niazadeh. Online auctions and multi-scale online learning. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 497–514, 2017.
- Sébastien Bubeck, Yuanzhi Li, Haipeng Luo, and Chen-Yu Wei. Improved path-length regret bounds for bandits. In *Conference On Learning Theory*, 2019.
- Liyu Chen and Haipeng Luo. Finding the stochastic shortest path with low regret: The adversarial cost and unknown transition case. In *International Conference on Machine Learning*, 2021.
- Liyu Chen, Haipeng Luo, and Chen-Yu Wei. Impossible tuning made possible: A new expert algorithm and its applications. In *Conference On Learning Theory*, 2021.
- Alon Cohen, Haim Kaplan, Yishay Mansour, and Aviv Rosenberg. Near-optimal regret bounds for stochastic shortest path. In *Proceedings of the 37th International Conference on Machine Learning*, pages 8210–8219, 2020.

- Ashok Cutkosky and Francesco Orabona. Black-box reductions for parameter-free online learning in banach spaces. In *Conference on Learning Theory (COLT)*, pages 1493–1529, 2018.
- Dylan J Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems*, 2016.
- Dylan J Foster, Satyen Kale, Mehryar Mohri, and Karthik Sridharan. Parameter-free online learning via model selection. In *Advances in Neural Information Processing Systems*, pages 6020–6030, 2017.
- Dylan J. Foster, Akshay Krishnamurthy, and Haipeng Luo. Model selection for contextual bandits. In *Advances in Neural Information Processing Systems*, 2019.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Sébastien Gerchinovitz and Tor Lattimore. Refined lower bounds for adversarial bandits. In *Advances in Neural Information Processing Systems*, pages 1198–1206, 2016.
- Chi Jin, Zeyuan Allen-Zhu, Sébastien Bubeck, and Michael I Jordan. Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873, 2018.
- Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu. Learning adversarial Markov decision processes with bandit feedback and unknown transition. In *Proceedings of the 37th International Conference on Machine Learning*, pages 4860–4869, 2020.
- Wojciech Kotłowski and Gergely Neu. Bandit principal component analysis. In *Conference On Learning Theory*, 2019.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, and Mengxiao Zhang. Bias no more: high-probability data-dependent regret bounds for adversarial bandits and MDPs. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Chung-Wei Lee, Haipeng Luo, and Mengxiao Zhang. A closer look at small-loss bounds for bandits with graph feedback. In *Conference on Learning Theory*, 2020b.
- Haipeng Luo, Chen-Yu Wei, and Kai Zheng. Efficient online portfolio with logarithmic regret. In *Advances in Neural Information Processing Systems*, 2018.
- Gergely Neu, Andras Gyorgy, and Csaba Szepesvári. The adversarial stochastic shortest path problem with unknown transition probabilities. In *Artificial Intelligence and Statistics*, pages 805–813, 2012.
- Aviv Rosenberg and Yishay Mansour. Online convex optimization in adversarial Markov decision processes. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5478–5486, 2019.

- Aviv Rosenberg and Yishay Mansour. Stochastic shortest path with adversarially changing costs. *arXiv preprint arXiv:2006.11561*, 2020.
- Lior Shani, Yonathan Efroni, Aviv Rosenberg, and Shie Mannor. Optimistic policy optimization with bandit feedback. In *Proceedings of the 37th International Conference on Machine Learning*, pages 8604–8613, 2020.
- Jacob Steinhardt and Percy Liang. Adaptivity and optimism: An improved exponentiated gradient algorithm. In *International Conference on Machine Learning*, pages 1593–1601, 2014.
- Jean Tarbouriech, Evrard Garcelon, Michal Valko, Matteo Pirotta, and Alessandro Lazaric. No-regret exploration in goal-oriented reinforcement learning. In *International Conference on Machine Learning*, pages 9428–9437. PMLR, 2020.
- Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. In *Conference On Learning Theory*, pages 1263–1291, 2018.
- Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *Proceedings of the 36th International Conference on Machine Learning*, pages 7304–7312, 2019.
- Alexander Zimin and Gergely Neu. Online learning in episodic Markovian decision processes by relative entropy policy search. In *Advances in neural information processing systems*, pages 1583–1591, 2013.



## Appendix A. Omitted details for Section 3

In this section, we provide all proofs for Section 3.

### A.1. Proof of Lemma 1

**Proof** By standard OMD analysis (see for example Eq. (12) of (Rosenberg and Mansour, 2020)), for any  $q \in \Delta(T)$  we have:

$$\sum_{k=1}^K \langle q_k - q, c_k \rangle \leq D_\psi(q, q_1) + \sum_{k=1}^K \langle q_k - q'_{k+1}, c_k \rangle, \quad (3)$$

where  $q'_{k+1} = \operatorname{argmin}_{q \in \mathbb{R}^\Gamma} \langle q, c_k \rangle + D_\psi(q, q_k)$ , or equivalently, with the particular choice of the regularizer,  $q'_{k+1}(s, a) = q_k(s, a)e^{-\eta c_k(s, a)}$ . Applying the inequality  $1 - e^{-x} \leq x$ , we obtain

$$\sum_{k=1}^K \langle q_k - q'_{k+1}, c_k \rangle \leq \eta \sum_{k=1}^K \sum_{(s, a)} q_k(s, a) c_k^2(s, a) \leq \eta \sum_{k=1}^K \langle q_k, c_k \rangle.$$

Substituting this back into Eq. (3), choosing  $q = q_{\pi^*}$  (recall the condition  $q_{\pi^*} \in \Delta(T)$  of the lemma), and rearranging, we arrive at

$$\begin{aligned} \sum_{k=1}^K \langle q_k - q_{\pi^*}, c_k \rangle &\leq \frac{1}{1 - \eta} \left( D_\psi(q_{\pi^*}, q_1) + \eta \sum_{k=1}^K \langle q_{\pi^*}, c_k \rangle \right) \\ &\leq 2D_\psi(q_{\pi^*}, q_1) + 2\eta \sum_{k=1}^K \langle q_{\pi^*}, c_k \rangle. \end{aligned} \quad (4)$$

It remains to bound the last two terms. For the first one, since  $q_1$  minimizes  $\psi$  over  $\Delta(T)$ , we have  $\langle \nabla \psi(q_1), q_{\pi^*} - q_1 \rangle \geq 0$ , and thus

$$\begin{aligned} D_\psi(q_{\pi^*}, q_1) &\leq \psi(q_{\pi^*}) - \psi(q_1) = \frac{1}{\eta} \sum_{(s, a)} q_{\pi^*}(s, a) \ln q_{\pi^*}(s, a) - \frac{1}{\eta} \sum_{(s, a)} q_1(s, a) \ln q_1(s, a) \\ &\leq \frac{1}{\eta} \sum_{(s, a)} q_{\pi^*}(s, a) \ln T - \frac{T}{\eta} \sum_{(s, a)} \frac{q_1(s, a)}{T} \ln \frac{q_1(s, a)}{T} \\ &\leq \frac{T \ln(T)}{\eta} + \frac{T \ln(SA)}{\eta} = \frac{T \ln(SAT)}{\eta}. \end{aligned}$$

For the second one, we use the fact  $\sum_{k=1}^K \langle q_{\pi^*}, c_k \rangle \leq \sum_{k=1}^K \langle q_{\pi^f}, c_k \rangle \leq DK$ . Put together, this implies

$$\sum_{k=1}^K \langle q_k - q_{\pi^*}, c_k \rangle \leq \frac{2T \ln(SAT)}{\eta} + 2\eta DK.$$

With the optimal  $\eta = \min \left\{ \frac{1}{2}, \sqrt{\frac{T \ln(SAT)}{DK}} \right\}$ , we have thus shown

$$\mathbb{E}[R_k] = \mathbb{E} \left[ \sum_{k=1}^K \langle q_k - q_{\pi^*}, c_k \rangle \right] = \mathcal{O} \left( \sqrt{DTK \ln(SAT)} + T \ln(SAT) \right) = \tilde{\mathcal{O}} \left( \sqrt{DTK} \right),$$

completing the proof. ■

## A.2. Proof of Theorem 2

**Proof** First, note that the value of  $j_0$  is such that the smallest parameter  $b(1)$  is larger than  $T^{\pi^f}(s_0)$  and thus  $\Delta(b(j))$  is non-empty for all  $j \in [N]$ , making all  $N$  instances of Algorithm 1 well-defined. Next, let  $j^*$  be the index of the instance with size parameter closest to the unknown parameter  $T_*$ , that is,  $\frac{b(j^*)}{2} \leq T_* \leq b(j^*)$ . Such  $j^*$  must exist since  $b(N) \geq K$  and we only need to consider the case  $T_* \leq K$  (otherwise the claimed regret bound is vacuous). Now we decompose the regret as two parts, the regret of the meta algorithm to instance  $j^*$ , and the regret of instance  $j^*$  to the best policy:

$$\begin{aligned} \mathbb{E}[R_K] &= \mathbb{E} \left[ \sum_{k=1}^K \sum_{j=1}^N p_k(j) \langle q_k^j, c_k \rangle - \sum_{k=1}^K \langle q_{\pi^*}, c_k \rangle \right] \\ &= \mathbb{E} \left[ \sum_{k=1}^K \sum_{j=1}^N p_k(j) \langle q_k^j, c_k \rangle - \langle q_k^{j^*}, c_k \rangle \right] + \mathbb{E} \left[ \sum_{k=1}^K \langle q_k^{j^*} - q_{\pi^*}, c_k \rangle \right] \\ &= \mathbb{E} \left[ \sum_{k=1}^K \langle p_k - e_{j^*}, \ell_k \rangle \right] + \mathbb{E} \left[ \sum_{k=1}^K \langle q_k^{j^*} - q_{\pi^*}, c_k \rangle \right], \end{aligned}$$

where  $e_{j^*} \in \Omega$  is the basis vector with the  $j^*$ -th coordinate being 1. By the regret guarantee of Algorithm 1 (Lemma 1), the second term above is bounded by  $\tilde{O}(\sqrt{Db(j^*)K}) = \tilde{O}(\sqrt{DT_*K})$ , which also means

$$\begin{aligned} \mathbb{E} \left[ \sum_{k=1}^K \ell_k(j^*) \right] &\leq \mathbb{E} \left[ \sum_{k=1}^K \langle q_{\pi^*}, c_k \rangle \right] + \tilde{O}(\sqrt{DT_*K}) \\ &\leq \mathbb{E} \left[ \sum_{k=1}^K \langle q_{\pi^f}, c_k \rangle \right] + \tilde{O}(\sqrt{DT_*K}) \\ &\leq DK + \tilde{O}(\sqrt{DT_*K}). \end{aligned}$$

Using Lemma 12, the first term is bounded as

$$\begin{aligned} \mathbb{E} \left[ \sum_{k=1}^K \langle p_k - e_{j^*}, \ell_k \rangle \right] &= \tilde{O} \left( \frac{1}{\eta_{j^*}} + \eta_{j^*} b(j^*) \mathbb{E} \left[ \sum_{k=1}^K \ell_k(j^*) \right] \right) \\ &= \tilde{O} \left( \frac{1}{\eta_{j^*}} + \eta_{j^*} DT_* K + \eta_{j^*} T_* \sqrt{DT_* K} \right) \\ &= \tilde{O} \left( \frac{1}{\eta_{j^*}} + \eta_{j^*} DT_* K \right). \end{aligned}$$

Finally plugging in the definition of  $\eta_{j^*}$  finishes the proof. ■

The lemma below is an improved guarantee for the multi-scale expert problem, which might be of independent interest.

**Lemma 12** For any  $j^* \in [N]$ , [Algorithm 2](#) ensures

$$\sum_{k=1}^K \langle p_k - e_{j^*}, \ell_k \rangle = \frac{2 + \ln \left( N \sqrt{\frac{b(j^*)}{b(1)}} \right)}{\eta_{j^*}} + 4\eta_{j^*} b(j^*) \sum_{k=1}^K \ell_k(j^*).$$

**Proof** Similar to [Eq. \(3\)](#), by standard OMD analysis (see also [\(Bubeck et al., 2017, Lemma 6\)](#)) we have:

$$\sum_{k=1}^K \langle p_k - e_{j^*}, \ell_k + a_k \rangle \leq D_\psi(e_{j^*}, p_1) + \sum_{k=1}^K \langle p_k - p'_{k+1}, \ell_k + a_k \rangle \quad (5)$$

where  $p'_{k+1}(j) = p_k(j)e^{-\eta_j(\ell_k(j) + a_k(j))}$ . Using the inequality  $1 - e^{-x} \leq x$  and the fact  $a_k(j) \leq 4\eta_j b(j)\ell_k(j) \leq \ell_k(j)$  (since  $\ell_k(j) \leq \sum_{(s,a)} q_k^j(s, a) \leq b(j)$  and  $\eta_j \leq \frac{1}{4b(j)}$ ), we obtain

$$\langle p_k - p'_{k+1}, \ell_k + a_k \rangle \leq \sum_{j=1}^N \eta_j p_k(j) (\ell_k(j) + a_k(j))^2 \leq 4 \sum_{j=1}^N \eta_j p_k(j) \ell_k^2(j) = \langle p_k, a_k \rangle.$$

Plugging this back into [Eq. \(5\)](#) and rearranging leads to

$$\begin{aligned} \sum_{k=1}^K \langle p_k - e_{j^*}, \ell_k \rangle &\leq D_\psi(e_{j^*}, p_1) + \sum_{k=1}^K a_k(j^*) = D_\psi(e_{j^*}, p_1) + 4\eta_{j^*} \sum_{k=1}^K \ell_k^2(j^*) \\ &\leq D_\psi(e_{j^*}, p_1) + 4\eta_{j^*} b(j^*) \sum_{k=1}^K \ell_k(j^*). \end{aligned}$$

It remains to bound  $D_\psi(e_{j^*}, p_1)$ , which by definition is

$$\sum_{j=1}^N \frac{1}{\eta_j} \left( e_{j^*}(j) \ln \frac{e_{j^*}(j)}{p_1(j)} - e_{j^*}(j) + p_1(j) \right) \leq \frac{1}{\eta_{j^*}} \ln \frac{1}{p_1(j^*)} + \sum_{j=1}^N \frac{p_1(j)}{\eta_j}.$$

Using the definition of  $p_1$ , when  $j^* \neq 1$  we have

$$\frac{1}{\eta_{j^*}} \ln \frac{1}{p_1(j^*)} = \frac{1}{\eta_{j^*}} \ln \left( \frac{N\eta_1}{\eta_{j^*}} \right) = \frac{\ln \left( N \sqrt{\frac{b(j^*)}{b(1)}} \right)}{\eta_{j^*}};$$

when  $j^* = 1$ , the same holds as an upper bound since  $p_1(1) \geq 1/N$ . Finally, the second term can be bounded as

$$\sum_{j=1}^N \frac{p_1(j)}{\eta_j} = \frac{p_1(1)}{\eta_1} + \sum_{j \neq 1} \frac{1}{N\eta_1} \leq \frac{2}{\eta_1} \leq \frac{2}{\eta_{j^*}},$$

which finishes the proof. ■

### A.3. Proof of Theorem 3

**Proof** By Yao's minimax principle, in order to obtain a regret lower bound, it suffices to show that there exists a distribution of SSP instances that forces any deterministic learner to suffer a regret bound of  $\Omega(\sqrt{DT_*K})$  in expectation. Below we describe such a distribution (the MDP is fixed but the costs are stochastic).

- The state space is  $\mathcal{S} = \{s_0, s_1, \dots, s_N, f\}$  for any  $N \geq 2$ .
- At state  $s_0$ , there are  $N$  available actions  $a_1, \dots, a_N$ ; at each state of  $s_1, \dots, s_N$ , there are two available actions  $a_g$  and  $a_f$ ; and at state  $f$ , there is only one action  $a_g$ .
- At state  $s_0$ , taking action  $a_j$  transits to state  $s_j$  deterministically for all  $j \in [N]$ . At any state  $s_j$  ( $j \in [N]$ ), taking action  $a_f$  transits to state  $f$  deterministically, while taking action  $a_g$  transits to the goal state  $g$  with probability  $1/T_*$  and stays at the same state with probability  $1 - 1/T_*$ . Finally, at state  $f$ , taking action  $a_g$  transits to the goal state  $g$  with probability  $1/D$  and stays with probability  $1 - 1/D$ .
- The cost at state  $s_0$  is always zero, that is,  $c_k(s_0, a) = 0$  for all  $k$  and  $a$ ; the cost of action  $a_f$  is also always zero, that is,  $c_k(s, a_f) = 0$  for all  $k$  and  $s \in \{s_1, \dots, s_N\}$ ; the cost at state  $f$  is always one, that is,  $c_k(f, a_g) = 1$  for all  $k$ ; finally, the cost of taking action  $a_g$  at state  $s \in \{s_1, \dots, s_N\}$  is generated stochastically as follows: first, a good state  $j^* \in [N]$  is sampled uniformly at random ahead of time and then fixed throughout the  $K$  episodes; then, in each episode  $k$ ,  $c_k(s, a_g)$  is an independent sample of Bernoulli( $\frac{D}{2T_*}$ ) if  $s = s_{j^*}$ , and an independent sample of Bernoulli( $\frac{D}{2T_*} + \epsilon$ ) if  $s \neq s_{j^*}$ , for some  $\epsilon \leq \frac{D}{2T_*}$  to be specified later.

It is clear that in all these SSP instances, the diameter is  $D + 2$  (since one can reach the goal state via the fast state  $f$  within at most  $D + 2$  steps in expectation), and the hitting time of the optimal policy is indeed  $T_* + 1$  (in fact, the hitting time of any stationary deterministic policy is either  $T_* + 1$  or  $D + 2 \leq T_* + 1$ ). It remains to argue  $\mathbb{E}[R_K] = \Omega(\sqrt{DT_*K})$  for any deterministic learner, where the expectation is over the randomness of the costs. To do so, let  $\mathbb{E}_j$  denote the conditional expectation given that the good state  $j^*$  is  $j$ . Then we have

$$\begin{aligned} \mathbb{E}[R_K] &= \frac{1}{N} \sum_{j=1}^N \left( \mathbb{E}_j \left[ \sum_{k=1}^K \sum_{i=1}^{I_k} c_k(s_k^i, a_k^i) - \min_{\pi \in \Pi_{\text{proper}}} \sum_{k=1}^K J_k^\pi(s_0) \right] \right) \\ &\geq \frac{1}{N} \sum_{j=1}^N \left( \mathbb{E}_j \left[ \sum_{k=1}^K \sum_{i=1}^{I_k} c_k(s_k^i, a_k^i) - \sum_{k=1}^K J_k^{\pi_j}(s_0) \right] \right), \end{aligned}$$

where  $\pi_j$  is the policy that picks action  $a_j$  at state  $s_0$  and  $a_g$  at state  $s_j$  (other states are irrelevant). Note that it takes  $T_*$  steps in expectation for  $\pi_j$  to reach  $g$  from  $s_j$  and each step incur expected cost  $\frac{D}{2T_*}$ , which means  $\mathbb{E}_j[J_k^{\pi_j}(s_0)] = \frac{D}{2T_*} \times T_* = \frac{D}{2}$ . On the other hand, the learner is always better off not visiting  $f$  at all, since starting from state  $f$ , the expected cost before reaching  $g$  is  $D$ , while the expected cost of reaching the goal state via any other states is at most  $(\frac{D}{2T_*} + \epsilon) \times T_* \leq D$ . Therefore, depending on whether the learner selects the good action  $a_{j^*}$  or not at the first step, we further lower bound the expected regret as

$$\mathbb{E}[R_K] \geq \frac{1}{N} \sum_{j=1}^N \sum_{k=1}^K \mathbb{E}_j \left[ \frac{D}{2} + T_* \epsilon \mathbb{1}\{a_k^1 \neq a_j\} - \frac{D}{2} \right] = T_* K \epsilon - \frac{T_* \epsilon}{N} \sum_{j=1}^N \mathbb{E}_j[K_j],$$

where  $K_j = \sum_{k=1}^K \mathbb{I}\{a_k^1 = a_j\}$ .

It thus suffices to upper bound  $\sum_{j=1}^N \mathbb{E}_j[K_j]$ . To do so, consider a reference environment without a good state, that is,  $c_k(s, a_g)$  is an independent sample of  $\text{Bernoulli}(\frac{D}{2T_\star} + \epsilon)$  for all  $k$  and all  $s \in \{s_1, \dots, s_N\}$ . Denote by  $\mathbb{E}_0$  the expectation with respect to this reference environment, and by  $P_0$  the distribution of the learner's observation in this environment ( $P_j$  is defined similarly). Then with the fact  $K_j \leq K$  and Pinsker's inequality, we have

$$\mathbb{E}_j[K_j] - \mathbb{E}_0[K_j] \leq K \|P_j - P_0\|_1 \leq K \sqrt{2\text{KL}(P_0, P_j)}.$$

By the divergence decomposition lemma (see e.g. (Lattimore and Szepesvári, 2020, Lemma 15.1)) and the nature of the full-information setting, we further have

$$\begin{aligned} \text{KL}(P_0, P_j) &= \sum_{j'=1}^N \mathbb{E}_0[K_{j'}] \times \text{KL}\left(\text{Bernoulli}\left(\frac{D}{2T_\star} + \epsilon\right), \text{Bernoulli}\left(\frac{D}{2T_\star}\right)\right) \\ &= K \times \text{KL}\left(\text{Bernoulli}\left(\frac{D}{2T_\star} + \epsilon\right), \text{Bernoulli}\left(\frac{D}{2T_\star}\right)\right) \\ &\leq \frac{K\epsilon^2}{\alpha(1-\alpha)}, \end{aligned}$$

where the last step is by (Gerchinovitz and Lattimore, 2016, Lemma 6) with  $\alpha = \frac{D}{2T_\star}$ . Therefore, we have

$$\sum_{j=1}^N \mathbb{E}_j[K_j] \leq \sum_{j=1}^N \mathbb{E}_0[K_j] + NK \sqrt{\frac{2K\epsilon^2}{\alpha(1-\alpha)}} = K + NK \sqrt{\frac{2K\epsilon^2}{\alpha(1-\alpha)}}.$$

This is enough to show the claimed lower bound:

$$\begin{aligned} \mathbb{E}[R_K] &\geq T_\star \epsilon K - \frac{T_\star \epsilon}{N} \sum_{j=1}^N \mathbb{E}_j[K_j] \\ &\geq T_\star \epsilon K - \frac{T_\star \epsilon}{N} \left[ K + NK \sqrt{\frac{2K\epsilon^2}{\alpha(1-\alpha)}} \right] \\ &= T_\star \epsilon K \left[ 1 - \frac{1}{N} - \epsilon \sqrt{\frac{2K}{\alpha(1-\alpha)}} \right] \\ &\geq T_\star \epsilon K \left[ \frac{1}{2} - \epsilon \sqrt{\frac{2K}{\alpha(1-\alpha)}} \right] \\ &= \frac{T_\star K}{16} \sqrt{\frac{\alpha(1-\alpha)}{2K}} = \Omega\left(\sqrt{DT_\star K}\right), \end{aligned}$$

where in the last line we choose  $\epsilon = \frac{1}{4} \sqrt{\frac{\alpha(1-\alpha)}{2K}} \leq \frac{1}{8} \sqrt{\frac{D}{T_\star K}} \leq \frac{D}{2T_\star}$  to maximize the lower bound.  $\blacksquare$

#### A.4. Proof of Lemma 4

**Proof** With the inequality  $(\sum_{i=1}^I a_i)^2 \leq 2 \sum_i a_i (\sum_{i'=i}^I a_{i'})$ , we proceed as

$$\begin{aligned}
 & \mathbb{E} \left[ \left( \sum_{(s,a)} N_k(s,a) c_k(s,a) \right)^2 \right] \\
 &= \mathbb{E} \left[ \left( \sum_{i=1}^{I_k} \sum_{(s,a)} \mathbb{I}\{s_k^i = s, a_k^i = a\} c_k(s,a) \right)^2 \right] \\
 &\leq 2 \mathbb{E} \left[ \sum_{i=1}^{I_k} \sum_{(s,a)} \mathbb{I}\{s_k^i = s, a_k^i = a\} \left( \sum_{i'=i}^{I_k} \sum_{(s',a')} \mathbb{I}\{s_k^{i'} = s', a_k^{i'} = a'\} c_k(s',a') \right) \right] \\
 &= 2 \mathbb{E} \left[ \sum_{i=1}^{I_k} \sum_{s \in \mathcal{S}} \mathbb{I}\{s_k^i = s\} \mathbb{E} \left[ \sum_{i'=i}^{I_k} \sum_{(s',a')} \mathbb{I}\{s_k^{i'} = s', a_k^{i'} = a'\} c_k(s',a') \middle| s_k^i = s \right] \right] \\
 &= 2 \mathbb{E} \left[ \sum_{i=1}^{I_k} \sum_{s \in \mathcal{S}} \mathbb{I}\{s_k^i = s\} J_k^\pi(s) \right] \\
 &= 2 \sum_{s \in \mathcal{S}} q_\pi(s) J_k^\pi(s) = 2 \langle q_\pi, J_k^\pi \rangle,
 \end{aligned}$$

completing the proof. ■

#### A.5. Proof of Lemma 6

**Proof** Applying Lemma 4 (to the loop-free instance), we have

$$\mathbb{E} \left[ \left\langle \tilde{N}_k, c_k \right\rangle^2 \right] \leq 2 \sum_{\tilde{s} \in \tilde{\mathcal{S}}} q_{\tilde{\pi}}(\tilde{s}) J_k^{\tilde{\pi}}(\tilde{s}) = 2 \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}}(s,h) J_k^{\tilde{\pi}}(s,h).$$

Denote  $q_{\tilde{\pi},(s,h)}$  as the occupancy measure of policy  $\tilde{\pi}$  with initial state  $(s,h)$ , so that

$$J_k^{\tilde{\pi}}(s,h) = \sum_{(s',a') \in \tilde{\Gamma}} \sum_{h' \geq h} q_{\tilde{\pi},(s,h)}(s',a',h') c_k(s',a',h').$$

Then we continue with the following equalities:

$$\begin{aligned}
 & \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}}(s,h) J_k^{\tilde{\pi}}(s,h) \\
 &= \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}}(s,h) \sum_{(s',a') \in \tilde{\Gamma}} \sum_{h' \geq h} q_{\tilde{\pi},(s,h)}(s',a',h') c_k(s',a',h')
 \end{aligned}$$

$$\begin{aligned}
 &= \sum_{h=1}^H \sum_{(s', a') \in \tilde{\Gamma}} \sum_{h' \geq h} \left( \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}}(s, h) q_{\tilde{\pi}, (s, h)}(s', a', h') \right) c_k(s', a', h') \\
 &= \sum_{h=1}^H \sum_{(s', a') \in \tilde{\Gamma}} \sum_{h' \geq h} q_{\tilde{\pi}}(s', a', h') c_k(s', a', h') \\
 &= \sum_{h=1}^H \sum_{(s, a) \in \tilde{\Gamma}} h \cdot q_{\tilde{\pi}}(s, a, h) c_k(s, a, h) = \langle q_{\tilde{\pi}}, \mathbf{h} \circ c_k \rangle. \tag{6}
 \end{aligned}$$

where in the third line we use the equality  $\sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}}(s, h) q_{\tilde{\pi}, (s, h)}(s', a', h') = q_{\tilde{\pi}}(s', a', h')$  by definition (since both sides are the probability of visiting  $(s', a', h')$ ). This completes the proof. ■

### A.6. Proof of Lemma 7

**Proof** We first prove the second statement  $\mathbb{E}[R_K] \leq \mathbb{E}[\text{REG}] + \tilde{\mathcal{O}}(1)$ . Since the fast policy reaches the goal state within  $D$  steps in expectation starting from any state, by the definition of  $\sigma(\tilde{\pi})$  and  $\tilde{M}$ , we have  $J_k^{\sigma(\tilde{\pi})}(s_0) \leq J_k^{\tilde{\pi}}(\tilde{s}_0)$  for any  $\tilde{\pi}$ , that is, the expected cost of executing  $\sigma(\tilde{\pi})$  in  $M$  is not larger than that of executing  $\tilde{\pi}$  in  $\tilde{M}$ . On the other hand, since the probability of not reaching the goal state within  $H_1$  steps when executing  $\pi^*$  is at most:  $2e^{-\frac{H_1}{4T_{\max}}} \leq \frac{2}{K^2}$  by Lemma 13 and the choice of  $H_1$ , the expected cost of  $\pi^*$  in  $M$  and the expected cost of  $\tilde{\pi}^*$  in  $\tilde{M}$  is very similar:

$$J_k^{\tilde{\pi}^*}(\tilde{s}_0) \leq J_k^{\pi^*}(s_0) + \frac{2H_2}{K^2} = J_k^{\pi^*}(s_0) + \tilde{\mathcal{O}}\left(\frac{1}{K}\right). \tag{7}$$

This proves the second statement:

$$\begin{aligned}
 \mathbb{E}[R_K] &= \mathbb{E} \left[ \sum_{k=1}^K J_k^{\sigma(\tilde{\pi}_k)}(s_0) - J_k^{\pi^*}(s_0) \right] \\
 &\leq \mathbb{E} \left[ \sum_{k=1}^K J_k^{\tilde{\pi}_k}(\tilde{s}_0) - J_k^{\tilde{\pi}^*}(\tilde{s}_0) \right] + \tilde{\mathcal{O}}(1) = \mathbb{E}[\text{REG}] + \tilde{\mathcal{O}}(1).
 \end{aligned}$$

To prove the first statement, we apply Lemma 13 again to show that for each episode  $k$ , the probability of the learner not reaching  $g$  within  $H$  steps is at most  $2e^{-\frac{H_2}{4D}} = \frac{\delta}{2K}$ . With a union bound, this means, with probability at least  $1 - \frac{\delta}{2}$ , the learner reaches the goal within  $H$  steps for all episodes and thus her actual loss in  $M$  is not larger than that in  $\tilde{M}$ :  $\sum_{k=1}^K \langle N_k, c_k \rangle \leq \sum_{k=1}^K \langle \tilde{N}_k, c_k \rangle$ . Together with Eq. (7), this shows

$$R_K \leq \sum_{k=1}^K \langle \tilde{N}_k - q_{\tilde{\pi}^*}, c_k \rangle + \tilde{\mathcal{O}}(1).$$

It thus remains to bound the deviation  $\sum_{k=1}^K \langle \tilde{N}_k - q_k, c_k \rangle$ , which is the sum of a martingale difference sequence. We apply Freedman's inequality Lemma 21 directly: the variable  $\langle \tilde{N}_k, c_k \rangle$

is bounded by  $H$  always, and its conditional variance is bounded by  $2 \langle q_k, \mathbf{h} \circ c_k \rangle$  as shown in [Lemma 6](#), which means for any  $\lambda \in (0, 2/H]$ ,

$$\sum_{k=1}^K \left\langle \tilde{N}_k - q_k, c_k \right\rangle \leq \lambda \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle + \frac{2 \ln(2/\delta)}{\lambda}$$

holds with probability at least  $1 - \frac{\delta}{2}$ . Applying another union bound finishes the proof.  $\blacksquare$

**Lemma 13** (*Rosenberg and Mansour, 2020, Lemma E.1*) *Let  $\pi$  be a policy with expected hitting time at most  $\tau$  starting from any state. Then, the probability that  $\pi$  takes more than  $m$  steps to reach the goal state is at most  $2e^{-\frac{m}{4\tau}}$ .*

### A.7. Proof of [Theorem 8](#)

For completeness, we first spell out the definition of  $\tilde{\Delta}(T)$ , which is the exact counterpart of  $\Delta(T)$  defined in [Eq. \(2\)](#) for  $\tilde{M}$  (the first equality below), but can be simplified using the special structure of  $\tilde{P}$  (the second equality below).

$$\begin{aligned} \tilde{\Delta}(T) &= \left\{ q \in [0, 1]^{\tilde{\Gamma} \times [H]} : \sum_{h=1}^H \sum_{(s,a) \in \tilde{\Gamma}} q(s, a, h) \leq T, \right. \\ &\quad \left. \sum_{a \in \tilde{\mathcal{A}}(s,h)} q(s, a, h) - \sum_{h'=1}^H \sum_{(s',a') \in \tilde{\Gamma}} \tilde{P}((s, h) | (s', h'), a') q(s', a', h') = \mathbb{I}\{(s, h) = \tilde{s}_0\}, \forall (s, h) \in \tilde{\mathcal{S}} \right\} \\ &= \left\{ q \in [0, 1]^{\tilde{\Gamma} \times [H]} : \sum_{h=1}^H \sum_{(s,a) \in \tilde{\Gamma}} q(s, a, h) \leq T, \quad \sum_{a \in \mathcal{A}_{s_0}} q(s_0, a, 1) = 1, \right. \\ &\quad q(s, a, 1) = 0, \forall s \neq s_0 \text{ and } a \in \mathcal{A}_s, \quad q(s, a, h) = 0, \forall (s, a) \in \Gamma \text{ and } h > H_1, \\ &\quad \left. q(s_f, a_f, h) = \mathbb{I}\{h > H_1\} \sum_{(s',a') \in \Gamma} q(s', a', H_1), \right. \\ &\quad \left. \sum_{a \in \mathcal{A}_s} q(s, a, h) = \sum_{(s',a') \in \Gamma} P(s|s', a') q(s', a', h-1), \forall s \in \mathcal{S} \text{ and } 1 < h \leq H_1. \right\} \quad (8) \end{aligned}$$

Note that  $q_{\tilde{\pi}^*}$  belongs to  $\tilde{\Delta}(T_\star + 1)$  as shown in the following lemma.

**Lemma 14** *The policy  $\tilde{\pi}^*$  satisfies  $T^{\tilde{\pi}^*}(\tilde{s}_0) = \sum_{h=1}^H \sum_{(s,a) \in \tilde{\Gamma}} q_{\tilde{\pi}^*}(s, a, h) \leq T_\star + 1$  and thus  $q_{\tilde{\pi}^*} \in \tilde{\Delta}(T_\star + 1)$ .*

**Proof** This is a direct application of the fact  $T^{\pi^*}(s_0) = T_\star$  and [Lemma 13](#): the probability of not reaching the goal state within  $H_1$  steps when executing  $\pi^*$  is at most:  $2e^{-\frac{H_1}{4T_{\max}}} \leq \frac{2}{K^2}$ . Therefore,  $T^{\tilde{\pi}^*}(\tilde{s}_0) \leq T^{\pi^*}(s_0) + \frac{2H_2}{K^2} \leq T_\star + 1$ , finishing the proof.  $\blacksquare$

We also need the following lemma.



**Lemma 15** *The policy  $\tilde{\pi}^*$  satisfies  $\sum_{k=1}^K \langle q_{\tilde{\pi}^*}, \mathbf{h} \circ c_k \rangle = \mathcal{O}(DT_*K)$ .*

**Proof** We proceed as follows:

$$\begin{aligned}
 \sum_{k=1}^K \langle q_{\tilde{\pi}^*}, \mathbf{h} \circ c_k \rangle &= \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}^*}(s, h) J_k^{\tilde{\pi}^*}(s, h) \\
 &= \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}^*}(s, h) \sum_{k=1}^K J_k^{\tilde{\pi}^*}(s, h) \\
 &\leq \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}^*}(s, h) \sum_{k=1}^K J_k^{\tilde{\pi}^*}(s, 1) \\
 &\leq \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}^*}(s, h) \left( \tilde{\mathcal{O}}(1) + \sum_{k=1}^K J_k^{\pi^*}(s) \right) \\
 &\leq \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}^*}(s, h) \left( \tilde{\mathcal{O}}(1) + \sum_{k=1}^K J_k^{\pi^f}(s) \right) \\
 &\leq \left( \tilde{\mathcal{O}}(1) + DK \right) \sum_{h=1}^H \sum_{s \in \mathcal{S} \cup \{s_f\}} q_{\tilde{\pi}^*}(s, h) \\
 &\leq \tilde{\mathcal{O}}(DT_*K),
 \end{aligned}$$

where the first line is by [Eq. \(6\)](#), the fourth line is by the same reasoning of [Eq. \(7\)](#), and the last line is by [Lemma 14](#).  $\blacksquare$

We are now ready to prove [Theorem 8](#).

**Proof** Define  $\phi^* = q_{\tilde{\pi}^*} + \lambda \mathbf{h} \circ q_{\tilde{\pi}^*}$ . which belongs to the set  $\Omega$  by [Lemma 14](#) and the condition  $T \geq T_* + 1$ . By the exact same reasoning of [Eq. \(4\)](#) in the proof of [Lemma 1](#), OMD ensures

$$\sum_{k=1}^K \langle \phi_k - \phi^*, c_k \rangle \leq 2D_\psi(\phi^*, \phi_1) + 2\eta \sum_{k=1}^K \langle \phi^*, c_k \rangle.$$

The last two terms can also be bounded in a similar way as in the proof of [Lemma 1](#): for the first term, since  $\phi_1$  minimizes  $\psi$  over  $\Omega$ , we have  $\langle \nabla \psi(\phi_1), \phi^* - \phi_1 \rangle \geq 0$ , and thus with the fact  $\sum_{h=1}^H \sum_{(s,a)} \phi(s, a, h) \leq T + \lambda HT \leq 2T$  for any  $\phi \in \Omega$  we obtain

$$\begin{aligned}
 D_\psi(\phi^*, \phi_1) &\leq \psi(\phi^*) - \psi(\phi_1) \\
 &= \frac{1}{\eta} \sum_{h=1}^H \sum_{(s,a)} \phi^*(s, a, h) \ln \phi^*(s, a, h) - \frac{1}{\eta} \sum_{h=1}^H \sum_{(s,a)} \phi_1(s, a, h) \ln \phi_1(s, a, h) \\
 &\leq \frac{2T \ln(2T)}{\eta} - \frac{2T}{\eta} \sum_{h=1}^H \sum_{(s,a)} \frac{\phi_1(s, a, h)}{2T} \ln \frac{\phi_1(s, a, h)}{2T}
 \end{aligned}$$

$$\begin{aligned} &\leq \frac{2T \ln(2T)}{\eta} + \frac{2T \ln(|\tilde{\Gamma}|H)}{\eta} \\ &= \mathcal{O}\left(\frac{T \ln(SAHT)}{\eta}\right); \end{aligned}$$

for the second term, we have

$$\begin{aligned} \sum_{k=1}^K \langle \phi_{\tilde{\pi}^*}, c_k \rangle &\leq 2 \sum_{k=1}^K \langle q_{\tilde{\pi}^*}, c_k \rangle \leq 2 \sum_{k=1}^K J_k^{\pi^*}(s_0) + \tilde{\mathcal{O}}(1) \\ &\leq 2 \sum_{k=1}^K J_k^{\pi^f}(s_0) + \tilde{\mathcal{O}}(1) \leq 2DK + \tilde{\mathcal{O}}(1), \end{aligned}$$

where the second inequality is by [Eq. \(7\)](#). Combining the above and plugging the choice of  $\eta$ , we arrive at

$$\sum_{k=1}^K \langle \phi_k - \phi^*, c_k \rangle \leq \mathcal{O}\left(\frac{T \ln(SAHT)}{\eta}\right) + 2\eta DK + \tilde{\mathcal{O}}(1) = \tilde{\mathcal{O}}\left(\sqrt{DTK}\right).$$

Finally, we apply [Lemma 7](#): with probability at least  $1 - \delta$ ,

$$\begin{aligned} R_K &\leq \sum_{k=1}^K \langle q_k - q_{\tilde{\pi}^*}, c_k \rangle + \lambda \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle + \frac{2 \ln(2/\delta)}{\lambda} + \tilde{\mathcal{O}}(1) \\ &= \sum_{k=1}^K \langle \phi_k - \phi^*, c_k \rangle + \lambda \sum_{k=1}^K \langle q_{\tilde{\pi}^*}, \mathbf{h} \circ c_k \rangle + \frac{2 \ln(2/\delta)}{\lambda} + \tilde{\mathcal{O}}(1) \\ &= \tilde{\mathcal{O}}\left(\sqrt{DTK}\right) + \lambda \sum_{k=1}^K \langle q_{\tilde{\pi}^*}, \mathbf{h} \circ c_k \rangle + \frac{2 \ln(2/\delta)}{\lambda} \\ &= \tilde{\mathcal{O}}\left(\sqrt{DTK}\right) + \tilde{\mathcal{O}}(\lambda DTK) + \frac{2 \ln(2/\delta)}{\lambda} \tag{Lemma 15} \\ &= \tilde{\mathcal{O}}\left(\sqrt{DTK \ln(1/\delta)}\right), \tag{by the choice of } \lambda \end{aligned}$$

which finishes the proof. ■

## Appendix B. Omitted details for [Section 4](#)

In this section, we provide all omitted algorithms and proofs for [Section 4](#).

### B.1. Optimal Expected Regret

**Proof** [of [Theorem 9](#)] Using the second statement of [Lemma 7](#), we have

$$\mathbb{E}[R_K] = \mathbb{E}\left[\sum_{k=1}^K \langle q_k - q_{\tilde{\pi}^*}, c_k \rangle\right] + \tilde{\mathcal{O}}(1).$$

As in all analysis for OMD with log-barrier regularizer, we consider a slightly perturbed benchmark  $q^* = (1 - \frac{1}{TK})q_{\tilde{\pi}^*} + \frac{1}{TK}q_1$  which is in  $\tilde{\Delta}(T)$  by the convexity of  $\tilde{\Delta}(T)$ , the condition  $T \geq T_* + 1$ , and [Lemma 14](#). We then have

$$\begin{aligned} \mathbb{E}[R_K] &\leq \mathbb{E} \left[ \sum_{k=1}^K \langle q_k - q^*, c_k \rangle \right] + \frac{1}{TK-1} \mathbb{E} \left[ \sum_{k=1}^K \langle q_1, c_k \rangle \right] + \tilde{O}(1) \\ &= \mathbb{E} \left[ \sum_{k=1}^K \langle q_k - q^*, c_k \rangle \right] + \tilde{O}(1). \end{aligned}$$

It remains to bound  $\mathbb{E} \left[ \sum_{k=1}^K \langle q_k - q^*, c_k \rangle \right]$ . Let  $\phi^* = q^* + \lambda \mathbf{h} \circ q^* \in \Omega$ . By the non-negativity and the unbiasedness of the cost estimator, the obliviousness of the adversary, and the same argument of ([Agarwal et al., 2017](#), Lemma 12), OMD with log-barrier regularizer ensures

$$\mathbb{E} \left[ \sum_{k=1}^K \langle \phi_k - \phi^*, c_k \rangle \right] = \mathbb{E} \left[ \sum_{k=1}^K \langle \phi_k - \phi^*, \hat{c}_k \rangle \right] \leq D_\psi(\phi^*, \phi_1) + \eta \mathbb{E} \left[ \sum_{k=1}^K \sum_{(s,a)} \phi_k^2(s,a) \hat{c}_k^2(s,a) \right].$$

For the first term, as  $\phi_1$  minimizes  $\psi$ , we have  $\langle \nabla \psi(\phi_1), \phi^* - \phi_1 \rangle \geq 0$  and thus

$$D_\psi(\phi^*, \phi_1) \leq \frac{1}{\eta} \sum_{(s,a)} \ln \frac{\phi_1(s,a)}{\phi^*(s,a)} = \frac{SA}{\eta} \ln(HT) = \tilde{O} \left( \frac{SA}{\eta} \right).$$

For the second term, we note that

$$\begin{aligned} \mathbb{E} \left[ \sum_{(s,a)} \phi_k^2(s,a) \hat{c}_k^2(s,a) \right] &\leq 4 \mathbb{E} \left[ \sum_{(s,a)} q_k^2(s,a) \hat{c}_k^2(s,a) \right] = 4 \mathbb{E} \left[ \sum_{(s,a)} \tilde{N}_k^2(s,a) c_k^2(s,a) \right] \\ &\leq 4 \mathbb{E} \left[ \langle \tilde{N}_k, c_k \rangle^2 \right] \leq 8 \mathbb{E} [\langle q_k, \mathbf{h} \circ c_k \rangle], \end{aligned}$$

where the last step is by [Lemma 6](#). Combining everything, we have shown

$$\mathbb{E} \left[ \sum_{k=1}^K \langle \phi_k - \phi^*, c_k \rangle \right] = \tilde{O} \left( \frac{SA}{\eta} \right) + 8\eta \mathbb{E} \left[ \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle \right],$$

and thus

$$\begin{aligned} \mathbb{E} \left[ \sum_{k=1}^K \langle q_k - q^*, c_k \rangle \right] &= \mathbb{E} \left[ \sum_{k=1}^K \langle \phi_k - \phi^*, c_k \rangle \right] + \lambda \mathbb{E} \left[ \sum_{k=1}^K \langle q^*, \mathbf{h} \circ c_k \rangle \right] - \lambda \mathbb{E} \left[ \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle \right] \\ &= \tilde{O} \left( \frac{SA}{\eta} \right) + 8\eta \mathbb{E} \left[ \sum_{k=1}^K \langle q^*, \mathbf{h} \circ c_k \rangle \right] \quad (\lambda = 8\eta) \\ &= \tilde{O} \left( \frac{SA}{\eta} + \eta DTK \right). \quad (\text{Lemma 15}) \end{aligned}$$

Plugging the choice of  $\eta$  finishes the proof.  $\blacksquare$

## B.2. Proof of Theorem 10

**Proof** By Yao's minimax principle, in order to obtain a regret lower bound, it suffices to show that there exists a distribution of SSP instances that forces any deterministic learner to suffer a regret bound of  $\Omega(\sqrt{DT_*SAK})$  in expectation. We use the exact same construction as in Theorem 3 with  $N = S - 2$  (note that the average number of actions  $A$  is  $\mathcal{O}(1)$ ). The proof is the same up to the point where we show

$$\mathbb{E}[R_K] \geq T_*K\epsilon - \frac{T_*\epsilon}{N} \sum_{j=1}^N \mathbb{E}_j[K_j],$$

with  $K_j = \sum_{k=1}^K \mathbb{I}\{a_k^1 = a_j\}$ , and

$$\mathbb{E}_j[K_j] - \mathbb{E}_0[K_j] \leq K\|P_j - P_0\|_1 \leq K\sqrt{2\text{KL}(P_0, P_j)}.$$

What is different is the usage of the divergence decomposition lemma (see e.g. (Lattimore and Szepesvári, 2020, Lemma 15.1)) due to the different observation model:

$$\begin{aligned} \text{KL}(P_0, P_j) &= \sum_{j'=1}^N \mathbb{E}_0[K_{j'}] \times \text{KL}\left(\text{Bernoulli}\left(\frac{D}{2T_*} + \epsilon\right), \text{Bernoulli}\left(\frac{D}{2T_*} + \epsilon\mathbb{I}\{j' \neq j\}\right)\right) \\ &= \mathbb{E}_0[K_j] \times \text{KL}\left(\text{Bernoulli}\left(\frac{D}{2T_*} + \epsilon\right), \text{Bernoulli}\left(\frac{D}{2T_*}\right)\right) \\ &\leq \frac{\mathbb{E}_0[K_j]\epsilon^2}{\alpha(1-\alpha)}, \end{aligned}$$

where the last step is again by (Gerchinovitz and Lattimore, 2016, Lemma 6) with  $\alpha = \frac{D}{2T_*}$ . Therefore, we can upper bound  $\sum_{j=1}^N \mathbb{E}_j[K_j]$  as:

$$\begin{aligned} \sum_{j=1}^N \mathbb{E}_j[K_j] &\leq \sum_{j=1}^N \mathbb{E}_0[K_j] + K\sqrt{\frac{2\epsilon^2}{\alpha(1-\alpha)}} \sum_{j=1}^N \sqrt{\mathbb{E}_0[K_j]} \\ &\leq \sum_{j=1}^N \mathbb{E}_0[K_j] + K\sqrt{\frac{2N\epsilon^2}{\alpha(1-\alpha)}} \sum_{j=1}^N \mathbb{E}_0[K_j] \quad (\text{Cauchy-Schwarz inequality}) \\ &= K + K\sqrt{\frac{2NK\epsilon^2}{\alpha(1-\alpha)}}. \quad (\sum_{j=1}^N \mathbb{E}_0[K_j] = K) \end{aligned}$$

This shows the following lower bound:

$$\begin{aligned} \mathbb{E}[R_K] &\geq T_*\epsilon K - \frac{T_*\epsilon}{N} \sum_{j=1}^N \mathbb{E}_j[K_j] \\ &\geq T_*\epsilon K - \frac{T_*\epsilon}{N} \left[ K + K\sqrt{\frac{2NK\epsilon^2}{\alpha(1-\alpha)}} \right] \end{aligned}$$

$$\begin{aligned}
 &= T_* \epsilon K \left[ 1 - \frac{1}{N} - \epsilon \sqrt{\frac{2K}{N\alpha(1-\alpha)}} \right] \\
 &\geq T_* \epsilon K \left[ \frac{1}{2} - \epsilon \sqrt{\frac{2K}{N\alpha(1-\alpha)}} \right] \\
 &= \frac{T_* K}{16} \sqrt{\frac{N\alpha(1-\alpha)}{2K}} = \Omega\left(\sqrt{DT_* NK}\right) = \Omega\left(\sqrt{DT_* SAK}\right),
 \end{aligned}$$

where in the last step we set  $\epsilon = \frac{1}{4} \sqrt{\frac{N\alpha(1-\alpha)}{2K}} \leq \frac{1}{8} \sqrt{\frac{SD}{T_* K}} \leq \frac{D}{2T_*}$  to maximize the lower bound. ■

### B.3. Optimal High-probability Regret

We present our algorithm with optimal high-probability regret in [Algorithm 5](#). The key difference compared to [Algorithm 4](#) is the use of the extra bias term  $\widehat{b}_k$  in the OMD update and the time-varying individual learning  $\eta_k(s, a)$  for each state-action pair together with an increasing learning rate schedule (see the last for loop). Note that, similar to ([Lee et al., 2020a](#)), the decision set  $\Omega$  has the extra constraint  $q(s, a) \geq \frac{1}{TK^4}$  compared to [Algorithm 3](#) and [Algorithm 4](#), and it is always non-empty as long as  $K$  is large enough and every state is reachable within  $H$  steps starting from  $s_0$  (states not satisfying this can simply be removed without affecting  $\widetilde{M}$ ).

Below we present the proof of [Theorem 11](#). It decomposes the regret into several terms, each of which is bounded by a lemma included after the proof.

**Proof [of [Theorem 11](#)]** We apply the first statement of [Lemma 7](#): with probability  $1 - \delta$ ,

$$R_K \leq \sum_{k=1}^K \left\langle \widetilde{N}_k - q_{\pi^*}, c_k \right\rangle + \widetilde{O}(1).$$

Similar to the proof of [Theorem 9](#), we define a slightly perturbed benchmark  $q^* = (1 - \frac{1}{TK})q_{\pi^*} + \frac{1}{TK}q_0 \in \widetilde{\Delta}(T)$  for some  $q_0 \in \widetilde{\Delta}(T)$  with  $q_0(s, a) \geq \frac{1}{K^3}$  for all  $(s, a) \in \widetilde{\Gamma}$  (which again exists as long as  $K$  is large enough), so that  $R_K \leq \sum_{k=1}^K \left\langle \widetilde{N}_k - q^*, c_k \right\rangle + \widetilde{O}(1)$  still holds. Also define  $\phi^* = q^* + \lambda \mathbf{h} \circ q^* \in \Omega$  and  $b_k \in \mathbb{R}^{\widetilde{\Gamma}}$  such that  $b_k(s, a) = \frac{\sum_h h q_k(s, a, h) c_k(s, a)}{q_k(s, a)}$ , which clearly satisfies  $\mathbb{E}_k[\widehat{b}_k] = b_k$ . We then decompose  $\sum_{k=1}^K \left\langle \widetilde{N}_k - q^*, c_k \right\rangle$  as

$$\begin{aligned}
 &\sum_{k=1}^K \left\langle \widetilde{N}_k - q^*, c_k \right\rangle \\
 &= \sum_{k=1}^K \left\langle q_k, \widehat{c}_k \right\rangle - \sum_{k=1}^K \left\langle q^*, c_k \right\rangle \quad (\langle \widetilde{N}_k, c_k \rangle = \langle q_k, \widehat{c}_k \rangle) \\
 &= \sum_{k=1}^K \left\langle \phi_k - \phi^*, \widehat{c}_k \right\rangle + \sum_{k=1}^K \left\langle \phi^*, \widehat{c}_k - c_k \right\rangle + \lambda \sum_{k=1}^K \left\langle \mathbf{h} \circ q^*, c_k \right\rangle - \lambda \sum_{k=1}^K \left\langle \mathbf{h} \circ q_k, \widehat{c}_k \right\rangle
 \end{aligned}$$

---

**Algorithm 5** Log-barrier Policy Search for SSP (High Probability)
 

---

**Input:** Upper bound on expected hitting time  $T$ , horizon parameter  $H_1$ , and confidence level  $\delta$

**Parameters:**  $H_2 = \lceil 4D \ln \frac{4K}{\delta} \rceil$ ,  $H = H_1 + H_2$ ,  $C = \lceil \log_2(TK^4) \rceil \lceil \log_2(T^2K^9) \rceil$ ,  $\beta = e^{\frac{1}{7 \ln K}}$ ,  $\eta = \sqrt{\frac{SA \ln^{1/\delta}}{DT_* K}}$ ,  $\gamma = 100\eta \ln K \left(1 + C \sqrt{8 \ln \frac{CSA}{\delta}}\right)^2$ ,  $\lambda = 40\eta + 2\gamma$ .

**Define:** regularizer  $\psi_k(\phi) = \sum_{(s,a) \in \tilde{\Gamma}} \frac{1}{\eta_k(s,a)} \ln \frac{1}{\phi(s,a)}$  where  $\phi(s,a) = \sum_{h=1}^H \phi(s,a,h)$

**Define:** decision set  $\Omega = \{\phi = q + \lambda \mathbf{h} \circ q : q \in \tilde{\Delta}(T), q(s,a) \geq \frac{1}{TK^4}, \forall (s,a) \in \tilde{\Gamma}\}$

**Initialization:**  $\phi_1 = q_1 + \lambda \mathbf{h} \circ q_1 = \operatorname{argmin}_{\phi \in \Omega} \psi_1(\phi)$ .

**Initialization:** for all  $(s,a) \in \tilde{\Gamma}$ ,  $\eta_1(s,a) = \eta$ ,  $\rho_1(s,a) = 2T$ .

**for**  $k = 1, \dots, K$  **do**

Execute  $\sigma(\tilde{\pi}_k)$  where  $\tilde{\pi}_k$  is such that  $\tilde{\pi}_k(a|(s,h)) \propto q_k(s,a,h)$ .

Construct cost estimator  $\hat{c}_k \in \mathbb{R}_{\geq 0}^{\tilde{\Gamma}}$  such that  $\hat{c}_k(s,a) = \frac{\tilde{N}_k(s,a)c_k(s,a)}{q_k(s,a)}$  where  $\tilde{N}_k(s,a) = \sum_h \tilde{N}_k(s,a,h)$  and  $q_k(s,a) = \sum_h q_k(s,a,h)$  ( $\tilde{N}_k$  is defined after [Lemma 6](#)).

Construct bias term  $\hat{b}_k \in \mathbb{R}_{\geq 0}^{\tilde{\Gamma}}$  such that  $\hat{b}_k(s,a) = \frac{\sum_h h q_k(s,a,h) \hat{c}_k(s,a)}{q_k(s,a)}$ .

Update

$$\phi_{k+1} = q_{k+1} + \lambda \mathbf{h} \circ q_{k+1} = \operatorname{argmin}_{\phi \in \Omega} \sum_{(s,a)} \phi(s,a) \left( \hat{c}_k(s,a) - \gamma \hat{b}_k(s,a) \right) + D_{\psi_k}(\phi, \phi_k).$$

**for**  $\forall (s,a) \in \tilde{\Gamma}$  **do**

**if**  $\frac{1}{\phi_{k+1}(s,a)} > \rho_k(s,a)$  **then**

$$\left[ \rho_{k+1}(s,a) = \frac{2}{\phi_{k+1}(s,a)}, \eta_{k+1}(s,a) = \beta \eta_k(s,a). \right]$$

**else**

$$\left[ \rho_{k+1}(s,a) = \rho_k(s,a), \eta_{k+1}(s,a) = \eta_k(s,a). \right]$$

$$\begin{aligned} &= \sum_{k=1}^K \langle \phi_k - \phi^*, \hat{c}_k \rangle + \sum_{k=1}^K \langle \phi^*, \hat{c}_k - c_k \rangle + \tilde{O}(\lambda DTK) - \lambda \sum_{k=1}^K \langle \mathbf{h} \circ q_k, \hat{c}_k \rangle \quad (\text{Lemma 15}) \\ &= \sum_{k=1}^K \langle \phi_k - \phi^*, \hat{c}_k \rangle + \tilde{O}(\lambda DTK) + \text{DEV}_1 + \text{DEV}_2 - \lambda \sum_{k=1}^K \langle \mathbf{h} \circ q_k, c_k \rangle \\ &\quad (\text{define } \text{DEV}_1 = \sum_{k=1}^K \langle \phi^*, \hat{c}_k - c_k \rangle \text{ and } \text{DEV}_2 = \lambda \sum_{k=1}^K \langle \mathbf{h} \circ q_k, c_k - \hat{c}_k \rangle) \\ &= \text{REG}_\phi + \tilde{O}(\lambda DTK) + \text{DEV}_1 + \text{DEV}_2 + \gamma \sum_{k=1}^K \langle \phi_k - \phi^*, \hat{b}_k \rangle - \lambda \sum_{k=1}^K \langle \mathbf{h} \circ q_k, c_k \rangle \\ &\quad (\text{define } \text{REG}_\phi = \sum_{k=1}^K \langle \phi_k - \phi^*, \hat{c}_k - \gamma \hat{b}_k \rangle) \\ &= \text{REG}_\phi + \tilde{O}(\lambda DTK) + \text{DEV}_1 + \text{DEV}_2 + \text{DEV}_3 + \text{DEV}_4 \\ &\quad + \gamma \sum_{k=1}^K \langle \phi_k - \phi^*, b_k \rangle - \lambda \sum_{k=1}^K \langle \mathbf{h} \circ q_k, c_k \rangle \\ &\quad (\text{define } \text{DEV}_3 = \gamma \sum_{k=1}^K \langle \phi_k, \hat{b}_k - b_k \rangle \text{ and } \text{DEV}_4 = \gamma \sum_{k=1}^K \langle \phi^*, b_k - \hat{b}_k \rangle) \end{aligned}$$

$$\begin{aligned}
 &\leq \text{REG}_\phi + \tilde{\mathcal{O}}(\lambda DTK) + \text{DEV}_1 + \text{DEV}_2 + \text{DEV}_3 + \text{DEV}_4 \\
 &\quad + 2\gamma \sum_{k=1}^K \langle q_k, b_k \rangle - \gamma \sum_{k=1}^K \langle \phi^*, b_k \rangle - \lambda \sum_{k=1}^K \langle \mathbf{h} \circ q_k, c_k \rangle \\
 &= \text{REG}_\phi + \tilde{\mathcal{O}}(\lambda DTK) + \text{DEV}_1 + \text{DEV}_2 + \text{DEV}_3 + \text{DEV}_4 \\
 &\quad + (2\gamma - \lambda) \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle - \gamma \sum_{k=1}^K \langle \phi^*, b_k \rangle. \\
 &\hspace{20em} (\langle q_k, b_k \rangle = \langle q_k, \mathbf{h} \circ c_k \rangle)
 \end{aligned}$$

The  $\text{REG}_\phi$  term can be upper bounded by the OMD analysis (see [Lemma 16](#)), and the four deviation terms  $\text{DEV}_1, \text{DEV}_2, \text{DEV}_3$ , and  $\text{DEV}_4$  are all sums of martingale difference sequences and can be bounded using Azuma's or Freedman's inequality (see [Lemma 17](#) and [Lemma 18](#)). Combining everything, we obtain

$$\begin{aligned}
 R_K &\leq \tilde{\mathcal{O}}\left(\frac{SA}{\eta}\right) - \frac{\langle \phi^*, \rho_K \rangle}{70\eta \ln K} + 40\eta \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle + \tilde{\mathcal{O}}(\lambda DTK) \\
 &\quad + \left(1 + C\sqrt{8 \ln\left(\frac{CSA}{\delta}\right)}\right) \left(\frac{\langle \phi^*, \rho_K \rangle}{\eta'} + \eta' \left\langle \phi^*, \sum_{k=1}^K b_k \right\rangle\right) + \left(4CH \ln\left(\frac{CSA}{\delta}\right)\right) \langle \phi^*, \rho_K \rangle \\
 &\quad + (2\gamma - \lambda) \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle - \gamma \sum_{k=1}^K \langle \phi^*, b_k \rangle \\
 &= \tilde{\mathcal{O}}\left(\frac{SA}{\eta} + \lambda DTK\right) + \left(\frac{1 + C\sqrt{8 \ln\left(\frac{CSA}{\delta}\right)}}{\eta'} + 4CH \ln\left(\frac{CSA}{\delta}\right) - \frac{1}{70\eta \ln K}\right) \langle \phi^*, \rho_K \rangle \\
 &\quad + (40\eta + 2\gamma - \lambda) \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle + \left(\left(1 + C\sqrt{8 \ln\left(\frac{CSA}{\delta}\right)}\right) \eta' - \gamma\right) \sum_{k=1}^K \langle \phi^*, b_k \rangle.
 \end{aligned}$$

Finally, note that  $\eta' \geq 0$  from [Lemma 17](#) and [Lemma 18](#) can be chosen arbitrarily. Setting  $\eta' = \gamma / \left(1 + C\sqrt{8 \ln\left(\frac{CSA}{\delta}\right)}\right)$ , and plugging the choice of  $\gamma = 100\eta \ln K \left(1 + C\sqrt{8 \ln\left(\frac{CSA}{\delta}\right)}\right)^2$  and  $\lambda = 40\eta + 2\gamma$ , one can see that the coefficients multiplying the last three terms  $\langle \phi^*, \rho_K \rangle$ ,  $\sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle$ , and  $\sum_{k=1}^K \langle \phi^*, b_k \rangle$  are all non-positive. Therefore, we arrive at

$$R_K = \tilde{\mathcal{O}}\left(\frac{SA}{\eta} + \eta DTK \ln(1/\delta)\right) = \tilde{\mathcal{O}}\left(\sqrt{DTSAK \ln(1/\delta)}\right),$$

where the last step is by the choice of  $\eta$ . ■

**Lemma 16** [Algorithm 5](#) ensures with probability at least  $1 - \delta$ :

$$\text{REG}_\phi \leq \tilde{\mathcal{O}}\left(\frac{SA}{\eta}\right) - \frac{\langle \phi^*, \rho_K \rangle}{70\eta \ln K} + 40\eta \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle + \tilde{\mathcal{O}}\left(H^2 \sqrt{SA} \ln(1/\delta)\right).$$

**Proof** Denote by  $n(s, a)$  the number of times the learning rate for  $(s, a)$  increases, such that  $\eta_K(s, a) = \eta\beta^{n(s, a)}$ , and by  $k_1, \dots, k_{n(s, a)}$  the episodes where  $\eta_k(s, a)$  is increased, such that  $\eta_{k_{t+1}}(s, a) = \beta \cdot \eta_{k_t}(s, a)$ . Since  $\rho_1(s, a) = 2T$  and

$$\rho_1(s, a)2^{n(s, a)-1} \leq \dots \leq \rho_{k_{n(s, a)}}(s, a) \leq \frac{1}{\phi_{k_{n(s, a)+1}}(s, a)} \leq \frac{1}{q_{k_{n(s, a)+1}}(s, a)} \leq TK^4,$$

we have  $n(s, a) \leq 1 + \log_2 \frac{K^4}{2} \leq 7 \log_2 K$ . Therefore,  $\eta_K(s, a) \leq \eta e^{\frac{7 \log_2 K}{7 \ln K}} \leq 5\eta$ .

Now, notice that

$$\gamma \widehat{b}_k(s, a) \leq \frac{\gamma H \sum_h q_k(s, a, h) \widehat{c}_k(s, a)}{q_k(s, a)} = \gamma H \widehat{c}_k(s, a) \leq \widehat{c}_k(s, a).$$

This means that the cost  $\widehat{c}_k - \gamma \widehat{b}_k$  we feed to OMD is always non-negative, and thus by the same argument of (Agarwal et al., 2017, Lemma 12), we have

$$\begin{aligned} \text{REG}_\phi &= \sum_{k=1}^K \left\langle \phi_k - \phi^*, \widehat{c}_k - \gamma \widehat{b}_k \right\rangle \\ &\leq \sum_{k=1}^K D_{\psi_k}(\phi^*, \phi_k) - D_{\psi_k}(\phi^*, \phi_{k+1}) + \sum_{k=1}^K \sum_{(s, a)} \eta_k(s, a) \phi_k^2(s, a) (\widehat{c}_k(s, a) - \gamma \widehat{b}_k(s, a))^2 \\ &\leq D_{\psi_1}(\phi^*, \phi_1) + \sum_{k=1}^{K-1} (D_{\psi_{k+1}}(\phi^*, \phi_{k+1}) - D_{\psi_k}(\phi^*, \phi_{k+1})) + 5\eta \sum_{k=1}^K \sum_{(s, a)} \phi_k^2(s, a) \widehat{c}_k^2(s, a) \\ &\leq D_{\psi_1}(\phi^*, \phi_1) + \sum_{k=1}^{K-1} (D_{\psi_{k+1}}(\phi^*, \phi_{k+1}) - D_{\psi_k}(\phi^*, \phi_{k+1})) + 20\eta \sum_{k=1}^K \sum_{(s, a)} q_k^2(s, a) \widehat{c}_k^2(s, a) \\ &= D_{\psi_1}(\phi^*, \phi_1) + \sum_{k=1}^{K-1} (D_{\psi_{k+1}}(\phi^*, \phi_{k+1}) - D_{\psi_k}(\phi^*, \phi_{k+1})) + 20\eta \sum_{k=1}^K \sum_{(s, a)} \widetilde{N}_k^2(s, a) c_k^2(s, a). \end{aligned}$$

For the first term, since  $\phi_1$  minimizes  $\psi_1$  and thus  $\langle \nabla \psi_1(\phi_1), \phi^* - \phi_1 \rangle \geq 0$ , we have

$$D_{\psi_1}(\phi^*, \phi_1) \leq \psi_1(\phi^*) - \psi_1(\phi_1) = \frac{1}{\eta} \sum_{(s, a)} \ln \frac{\phi_1(s, a)}{\phi^*(s, a)} \leq \frac{1}{\eta} \sum_{(s, a)} \ln \frac{2H}{q^*(s, a)} = \widetilde{\mathcal{O}} \left( \frac{SA}{\eta} \right).$$

For the second term, we define  $\chi(y) = y - 1 - \ln y$  and proceed similarly to (Agarwal et al., 2017):

$$\begin{aligned} &\sum_{k=1}^{K-1} D_{\psi_{k+1}}(\phi^*, \phi_{k+1}) - D_{\psi_k}(\phi^*, \phi_{k+1}) \\ &= \sum_{k=1}^{K-1} \sum_{(s, a)} \left( \frac{1}{\eta_{k+1}(s, a)} - \frac{1}{\eta_k(s, a)} \right) \chi \left( \frac{\phi^*(s, a)}{\phi_{k+1}(s, a)} \right) \\ &\leq \sum_{(s, a)} \frac{1 - \beta}{\eta \beta^{n(s, a)}} \chi \left( \frac{\phi^*(s, a)}{\phi_{k_{n(s, a)+1}}(s, a)} \right) \end{aligned}$$



$$\begin{aligned}
 &= \sum_{(s,a)} \frac{1-\beta}{\eta\beta^{n(s,a)}} \left( \frac{\phi^*(s,a)}{\phi_{k_n(s,a)+1}(s,a)} - 1 - \ln \frac{\phi^*(s,a)}{\phi_{k_n(s,a)+1}(s,a)} \right) \\
 &\leq -\frac{1}{35\eta \ln K} \sum_{(s,a)} \left( \phi^*(s,a) \frac{\rho_K(s,a)}{2} - 1 - \ln \frac{\phi^*(s,a)}{\phi_{k_n(s,a)+1}(s,a)} \right) \\
 &\leq \frac{SA(1+6\ln K)}{35\eta \ln K} - \frac{\langle \phi^*, \rho_K \rangle}{70\eta \ln K} = \tilde{\mathcal{O}} \left( \frac{SA}{\eta} \right) - \frac{\langle \phi^*, \rho_K \rangle}{70\eta \ln K},
 \end{aligned}$$

where in the last two lines we use the facts  $1-\beta \leq -\frac{1}{7\ln K}$ ,  $\beta^{n(s,a)} \leq 5$ ,  $\rho_K(s,a) = \frac{2}{\phi_{k_n(s,a)+1}(s,a)}$ , and  $\ln \frac{q^*(s,a)}{q_{k_n(s,a)+1}(s,a)} \leq \ln(HTK^4) \leq 6\ln K$ .

Finally, for the third term, since  $\sum_{(s,a)} \tilde{N}_k^2(s,a) c_k^2(s,a) \leq \left( \sum_{(s,a)} \tilde{N}_k(s,a) \right)^2 \leq H^2$ , we apply Azuma's inequality (Lemma 20) and obtain, with probability at least  $1-\delta$ :

$$\begin{aligned}
 \eta \sum_{k=1}^K \sum_{(s,a)} \tilde{N}_k^2(s,a) c_k^2(s,a) &\leq \eta \sum_{k=1}^K \mathbb{E}_k \left[ \sum_{(s,a)} \tilde{N}_k^2(s,a) c_k^2(s,a) \right] + \tilde{\mathcal{O}} \left( \eta H^2 \sqrt{K \ln(1/\delta)} \right) \\
 &\leq \eta \sum_{k=1}^K \mathbb{E}_k \left[ \langle \tilde{N}_k, c_k \rangle^2 \right] + \tilde{\mathcal{O}} \left( H^2 \sqrt{SA} \ln(1/\delta) \right) \\
 &\leq 2\eta \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle + \tilde{\mathcal{O}} \left( H^2 \sqrt{SA} \ln(1/\delta) \right). \quad (\text{Lemma 6})
 \end{aligned}$$

Combining everything shows

$$\text{REG}_\phi \leq \tilde{\mathcal{O}} \left( \frac{SA}{\eta} \right) - \frac{\langle \phi^*, \rho_K \rangle}{70\eta \ln K} + 40\eta \sum_{k=1}^K \langle q_k, \mathbf{h} \circ c_k \rangle + \tilde{\mathcal{O}} \left( H^2 \sqrt{SA} \ln(1/\delta) \right).$$

finishing the proof. ■

**Lemma 17** For any  $\eta' > 0$ , with probability at least  $1-\delta$ ,

$$\text{DEV}_1 \leq C \sqrt{8 \ln \left( \frac{CSA}{\delta} \right)} \left( \frac{\langle \phi^*, \rho_K \rangle}{\eta'} + \eta' \left\langle \phi^*, \sum_{k=1}^K b_k \right\rangle \right) + \left( 4CH \ln \left( \frac{CSA}{\delta} \right) \right) \langle \phi^*, \rho_K \rangle.$$

Also, with probability at least  $1-\delta$ ,  $\text{DEV}_2 = \tilde{\mathcal{O}} \left( H^2 \sqrt{SA} \ln(1/\delta) \right)$ .

**Proof** Define  $X_k(s,a) = \hat{c}_k(s,a) - c_k(s,a)$ . Note that

$$X_k(s,a) \leq \frac{H}{q_k(s,a)} \leq \frac{2H}{\phi_k(s,a)} \leq 2H\rho_k(s,a) \leq 4HTK^4,$$

and

$$\begin{aligned}
 \sum_{k=1}^K \mathbb{E}_k [X_k^2(s, a)] &\leq \sum_{k=1}^K \frac{\mathbb{E}_k [\tilde{N}_k^2(s, a) c_k^2(s, a)]}{q_k^2(s, a)} \\
 &\leq 2\rho_k(s, a) \sum_{k=1}^K \frac{\mathbb{E}_k [\tilde{N}_k^2(s, a) c_k^2(s, a)]}{q_k(s, a)} \\
 &= 4\rho_k(s, a) \sum_{k=1}^K b_k(s, a). \tag{Lemma 19}
 \end{aligned}$$

Therefore, by applying a strengthened Freedman's inequality (Lemma 22) with  $b = 4HTK^4$ ,  $B_k = 2H\rho_k(s, a)$ ,  $\max_k B_k = 2H\rho_K(s, a)$ , and  $V = 4\rho_k(s, a) \sum_{k=1}^K b_k(s, a)$ , we have with probability  $1 - \delta/(SA)$ ,

$$\begin{aligned}
 &\sum_{k=1}^K \hat{c}_k(s, a) - c_k(s, a) \\
 &\leq C \left( \sqrt{32\rho_K(s, a) \sum_{k=1}^K b_k(s, a) \ln\left(\frac{CSA}{\delta}\right)} + 4H\rho_K(s, a) \ln\left(\frac{CSA}{\delta}\right) \right) \\
 &\leq C \sqrt{8 \ln\left(\frac{CSA}{\delta}\right)} \left( \frac{\rho_K(s, a)}{\eta'} + \eta' \sum_{k=1}^K b_k(s, a) \right) + 4CH\rho_K(s, a) \ln\left(\frac{CSA}{\delta}\right),
 \end{aligned}$$

where the last step is by AM-GM inequality. Further using a union bound shows that the above holds for all  $(s, a) \in \tilde{\Gamma}$  with probability  $1 - \delta$  and thus

$$\begin{aligned}
 \text{DEV}_1 &= \sum_{k=1}^K \langle \phi^*, \hat{c}_k - c_k \rangle \\
 &\leq C \sqrt{8 \ln\left(\frac{CSA}{\delta}\right)} \left( \frac{\langle \phi^*, \rho_K \rangle}{\eta'} + \eta' \left\langle \phi^*, \sum_{k=1}^K b_k \right\rangle \right) + \left( 4CH \ln\left(\frac{CSA}{\delta}\right) \right) \langle \phi^*, \rho_K \rangle.
 \end{aligned}$$

To bound  $\text{DEV}_2$ , simply note that  $|\langle \mathbf{h} \circ q_k, c_k - \hat{c}_k \rangle| \leq 2H^2$  and apply Azuma's inequality (Lemma 20): with probability  $1 - \delta$ ,

$$\text{DEV}_2 = \lambda \sum_{k=1}^K \langle \mathbf{h} \circ q_k, c_k - \hat{c}_k \rangle = \mathcal{O}\left(\lambda H^2 \sqrt{K \ln(1/\delta)}\right) = \tilde{\mathcal{O}}\left(H^2 \sqrt{SA} \ln(1/\delta)\right).$$

This completes the proof. ■

**Lemma 18** *With probability at least  $1 - \delta$ , we have  $\text{DEV}_3 = \tilde{\mathcal{O}}\left(H^2 \sqrt{SA} \ln(1/\delta)\right)$ . Also, for any  $\eta' > 0$ , with probability at least  $1 - \delta$ , we have*

$$\text{DEV}_4 \leq \frac{\langle \phi^*, \rho_K \rangle}{\eta'} + \eta' \left\langle \phi^*, \sum_{k=1}^K b_k \right\rangle + \tilde{\mathcal{O}}(1).$$

**Proof** To bound  $\text{DEV}_3$ , simply note that

$$\left| \langle \phi_k, \widehat{b}_k - b_k \rangle \right| \leq 4H \langle q_k, \widehat{c}_k \rangle \leq 4H \left( \sum_{(s,a)} \widetilde{N}_k(s,a) \right) \leq 4H^2$$

and apply Azuma's inequality: with probability  $1 - \delta$ ,

$$\text{DEV}_3 = \gamma \sum_{k=1}^K \langle \phi_k, \widehat{b}_k - b_k \rangle = \tilde{\mathcal{O}} \left( \gamma H^2 \sqrt{K \ln(1/\delta)} \right) = \tilde{\mathcal{O}} \left( H^2 \sqrt{SA} \ln(1/\delta) \right).$$

To bound  $\text{DEV}_4 = \gamma \sum_{k=1}^K \langle \phi^*, b_k - \widehat{b}_k \rangle$ , we note that  $b_k(s,a) - \widehat{b}_k(s,a) \leq b_k(s,a) \leq H$ , and

$$\begin{aligned} \mathbb{E}_k \left[ b_k(s,a) - \widehat{b}_k(s,a) \right]^2 &\leq \frac{\mathbb{E}_k \left[ (\sum_h h \cdot q_k(s,a,h) \widehat{c}_k(s,a))^2 \right]}{q_k^2(s,a)} \\ &\leq \frac{H^2 \mathbb{E}_k \left[ \widetilde{N}_k^2(s,a) c_k^2(s,a) \right]}{q_k^2(s,a)} \\ &\leq 2H^2 \rho_K(s,a) \frac{\mathbb{E}_k \left[ \widetilde{N}_k^2(s,a) c_k^2(s,a) \right]}{q_k(s,a)} \\ &\leq 4H^2 \rho_K(s,a) b_k(s,a). \end{aligned} \quad (\text{Lemma 19})$$

Hence, applying a strengthened Freedman's inequality (Lemma 22) with  $b = B_i = H$ ,  $V = 4H^2 \rho_K(s,a) \sum_{k=1}^K b_k(s,a)$ , and  $C' = \lceil \log_2 H \rceil \lceil \log_2(H^2 K) \rceil$ , we have with probability at least  $1 - \delta/(SA)$ ,

$$\begin{aligned} &\sum_{k=1}^K b_k(s,a) - \widehat{b}_k(s,a) \\ &\leq 2C'H \sqrt{\ln \left( \frac{C'SA}{\delta} \right)} \sqrt{8\rho_K(s,a) \sum_{k=1}^K b_k(s,a) + 2C'H \ln \left( \frac{C'SA}{\delta} \right)} \\ &= 2C'H \sqrt{\ln \left( \frac{C'SA}{\delta} \right)} \left( \frac{2\rho_K(s,a)}{\eta'} + \eta' \sum_k b_k(s,a) \right) + 2C'H \ln \left( \frac{C'SA}{\delta} \right), \end{aligned}$$

where the last step is by AM-GM inequality. Finally, applying a union bound shows that the above holds for all  $(s,a) \in \widetilde{\Gamma}$  with probability at least  $1 - \delta$  and thus

$$\text{DEV}_4 = \gamma \sum_{k=1}^K \langle \phi^*, b_k - \widehat{b}_k \rangle \leq \frac{\langle \phi^*, \rho_K \rangle}{\eta'} + \eta' \left\langle \phi^*, \sum_{k=1}^K b_k \right\rangle + \tilde{\mathcal{O}}(1),$$

where we bound  $\gamma C'H \sqrt{\ln \left( \frac{C'SA}{\delta} \right)}$  by a constant since  $\gamma$  is of order  $1/\sqrt{K}$  and is small enough when  $K$  is large.  $\blacksquare$

**Lemma 19** For any episode  $k$  and  $(s, a) \in \tilde{\Gamma}$ :  $\mathbb{E}_k \left[ \tilde{N}_k(s, a)^2 c_k(s, a)^2 \right] \leq 2q_k(s, a)b_k(s, a)$ .

**Proof** The proof is similar to those of [Lemma 4](#) and uses  $(\sum_{i=1}^I a_i)^2 \leq 2 \sum_i a_i (\sum_{i'=1}^I a_{i'})$ :

$$\begin{aligned}
 \mathbb{E}_k \left[ \tilde{N}_k(s, a)^2 c_k(s, a)^2 \right] &\leq \mathbb{E}_k \left[ \left( \sum_{h=1}^H \tilde{N}_k(s, a, h) \right)^2 c_k(s, a) \right] \\
 &\leq 2\mathbb{E}_k \left[ \left( \sum_{h=1}^H \tilde{N}_k(s, a, h) \right) \left( \sum_{h' \geq h}^H \tilde{N}_k(s, a, h') c_k(s, a) \right) \right] \\
 &\leq 2\mathbb{E}_k \left[ \sum_{h=1}^H \sum_{h' \geq h}^H \tilde{N}_k(s, a, h') c_k(s, a) \right] \quad (\tilde{N}_k(s, a, h) \in \{0, 1\}) \\
 &= 2 \sum_{h=1}^H \sum_{h' \geq h}^H q_k(s, a, h') c_k(s, a) \\
 &= 2 \sum_{h=1}^H h q_k(s, a, h) c_k(s, a) = 2q_k(s, a)b_k(s, a),
 \end{aligned}$$

where the last step is by the definition of  $b_k$ . ■

### Appendix C. Concentration Inequalities

**Lemma 20** (Azuma's inequality) Let  $X_{1:n}$  be a martingale difference sequence and  $|X_i| \leq B$  holds for  $i = 1, \dots, n$  and some fixed  $B > 0$ . Then, with probability at least  $1 - \delta$ :

$$\left| \sum_{i=1}^n X_i \right| \leq B \sqrt{2n \ln \frac{2}{\delta}}.$$

**Lemma 21** (A version of Freedman's inequality from [\(Beygelzimer et al., 2011\)](#)) Let  $X_{1:n}$  be a martingale difference sequence and  $X_i \leq B$  holds for  $i = 1, \dots, n$  and some fixed  $B > 0$ . Denote  $V = \sum_{i=1}^n \mathbb{E}_i[X_i^2]$ . Then, for any  $\lambda \in [0, 1/B]$ , with probability at least  $1 - \delta$ :

$$\sum_{i=1}^n X_i \leq \lambda V + \frac{\ln(1/\delta)}{\lambda}.$$

**Lemma 22** (Strengthened Freedman's inequality from [\(Lee et al., 2020a, Theorem 2.2\)](#)) Let  $X_{1:n}$  be a martingale difference sequence with respect to a filtration  $\mathcal{F}_1 \subseteq \dots \subseteq \mathcal{F}_n$  such that  $\mathbb{E}[X_i | \mathcal{F}_i] = 0$ . Suppose  $B_i \in [1, b]$  for a fixed constant  $b$  is  $\mathcal{F}_i$ -measurable and such that  $X_i \leq B_i$  holds almost surely. Then with probability at least  $1 - \delta$  we have

$$\sum_{i=1}^n X_i \leq C \left( \sqrt{8V \ln(C/\delta)} + 2B^* \ln(C/\delta) \right),$$

where  $V = \max \{1, \sum_{i=1}^n \mathbb{E}[X_i^2 | \mathcal{F}_i]\}$ ,  $B^* = \max_{i \in [n]} B_i$ , and  $C = \lceil \ln(b) \rceil \lceil \ln(nb^2) \rceil$ .