
Simultaneously Learning Stochastic and Adversarial Episodic MDPs with Known Transition

Tiancheng Jin

University of Southern California
tiancheng.jin@usc.edu

Haipeng Luo

University of Southern California
haipengl@usc.edu

Abstract

This work studies the problem of learning episodic Markov Decision Processes with known transition and bandit feedback. We develop the first algorithm with a “best-of-both-worlds” guarantee: it achieves $\mathcal{O}(\log T)$ regret when the losses are stochastic, and simultaneously enjoys worst-case robustness with $\tilde{\mathcal{O}}(\sqrt{T})$ regret even when the losses are adversarial, where T is the number of episodes. More generally, it achieves $\tilde{\mathcal{O}}(\sqrt{C})$ regret in an intermediate setting where the losses are corrupted by a total amount of C . Our algorithm is based on the Follow-the-Regularized-Leader method from Zimin and Neu [26], with a novel hybrid regularizer inspired by recent works of Zimmert et al. [27, 29] for the special case of multi-armed bandits. Crucially, our regularizer admits a non-diagonal Hessian with a highly complicated inverse. Analyzing such a regularizer and deriving a particular self-bounding regret guarantee is our key technical contribution and might be of independent interest.

1 Introduction

We study the problem of learning episodic Markov Decision Processes (MDPs). In this problem, a learner interacts with the environment through T episodes. In each episode, the learner starts from a fixed state, then sequentially selects one of the available actions and transits to the next state according to a fixed transition function for a fixed number of steps. The learner observes only the visited states and the loss for each visited state-action pair, and her goal is to minimize her regret, the difference between her total loss over T episodes and that of the optimal fixed policy in hindsight.

When the losses are adversarial and can change arbitrarily between episodes, the state-of-the-art is achieved by the UOB-REPS algorithm of [12] with near-optimal regret $\tilde{\mathcal{O}}(\sqrt{T})$ (ignoring dependence on other parameters). On the other hand, the majority of the literature focuses on the stochastic/i.i.d. loss setting where the loss for each state-action pair follows a fixed distribution. For example, Azar et al. [5] achieve the minimax regret $\tilde{\mathcal{O}}(\sqrt{T})$ in this case. Moreover, the recent work of Simchowitz and Jamieson [23] shows the first non-asymptotic gap-dependent regret bound of order $\mathcal{O}(\log T)$ for this problem, which is considerably more favorable than the worst-case $\tilde{\mathcal{O}}(\sqrt{T})$ regret.

A natural question then arises: *is it possible to achieve the best of both worlds with one single algorithm?* In other words, can we achieve $\mathcal{O}(\log T)$ regret when the losses are stochastic, and *simultaneously* enjoy worst-case robustness with $\tilde{\mathcal{O}}(\sqrt{T})$ regret when the losses are adversarial? Considering that the existing algorithms from [12] and [23] for these two settings are drastically different, it is highly unclear whether this is possible.

In this work, we answer the question affirmatively and develop the first algorithm with such a best-of-both-worlds guarantee, under the condition that the transition function is known. We emphasize that even in the case with known transition, the problem is still highly challenging. For example,

the adversarial case was studied in [26] and still requires using the Follow-the-Regularized-Leader (FTRL) or Online Mirror Descent framework from the online learning literature (see e.g., [10]) over the *occupancy measure* space, which the UOB-REPS algorithm [12] adopts as well. This is still significantly different from the algorithms designed for the stochastic setting.

Moreover, our algorithm achieves the logarithmic regret $\mathcal{O}(\log T)$ for a much broader range of situations besides the usual stochastic setting. In fact, neither independence nor identical distributions are required, as long as a certain gap condition similar to that of [27] (for multi-armed bandits) holds (see Eq. (2)). Even more generally, our algorithm achieves $\tilde{\mathcal{O}}(\log T + \sqrt{C})$ regret in an intermediate setting where the losses are corrupted by a total amount of C . This bound smoothly interpolates between the logarithmic regret for the stochastic setting and the worst-case $\tilde{\mathcal{O}}(\sqrt{T})$ regret for the adversarial setting as C increases from 0 to T .

Techniques. Our algorithm is mainly inspired by recent advances in achieving best-of-both-worlds guarantees for the special case of multi-armed bandits or semi-bandits [24, 27, 29]. These works show that, perhaps surprisingly, such guarantees can be obtained with the standard FTRL framework originally designed only for the adversarial case. All we need is a carefully designed regularizer and a particular analysis that relies on proving a certain kind of *self-bounding regret bounds*, which then automatically implies $\mathcal{O}(\log T)$ regret for the stochastic setting, and more generally $\tilde{\mathcal{O}}(\sqrt{C})$ regret for the case with C corruption.

We greatly extend this idea to the case of learning episodic MDPs. As mentioned, Zimin and Neu [26] already solved the adversarial case using FTRL over the occupancy measure space, in particular with Shannon entropy as the regularizer in the form $\sum_{s,a} q(s,a) \ln q(s,a)$, where $q(s,a)$ is the occupancy for state s and action a . Our key algorithmic contribution is to design a new regularizer based on the $1/2$ -Tsallis-entropy used in [27]. However, we argue that using only the Tsallis entropy, in the form of $-\sum_{s,a} \sqrt{q(s,a)}$, is not enough. Instead, inspired by the work of [29] for semi-bandits, we propose to use a *hybrid* regularizer in the form $-\sum_{s,a} (\sqrt{q(s,a)} + \sqrt{q(s) - q(s,a)})$ where $q(s) = \sum_a q(s,a)$. In fact, to stabilize the algorithm, we also need to add yet another regularizer in the form $-\sum_{s,a} \log q(s,a)$ (known as log-barrier), borrowing the idea from [7, 8, 16]. See Section 2.2 and Section 3 for more detailed discussions on the design of our regularizer.

More importantly, we emphasize that analyzing our new regularizer requires significantly new ideas, mainly because it admits a *non-diagonal Hessian* with a highly complicated inverse. Indeed, the key of the FTRL analysis lies in analyzing the quadratic norm of the loss estimator with respect to the inverse Hessian of the regularizer. As far as we know, almost all regularizers used in existing FTRL methods are decomposable over coordinates and thus admit a diagonal Hessian, making the analysis relatively straightforward (with some exceptions mentioned in related work below). Our approach is the first to apply and analyze an explicit non-decomposable regularizer with non-diagonal Hessian. Our analysis heavily relies on rewriting $q(s)$ in a different way and constructing the Hessian inverse recursively (see Section 4). The way we analyze our algorithm and derive a self-bounding regret bound for MDPs is the key technical contribution of this work and might be of independent interest.

While we only resolve the problem with known transition, we believe that our approach, providing the first best-of-both-worlds result for MDPs, sheds light on how to solve the general case with unknown transition.

Related work. We refer the reader to [23] for earlier works on gap-dependent logarithmic regret bounds for learning MDPs with stochastic losses, and to [12] for earlier works on learning MDPs with adversarial losses. Using very different techniques, the recent work of Lykouris et al. [18] also develops an algorithm for the stochastic setting that is robust to a certain amount of adversarial corruption to the environment (including both the transition and the losses). Their algorithm does not ensure a worst-case bound of $\mathcal{O}(\sqrt{T})$ and can only tolerate $o(\sqrt{T})$ amount of corruption, while our algorithm ensures $\mathcal{O}(\sqrt{T})$ regret always. On the other hand, their algorithm works even under unknown transition while ours cannot.

For the special case of multi-armed bandits (essentially our setting with one single state), the question of achieving best-of-both-worlds was first considered by Bubeck and Slivkins [6]. Since then, different improvements have been proposed over the years [22, 4, 21, 24, 17, 9, 27, 29]. As mentioned, our regularizer is largely based on the most recent advances from [27, 29], briefly reviewed in Section 2.2.

As far as we know, all existing works on using FTRL with a regularizer that admits a non-diagonal Hessian do not require calculating the Hessian inverse explicitly. For example, the SCRiBLE algorithm of [1, 2] for efficient bandit linear optimization and its variants (e.g., [20, 11]) use any self-concordant barrier [19] of the decision set as the regularizer, and the entire analysis only relies on certain properties of self-concordant barriers and does not even require knowing the explicit form of the regularizer. As another example, when applying FTRL over a space of matrices, the regularizer usually also does not decompose over the entries (see e.g., [14]). However, the Hessian inverse is often well-known from matrix calculus already. These previous attempts are all very different from our analysis where we need to explicitly work out the inverse of the non-diagonal Hessian.

2 Preliminaries

The problem of learning an episodic MDP through T episodes is defined by a tuple $(S, A, P, L, \{\ell_t\}_{t=1}^T)$, where S and A are the finite state and action space respectively, $P : S \times A \times S \rightarrow [0, 1]$ is the transition function such that $P(s'|s, a)$ is the probability of moving to state s' after executing action a at state s , L is the number of interactions within each episode, and $\ell_t : S \times A \rightarrow [0, 1]$ is the loss function for episode t , specifying the loss of each state-action pair.

Without loss of generality (see detailed discussions in [12]), the MDP is assumed to have the following layered structure. First, the state space S consists of $L + 1$ layers S_0, \dots, S_L such that $S = \bigcup_{k=0}^L S_k$ and $S_i \cap S_j = \emptyset$ for $i \neq j$. Second, S_0 and S_L are singletons, containing only the start state s_0 and the terminal state s_L respectively. Third, transitions are only possible between consecutive layers. In other words, if $P(s'|s, a) > 0$, then $s' \in S_{k+1}$ and $s \in S_k$ for some k .

Ahead of time, the environment decides $(S, A, P, L, \{\ell_t\}_{t=1}^T)$, with S, A, L , and P revealed to the learner. The interaction between the learner and the environment then proceeds in T episodes. For each episode t , the learner decides a stochastic policy $\pi_t : S \times A \rightarrow [0, 1]$, where $\pi_t(a|s)$ is the probability of selecting action a at state s , then executes the policy starting from the initial state s_0 until reaching the final state s_L , yielding and observing a sequence of L state-action-loss tuples $(s_0, a_0, \ell_t(s_0, a_0)), \dots, (s_{L-1}, a_{L-1}, \ell_t(s_{L-1}, a_{L-1}))$, where action a_k is drawn from $\pi_t(\cdot|s_k)$ and state s_{k+1} is drawn from $P(\cdot|s_k, a_k)$, for each $k = 0, \dots, L - 1$. Importantly, the learner only observes the loss of the visited action-state pairs and nothing else about the loss function ℓ_t , known as the *bandit feedback* setting.

For any policy π , with slight abuse of notation we denote its expected loss in episode t as $\ell_t(\pi) = \mathbb{E}[\sum_{k=0}^{L-1} \ell_t(s_k, a_k)]$, where the state-action pairs $(s_0, a_0), \dots, (s_{L-1}, a_{L-1})$ are generated according to the transition function P and the policy π . The expected regret of the learner with respect to a policy π is defined as $\text{Reg}_T(\pi) = \mathbb{E}[\sum_{t=1}^T \ell_t(\pi_t) - \sum_{t=1}^T \ell_t(\pi^*)]$, which is the difference between the total loss of the learner and that of policy π^* . The goal of the learner is to minimize her regret with respect to an optimal policy, denoted as $\text{Reg}_T = \max_{\pi} \text{Reg}_T(\pi)$. Throughout the paper, we use $\hat{\pi} : S \rightarrow A$ to denote a *deterministic* optimal policy, which is known to always exist.

Occupancy measures. To solve the problem using techniques from online learning, we need the concept of ‘‘occupancy measures’’ proposed in [26]. Specifically, fixing the MDP of interest, every stochastic policy π induces an occupancy measure $q^\pi : S \times A \rightarrow [0, 1]$ with $q^\pi(s, a)$ being the probability of visiting state-action pair (s, a) by following policy π . With this concept, the expected loss of a policy π in episode t can then be written as a simple linear function of q^π : $\ell_t(\pi) = \sum_{s \neq s_L, a \in A} q^\pi(s, a) \ell_t(s, a)$, which we denote as $\langle q^\pi, \ell_t \rangle$, and the regret can be written as $\text{Reg}_T = \mathbb{E}[\sum_{t=1}^T \langle q_t - \hat{q}, \ell_t \rangle]$ where $q_t = q^{\pi_t}$ and $\hat{q} = q^{\hat{\pi}}$.

In other words, the problem essentially becomes an instance of online linear optimization with bandit feedback, where in each episode the learner proposes an occupancy measure $q_t \in \Omega$. Here, Ω is the set of all possible occupancy measures and is known to be a polytope satisfying two constraints [26]: first, for every $k = 0, \dots, L - 1$, $\sum_{s \in S_k} \sum_{a \in A} q(s, a) = 1$; second, for every $k = 1, \dots, L - 1$ and every state $s \in S_k$,

$$\sum_{s' \in S_{k-1}} \sum_{a' \in A} q(s', a') P(s|s', a') = \sum_a q(s, a). \quad (1)$$

Having an occupancy measure $q_t \in \Omega$, one can directly find its induced policy via $\pi_t(a|s) = q_t(s, a) / \sum_{a' \in A} q_t(s, a')$.

More notation. We use $k(s)$ to denote the index of the layer to which state s belongs, and $\mathbb{I}\{\cdot\}$ to denote the indicator function whose value is 1 if the input holds true and 0 otherwise. For a positive definite matrix $M \in \mathbb{R}^{K \times K}$, $\|x\|_M \triangleq \sqrt{x^\top M x}$ is the quadratic norm of $x \in \mathbb{R}^K$ with respect to M . Throughout the paper, we use $\tilde{\mathcal{O}}(\cdot)$ to hide terms of order $\log T$.

2.1 Conditions on loss functions

So far we have not discussed how the environment decides the loss functions ℓ_1, \dots, ℓ_T . We consider two general settings. The first one is the *adversarial setting*, where there is no assumption at all on how the losses are generated — they can even be generated in a malicious way after seeing the learner’s algorithm (but not her random seeds).¹ The O-REPS algorithm of [26] achieves $\text{Reg}_T = \tilde{\mathcal{O}}(\sqrt{L|S||A|T})$ in this case, which was shown to be optimal.

The second general setting we consider subsumes many cases such as the stochastic case. Generalizing [28] (the full version of [27]) for bandits, we propose to summarize this setting by the following condition on the losses: there exist a gap function $\Delta : S \times A \rightarrow \mathbb{R}_+$, a mapping $\pi^* : S \rightarrow A$, and a constant $C \geq 0$ such that

$$\text{Reg}_T = \mathbb{E} \left[\sum_{t=1}^T \langle q_t - \hat{q}, \ell_t \rangle \right] \geq \mathbb{E} \left[\sum_{t=1}^T \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} q_t(s, a) \Delta(s, a) \right] - C \quad (2)$$

holds for all sequences of q_1, \dots, q_T . While seemingly strong and hard to interpret, this condition in fact subsumes many existing settings as explained below.

Stochastic losses. In the standard stochastic setting studied by most works in the literature, ℓ_1, \dots, ℓ_T are i.i.d. samples of a fixed and unknown distribution. In this case, by the performance difference lemma (see e.g., [13, Lemma 5.2.1]), Eq. (2) is satisfied with $\pi^* = \hat{\pi}$, $C = 0$, and $\Delta(s, a) = Q(s, a) - \min_{a' \neq a} Q(s, a')$ where Q is the Q function of the optimal policy $\hat{\pi}$.² In fact, Eq. (2) holds with equality in this case. Here, $\Delta(s, a)$ is the sub-optimality gap of action a at state s , and plays a key role in the optimal logarithmic regret for learning MDPs as shown in [23].

Stochastic losses with corruption. More generally, consider a setting where the environment first generates ℓ'_1, \dots, ℓ'_T as i.i.d. samples of an unknown distribution (call the corresponding MDP \mathcal{M}), and then corrupt them in an arbitrary way to arrive at the final loss functions ℓ_1, \dots, ℓ_T . Then Eq. (2) is satisfied with π^* being the optimal policy of \mathcal{M} , $C = 2 \sum_{t=1}^T \sum_{k < L} \max_{s \in S_{k,a}} |\ell_t(s, a) - \ell'_t(s, a)|$ being the total amount of corruption, and $\Delta(s, a) = Q(s, a) - \min_{a' \neq a} Q(s, a')$ where Q is the Q function of policy π^* with respect to \mathcal{M} . This is because: $\text{Reg}_T \geq \sum_{t=1}^T \langle q_t - q^{\pi^*}, \ell_t \rangle = \langle q_t - q^{\pi^*}, \ell'_t \rangle + \langle q_t, \ell_t - \ell'_t \rangle - \langle q^{\pi^*}, \ell_t - \ell'_t \rangle \geq \sum_{t=1}^T \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} q_t(s, a) \Delta(s, a) - C$, where in the last step we use the performance difference lemma again and the definition of C .

Compared to the corruption model studied by [18], the setting considered here is more general in the sense that C measures the total amount of corruption in the loss functions, instead of the number of corrupted episodes as in [18]. On the other hand, they allow corrupted transition as well when an episode is corrupted, while our transition is always fixed and known.

2.2 Follow-the-Regularized-Leader and self-bounding regret

FTRL is one of the standard frameworks to derive online learning algorithms for adversarial environments. In our context, FTRL computes $q_t = \text{argmin}_{q \in \Omega} \sum_{\tau < t} \langle q, \hat{\ell}_\tau \rangle + \psi_t(q)$ where $\hat{\ell}_\tau$ is some estimator for the loss function ℓ_τ and ψ_t is a regularizer usually of the form $\psi_t = \frac{1}{\eta_t} \psi$ for some learning rate $\eta_t > 0$ and a fixed convex function ψ . The O-REPS algorithm exactly falls into this

¹Technically, this is a setting with an *oblivious* adversary, where ℓ_t does not depend on the learner’s previous actions. However, this is only for simplicity and our results generalize to adaptive adversaries directly.

²One caveat here is that since we require $\Delta(s, a)$ to be non-zero, the optimal action for each state needs to be unique for Eq. (2) to hold. The work of [28] requires this uniqueness condition for multi-armed bandits as well, which is conjectured to be only an artifact of the analysis.

framework with $\widehat{\ell}_t$ being the importance-weighted estimator (more details in Section 3) and ψ being the entropy function.

While traditionally designed for adversarial environments, somewhat surprisingly FTRL was recently shown to be able to adapt to stochastic environments with $\mathcal{O}(\log T)$ regret as well, in the context of multi-armed bandits starting from the work by Wei and Luo [24], which was later greatly improved by Zimmert and Seldin [27, 28]. This approach is conceptually extremely clean and only relies on designing a regularizer that enables a certain kind of self-bounding regret bound. We briefly review the key ideas below since our approach is largely based on extending the same idea to MDPs.

Multi-armed bandit is a special case of our setting with $L = 1$. Thus, the concept of states does not play a role, and below we write $q(s_0, a)$ as $q(a)$ for conciseness. The regularizer used in [27] is the $1/2$ -Tsallis-entropy, originally proposed in [3], and is defined as $\psi(q) = -\sum_a \sqrt{q(a)}$. With a simple learning rate schedule $\eta_t = 1/\sqrt{t}$, it was shown that FTRL ensures the following adaptive regret bound for some constant $B > 0$,

$$\text{Reg}_T \leq \mathbb{E} \left[B \sum_{t=1}^T \sum_{a \neq a^*} \sqrt{\frac{q_t(a)}{t}} \right] \quad (3)$$

where a^* can be *any* action. The claim is now that, solely based on Eq. (3), one can derive the best-of-both-worlds guarantee already, without even further considering the details of the algorithm. To see this, first note that applying Cauchy-Schwarz inequality ($\sum_a \sqrt{q_t(s)} \leq \sqrt{|A|}$) immediately leads to a worst-case robustness guarantee of $\text{Reg}_T = \mathcal{O}(\sqrt{|A|T})$ (optimal for multi-armed bandits).

More importantly, suppose now Condition (2) holds (note again that there is only one state $s = s_0$ and we write $\Delta(s_0, a)$ as $\Delta(a)$). Then picking $a^* = \pi^*(s_0)$, one can arrive at the following self-bounding regret using Eq. (3):

$$\text{Reg}_T \leq \mathbb{E} \left[\sum_{t=1}^T \sum_{a \neq a^*} \frac{q_t(a)\Delta(a)}{2z} + \frac{zB^2}{2t\Delta(a)} \right] \leq \frac{\text{Reg}_T + C}{2z} + zB^2 \sum_{a \neq a^*} \frac{\log T}{\Delta(a)},$$

where the first step uses the AM-GM inequality and holds for any $z > 1/2$, and the second step uses Eq. (2) and the fact $\sum_{t=1}^T 1/t \leq 2 \log T$. Note that we have bounded the regret in terms of itself (hence the name self-bounding). Rearranging then gives $\text{Reg}_T \leq \frac{2z^2 B^2}{2z-1} \sum_{a \neq a^*} \frac{\log T}{\Delta(a)} + \frac{C}{2z-1}$. It just remains to pick the optimal z to minimize the bound. Specifically, using the shorthand $U = \frac{1}{2} B^2 \sum_{a \neq a^*} \frac{\log T}{\Delta(a)}$ and $x = 2z-1 > 0$, the bound can be simplified to $\text{Reg}_T \leq 2U + Ux + (C+U)/x$, and finally picking the best x to balance the last two terms gives $\text{Reg}_T \leq 2U + 2\sqrt{U(C+U)} \leq 4U + 2\sqrt{UC} = 2B^2 \sum_{a \neq a^*} \frac{\log T}{\Delta(a)} + \sqrt{2CB^2 \sum_{a \neq a^*} \frac{\log T}{\Delta(a)}}$.

In the case with stochastic losses ($C = 0$), the final bound is $\mathcal{O}(\sum_{a \neq a^*} \frac{\log T}{\Delta(a)})$, exactly matching the lower bound for stochastic multi-armed bandits [15]. More generally, in the corruption model, the regret is order $\mathcal{O}(\log T + \sqrt{C \log T})$, smoothly interpolating between the bounds for the stochastic setting and the adversarial setting as C increases from 0 to T .

Finally, we remark that although not mentioned explicitly, the follow-up work [29] reveals that using a *hybrid* regularizer in the form $\psi(q) = -\sum_a (\sqrt{q(a)} + \sqrt{1-q(a)})$, with Tsallis entropy applied to both q and its complement, also leads to the same bound Eq. (3), via an even simpler analysis. This is crucial for our algorithm design and analysis as explained in the next section.

3 Algorithm and Main Results

We are now ready to present our algorithm. Based on the discussions from Section 2.2, our goal is to design an algorithm with a regret bound akin to Eq. (3):

$$\text{Reg}_T \leq \mathbb{E} \left[B \sum_{t=1}^T \sum_{s \neq s_L} \sum_{a \neq \pi(s)} \sqrt{\frac{q_t(s, a)}{t}} \right] \quad (4)$$

Algorithm 1 FTRL with hybrid Tsallis entropy for learning stochastic and adversarial MDPs

Parameters: α, β, γ

Define: hybrid regularizer ϕ_H and log-barrier regularizer ϕ_L as in Eq. (5) and Eq. (6)

Define: valid occupancy measure space Ω (see Section 2), learning rate $\eta_t = \gamma/\sqrt{t}$

Initialize: $\widehat{L}_0(s, a) = 0$ for all (s, a)

for $t = 1$ **to** T **do**

 compute $q_t = \operatorname{argmin}_{q \in \Omega} \langle q, \widehat{L}_{t-1} \rangle + \psi_t(q)$ where $\psi_t(q) = \frac{1}{\eta_t} \phi_H(q) + \phi_L(q)$

 execute policy π_t where $\pi_t(a|s) = q_t(s, a)/q_t(s)$

 observe $(s_0, a_0, \ell_t(s_0, a_0)), \dots, (s_{L-1}, a_{L-1}, \ell_t(s_{L-1}, a_{L-1}))$

 construct estimator $\widehat{\ell}_t$ such that: $\forall (s, a), \widehat{\ell}_t(s, a) = \frac{\ell_t(s, a)}{q_t(s, a)} \mathbb{I}\{s_{k(s)} = s, a_{k(s)} = a\}$

 update $\widehat{L}_t = \widehat{L}_{t-1} + \widehat{\ell}_t$

for any mapping $\pi : S \rightarrow A$. This immediately implies a worst-case regret bound $\operatorname{Reg}_T = \mathcal{O}(\sqrt{|L|S|A|T})$ (by applying Cauchy-Schwarz inequality), matching the optimal bound of [26]. Moreover, repeating the same calculation and picking $\pi = \pi^*$, one can verify that under Condition (2) this leads to a similar bound of order $\mathcal{O}(\log T + \sqrt{C \log T})$ as in the case for multi-armed bandits.

To achieve so, a natural idea is to directly extend the $1/2$ -Tsallis-entropy regularizer to MDPs and take $\psi(q) = -\sum_{s,a} \sqrt{q(s, a)}$. However, generalizing the proofs of [27] one can only prove a weaker bound: $\operatorname{Reg}_T \leq B \sum_{t=1}^T \sum_{(s,a) \neq \xi(k(s))} \sqrt{\frac{q_t(s,a)}{t}}$ for any ξ mapping from a layer index to a state-action pair. Compared to the desired Eq. (4), one can see that instead of excluding one arbitrary action $\pi(s)$ for each state s in the bound, now we only exclude one arbitrary action for one single state specified by ξ in each layer. This is not enough to derive the same results as one can verify.

The hybrid regularizer $-\sum_{s,a} (\sqrt{q(s, a)} + \sqrt{1 - q(s, a)})$ [29] suffers the same issue. However, we propose a natural fix to this hybrid version by replacing 1 with $q(s) \triangleq \sum_a q(s, a)$, that is, the marginal probability of visiting state s under occupancy measure q .³ More concretely, we define (with “H” standing for “hybrid”)

$$\phi_H(q) = - \sum_{s \neq s_L, a \in A} \left(\sqrt{q(s, a)} + \alpha \sqrt{q(s) - q(s, a)} \right), \quad \text{where } q(s) = \sum_{a \in A} q(s, a) \quad (5)$$

for some parameter $\alpha > 0$, as the key component of our regularizer. Note that in the case of multi-armed bandits, this exactly recovers the original hybrid regularizer since $q(s_0) = 1$ and there is only one state. For MDPs, intuitively each state is dealing with a multi-armed bandit instance, but with total available probability $q(s)$ instead of 1, making ϕ_H a natural choice. However, also note another important distinction between ϕ_H and the ones discussed earlier: ϕ_H does not decompose over the action-state pairs, thus admitting a *non-diagonal Hessian*. This makes the analysis highly challenging since standard FTRL analysis requires analyzing the Hessian inverse of the regularizer. We will come back to this challenge in Section 4.

To further stabilize the algorithm and make sure that q_t and q_{t+1} are not too different (another important requirement of typical FTRL analysis), we apply another regularizer in addition to ϕ_H , defined as (with “L” standing for “log-barrier”):

$$\phi_L(q) = \beta \sum_{s \neq s_L, a \in A} \log \frac{1}{q(s, a)}, \quad (6)$$

for some parameter $\beta > 0$. Our final regularizer for time t is then $\psi_t(q) = \frac{1}{\eta_t} \phi_H(q) + \phi_L(q)$ where $\eta_t = \gamma/\sqrt{t}$ is a decreasing learning rate with parameter $\gamma > 0$. In all our results we pick $\beta = \mathcal{O}(L)$. Thus, the weight for ϕ_L is much smaller than that for ϕ_H . This idea of adding a small amount of log-barrier to stabilize the algorithm was first used in [7] and recently applied in several other works [8, 25, 16]. See more discussions in Section 4.

³We find it intriguing that while the Tsallis entropy regularizer and its hybrid version work equally well for multi-armed bandits, only the latter admits a natural way to be generalized to learning MDPs.

Our final algorithm is shown in [Algorithm 1](#). In each episode t , the algorithm follows standard FTRL and computes $q_t = \operatorname{argmin}_{q \in \Omega} \langle q, \widehat{L}_{t-1} \rangle + \psi_t(q)$ with $\widehat{L}_{t-1} = \sum_{s < t} \widehat{\ell}_s$ being the accumulated estimated loss. Then the policy π_t induced from q_t is executed, generating a sequence of state-action-loss tuples. A standard importance-weighted unbiased estimator $\widehat{\ell}_t$ is then constructed with $\widehat{\ell}_t(s, a)$ being the actual loss $\ell_t(s, a)$ divided by $q_t(s, a)$ if the state-action pair (s, a) was visited in this episode, and zero otherwise. Also note that [Algorithm 1](#) can be efficiently implemented since the key FTRL step is a convex optimization problem with $\mathcal{O}(L + |S||A|)$ linear constraints (solving it to an inaccuracy of $\mathcal{O}(1/T)$ is enough clearly).

3.1 Main Results

We move on to present the regret guarantees of our algorithm. As mentioned, the goal is to show [Eq. \(4\)](#), and the theorem below shows that our algorithm essentially achieves this (see [Appendix A](#) for the proof).

Theorem 1. *With $\alpha = 1/\sqrt{|A|}$, $\beta = 64L$, and $\gamma = 1$, [Algorithm 1](#) ensures that Reg_T is bounded by*

$$\sum_{t=1}^T \widetilde{\mathcal{O}} \left(\min \left\{ \mathbb{E} \left[B \sum_{s \neq s_L} \sum_{a \neq \pi(s)} \sqrt{\frac{q_t(s, a)}{t}} + D \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} \frac{q_t(s, a) + \check{q}(s, a)}{t}} \right], \frac{D}{\sqrt{t}} \right\} \right) \quad (7)$$

for any mapping $\pi : S \rightarrow A$, where $B = L^{\frac{5}{2}} + L\sqrt{|A|}$ and $D = \sqrt{L|S||A|}$.

Looking at the first term of the min operator, one sees that we have an extra term compared to the ideal bound [Eq. \(4\)](#). However, this only contributes to small terms in the final bounds as we explain below.⁴ Based on this theorem, we next present more concrete regret bounds for our algorithm. First, consider the adversarial setting where there is no further structure in the losses. Simply taking the second term of the min operator in [Eq. \(7\)](#) we obtain the following corollary.

Corollary 2. *With $\alpha = 1/\sqrt{|A|}$, $\beta = 64L$, and $\gamma = 1$, the regret of [Algorithm 1](#) is always bounded as*

$$\operatorname{Reg}_T \leq \widetilde{\mathcal{O}} \left(\sqrt{L|S||A|T} \right).$$

Again, this bound matches that of the O-REPS algorithm [\[26\]](#) and is known to be optimal. Note that using the first term of the min operator we can also derive an $\widetilde{\mathcal{O}}(\sqrt{T})$ bound, but the dependence on other parameters would be larger.

On the other hand, under [Condition Eq. \(2\)](#), we obtain the following bound.

Corollary 3. *Suppose [Condition \(2\)](#) holds. Then with $\alpha = 1/\sqrt{|A|}$, $\beta = 64L$, and $\gamma = 1$, the regret of [Algorithm 1](#) is bounded as $\operatorname{Reg}_T \leq \mathcal{O}(U + \sqrt{UC})$ where*

$$U = \frac{L|S||A| \log T}{\Delta_{\min}} + L^2 (L^3 + |A|) \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{\log T}{\Delta(s, a)} = \mathcal{O}(\log T)$$

and $\Delta_{\min} = \min_{s \neq s_L, a \neq \pi^*(s)} \Delta(s, a)$ is the minimal gap. Consequently, in the stochastic setting, we have $\operatorname{Reg}_T = \mathcal{O}(U)$.

See [Appendix A.1](#) for the proof, whose idea is similar to the discussions in [Section 2.2](#). Note that we are getting an extra term in U involving $1/\Delta_{\min}$, which in turn comes from the extra term in [Eq. \(7\)](#) mentioned earlier. For the stochastic setting, the best existing logarithmic bound is $\mathcal{O}(L^3 \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{\log T}{\Delta(s, a)})$ from [\[23\]](#) (which also has a dependence on $\frac{1}{\Delta_{\min}}$ but is ignored here for simplicity). Our bound has a worse factor $L^2(L^3 + |A|)$. We note that by tuning the parameters α and γ differently, one can obtain a better bound in this case. This set of parameters, however, leads to a sub-optimal adversarial bound. See [Appendix A.2](#) for details. Since our goal is to develop one single algorithm that adapts to different environments automatically, we stick to the same set of parameters in the theorem and corollaries above.

⁴The fact that this form of weaker bounds is also sufficient for the self-bounding property might be of interest for other bandit problems as well.

As mentioned, in the corruption setting, our bound $\tilde{\mathcal{O}}(\sqrt{C})$ smoothly interpolates between the logarithmic regret in the no corruption case and the $\tilde{\mathcal{O}}(\sqrt{T})$ worst-case regret in the adversarial case. On the other hand, the bound from [18] is of order $\tilde{\mathcal{O}}(C^2)$, only meaningful for $C = o(\sqrt{T})$.

4 Analysis Sketch

Analyzing our algorithm requires several new ideas, which we briefly mention in this section through a few steps, with all details deferred to the appendix.

First step. While we define $q(s)$ as $\sum_a q(s, a)$, the entire analysis relies on using an equivalent definition of $q(s)$ based on Eq. (1). The benefit of this alternative definition is that it contains the information of the transition function P and implicitly introduces a layered structure to the regularizer, which is important for constructing the Hessian and its inverse recursively. We emphasize that, however, this does *not* change the algorithm at all, because all occupancy measures in Ω ensure Eq. (1) by definition and the FTRL optimization over Ω is thus not affected.

Second step. Following the standard FTRL analysis one can obtain a regret bound in terms of $\|\widehat{\ell}_t\|_{\nabla^{-2}\phi_H(q_t)}$ for some q'_t between q_t and q_{t+1} . Moving from q'_t to q_t is exactly the part where the log-barrier ϕ_L is important. Specifically, following the idea of [16, Lemma 9], we prove that q_t and q_{t+1} are sufficiently close, and consequently Reg_T is mainly bounded by two terms: the penalty term $\sum_{t=1}^T (1/\eta_t - 1/\eta_{t-1}) (\phi_H(\hat{q}) - \phi_H(q_t))$ and the stability term $\sum_{t=1}^T \eta_t \|\widehat{\ell}_t\|_{\nabla^{-2}\phi_H(q_t)}^2$ (see Lemma 5).

Third step. Bounding the penalty term already requires a significant departure from the analysis for multi-armed bandits. Specifically, $\phi_H(\hat{q}) - \phi_H(q_t)$ can be written as

$$\sum_{s \neq s_L} \sqrt{q_t(s)} (h_s(\pi_t) - h_s(\hat{\pi})) + \sum_{s \neq s_L} \left(\sqrt{q_t(s)} - \sqrt{\hat{q}(s)} \right) h_s(\hat{\pi}) \quad (8)$$

where $h_s(\pi) = \sum_{a \in A} \sqrt{\pi(a|s)} + \alpha \sqrt{1 - \pi(a|s)}$ is basically the hybrid regularizer for multi-armed bandits (at state s) mentioned in Section 2.2. The first term in Eq. (8) can then be bounded as $(1 + \alpha) \sum_s \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)}$ for any mapping $\pi : S \rightarrow A$, in a similar way as in the multi-armed bandit analysis. However, the key difficulty is the second term, which does not appear for multi-armed bandits where there is only one state s_0 with $q(s_0) = 1$ for all q . This is highly challenging especially because we would like to arrive at a term with a summation over $a \neq \pi(s)$ as in the first term. Our main idea is to separately bound $\sqrt{q_t(s)} - \sqrt{q^\pi(s)}$ and $\sqrt{q^\pi(s)} - \sqrt{\hat{q}(s)}$ via a key induction lemma (Lemma 18) that connects them to similar terms in previous layers. We remark that this term is the source of the extra term $\sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} (q_t(s, a) + \hat{q}(s, a))} / t$ in our main regret bound Eq. (7). The complete proof for bounding the penalty term is in Appendix C.

Fourth step. Analyzing the stability term is yet another highly challenging part, since it requires working out the Hessian inverse $\nabla^{-2}\phi_H(q_t)$. The high-level idea of our proof is to first write the Hessian in a recursive form based on the layered structure introduced by writing $q(s)$ differently as mentioned in the first step. Then we apply Woodbury matrix identity to obtain a recursive form of the Hessian inverse. Finally, we argue that only certain parts of the Hessian inverse matter, and these parts enjoy nice properties allowing us to eventually bound $\mathbb{E} \left[\|\widehat{\ell}_t\|_{\nabla^{-2}\phi_H(q_t)}^2 \right]$ by $8eL^2 \left(\sqrt{L} + 1/\alpha L \right) \sum_{s \neq s_L} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)}$, again, for any mapping π . See Appendix D for the complete proof.

5 Conclusion

In this work, we provide the first algorithm for learning episodic MDPs that automatically adapts to different environments with favorable guarantees. Our algorithm applies a natural regularizer with a complicated non-diagonal Hessian to the FTRL framework, and our analysis for obtaining

a self-bounding regret bound requires several novel ideas. Apart from improving our bound in [Corollary 3](#), one key future direction is to remove the known transition assumption, which appears to demand new techniques since it is hard to also control the bias introduced by estimating the transition in terms of the adaptive bound in [Eq. \(7\)](#).

Broader Impact

This work is mostly theoretical, with no negative outcomes. Researchers working on theoretical aspects of online learning, bandit problems, and reinforcement learning (RL) may benefit from our results. Although our algorithm deals with the tabular setting and is not directly applicable to common RL applications with a large state and action space, it sheds light on how to increase robustness of a learning algorithm while adapting to specific instances, and serves as an important step towards developing more practical, adaptive, and robust RL algorithms, which in the long run might find its applications in the real world.

Acknowledgments and Disclosure of Funding

We thank Max Simchowitz for many helpful discussions. HL is supported by NSF Awards IIS-1755781 and IIS-1943607, and a Google Faculty Research Award.

References

- [1] Jacob D Abernethy, Elad Hazan, and Alexander Rakhlin. Competing in the dark: An efficient algorithm for bandit linear optimization. In *Proceedings of the Annual Conference on Learning Theory*, 2008.
- [2] Jacob D Abernethy, Elad Hazan, and Alexander Rakhlin. Interior-point methods for full-information and bandit online learning. *IEEE Transactions on Information Theory*, 2012.
- [3] Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits. In *Proceedings of the Annual Conference on Learning Theory*, 2009.
- [4] Peter Auer and Chao-Kai Chiang. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Proceedings of the Annual Conference on Learning Theory*, 2016.
- [5] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, 2017.
- [6] Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *Proceedings of the 23rd Annual Conference on Learning Theory*, 2012.
- [7] Sébastien Bubeck, Michael Cohen, and Yuanzhi Li. Sparsity, variance and curvature in multi-armed bandits. In *Algorithmic Learning Theory*, 2018.
- [8] Sébastien Bubeck, Yuanzhi Li, Haipeng Luo, and Chen-Yu Wei. Improved path-length regret bounds for bandits. In *Conference On Learning Theory*, 2019.
- [9] Anupam Gupta, Tomer Koren, and Kunal Talwar. Better algorithms for stochastic bandits with adversarial corruptions. In *Proceedings of the Annual Conference on Learning Theory*, 2019.
- [10] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [11] Elad Hazan and Kfir Levy. Bandit convex optimization: Towards tight bounds. In *Advances in Neural Information Processing Systems*, 2014.
- [12] Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu. Learning adversarial Markov decision processes with bandit feedback and unknown transition. In *International Conference on Machine Learning*, 2020.

- [13] Sham Machandranath Kakade. *On the sample complexity of reinforcement learning*. PhD thesis, University College London, 2003.
- [14] Wojciech Kotłowski and Gergely Neu. Bandit principal component analysis. In *Conference On Learning Theory*, 2019.
- [15] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [16] 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*, 2020.
- [17] Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *Proceedings of the Annual ACM SIGACT Symposium on Theory of Computing*, 2018.
- [18] Thodoris Lykouris, Max Simchowitz, Aleksandrs Slivkins, and Wen Sun. Corruption robust exploration in episodic reinforcement learning. *arXiv preprint arXiv:1911.08689*, 2019.
- [19] Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*. Siam, 1994.
- [20] Ankan Saha and Ambuj Tewari. Improved regret guarantees for online smooth convex optimization with bandit feedback. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2011.
- [21] Yevgeny Seldin and Gábor Lugosi. An improved parametrization and analysis of the exp3++ algorithm for stochastic and adversarial bandits. In *Proceedings of the Annual Conference on Learning Theory*, 2017.
- [22] Yevgeny Seldin and Aleksandrs Slivkins. One practical algorithm for both stochastic and adversarial bandits. In *Proceedings of the International Conference on Machine Learning*, 2014.
- [23] Max Simchowitz and Kevin G Jamieson. Non-asymptotic gap-dependent regret bounds for tabular MDPs. In *Advances in Neural Information Processing Systems*, pages 1151–1160, 2019.
- [24] Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. In *Proceedings of the Annual Conference On Learning Theory*, 2018.
- [25] Kai Zheng, Haipeng Luo, Ilias Diakonikolas, and Liwei Wang. Equipping experts/bandits with long-term memory. In *Advances in Neural Information Processing Systems*, pages 5927–5937, 2019.
- [26] Alexander Zimin and Gergely Neu. Online learning in episodic markovian decision processes by relative entropy policy search. In *Proceedings of the International Conference on Neural Information Processing Systems*, 2013.
- [27] Julian Zimmert and Yevgeny Seldin. An optimal algorithm for stochastic and adversarial bandits. In *The International Conference on Artificial Intelligence and Statistics*, 2019.
- [28] Julian Zimmert and Yevgeny Seldin. An optimal algorithm for stochastic and adversarial bandits. *arXiv preprint arXiv:1807.07623*, 2020.
- [29] Julian Zimmert, Haipeng Luo, and Chen-Yu Wei. Beating stochastic and adversarial semi-bandits optimally and simultaneously. In *Proceedings of the International Conference on Machine Learning*, 2019.

A Proof of Theorem 1 and Corollary 3

In this section, we provide the proof (outline) of [Theorem 1](#), the proof of [Corollary 3](#) ([Appendix A.1](#)), discussions on parameter tuning ([Appendix A.2](#)), and also some preliminaries on the Hessian of our regularizer which are useful for the rest of the appendix ([Appendix A.3](#)).

We prove the following version of [Theorem 1](#) with general value of the parameters α and γ , which facilitates further discussions on parameter tuning. Taking $\alpha = 1/\sqrt{|A|}$ and $\gamma = 1$ clearly recovers [Theorem 1](#).

Theorem 4. *With $\beta = 64L$, [Algorithm 1](#) guarantees:*

$$\text{Reg}_T = \mathcal{O} \left(\sum_{t=1}^T \min \left\{ \mathbb{E} \left[X \sum_{s \neq s_L} \sum_{a \neq \pi(s)} \sqrt{\frac{q_t(s, a)}{t}} + Y \sqrt{\frac{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a) + \dot{q}(s, a)}{t}} \right], \frac{Z}{\sqrt{t}} \right\} \right) + \mathcal{O}(L|S||A| \log T)$$

for any mapping $\pi : S \rightarrow A$, where coefficients X , Y and Z are defined as:

$$X = \frac{1 + \alpha}{\gamma} + \gamma L^2 \left(\sqrt{L} + \frac{1}{\alpha L} \right), Y = \frac{\sqrt{L|S|} (1 + \alpha|A|)}{\gamma}, Z = \sqrt{L|S||A|} \left(\frac{1 + \alpha\sqrt{|A|}}{\gamma} + \gamma \right).$$

The proof of this theorem relies on the following three important lemmas, which respectively correspond to the second, third, and fourth step of the analysis sketch in [Section 4](#) and are proven in [Appendix B](#), [Appendix C](#), and [Appendix D](#). The first one decomposes the regret into two terms (penalty and stability), with an additional small term of order $\mathcal{O}(L|S||A| \log T)$. The second one bounds the penalty term, while the third one bounds the stability term.

Lemma 5 (Regret decomposition). *With $\beta = 64L$, [Algorithm 1](#) ensures:*

$$\text{Reg}_T \leq \underbrace{\sum_{t=1}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \mathbb{E} [\phi_H(\hat{q}) - \phi_H(q_t)]}_{\text{penalty}} + 8 \underbrace{\sum_{t=1}^T \eta_t \mathbb{E} \left[\|\hat{\ell}_t\|_{\nabla^{-2}\phi_H(q_t)}^2 \right]}_{\text{stability}} + \mathcal{O}(L|S||A| \log T).$$

Lemma 6 (Penalty). *The hybrid regularizer ϕ_H defined in [Eq. \(5\)](#) ensures that $\phi_H(\hat{q}) - \phi_H(q_t)$ is bounded by*

$$(1 + \alpha) \sum_{s \neq s_L} \sum_{a \neq \pi(a)} \sqrt{q_t(s, a)} + (1 + \alpha|A|) \sqrt{|S|L} \min \left\{ 1, 2 \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a) + \dot{q}(s, a)} \right\},$$

for all $t = 1, \dots, T$, where π can be any mapping from S to A .

Lemma 7 (Stability). *[Algorithm 1](#) guarantees that $\mathbb{E} \left[\|\hat{\ell}_t\|_{\nabla^{-2}\phi_H(q_t)}^2 \right]$ is bounded by*

$$\min \left\{ 4\sqrt{L|S||A|}, \mathbb{E} \left[8eL^2 \left(\sqrt{L} + \frac{1}{\alpha \cdot L} \right) \sum_{s \neq s_L} \sum_{a \neq \pi(a)} \sqrt{q_t(s, a)} \right] \right\},$$

for all $t = 1, \dots, T$, where π can be any mapping from S to A .

Proof of [Theorem 4](#). For notational convenience, we denote

$$J_1(t) = \sqrt{\frac{1}{t}}, J_2(t) = \mathbb{E} \left[\sum_{s \neq s_L} \sum_{a \neq \pi(a)} \sqrt{\frac{q_t(s, a)}{t}} \right], J_3(t) = \mathbb{E} \left[\sqrt{\frac{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a) + \dot{q}(s, a)}{t}} \right].$$

By [Lemma 6](#) and the fact $1/\eta_t - 1/\eta_{t-1} = (\sqrt{t} - \sqrt{t-1})/\gamma = \frac{1}{\gamma(\sqrt{t} + \sqrt{t-1})} \leq 1/(\gamma\sqrt{t})$, the penalty term can be bounded by

$$\mathcal{O} \left(\sum_{t=1}^T \mathbb{E} \left[C_1 \sum_{s \neq s_L} \sum_{a \neq \pi(a)} \sqrt{\frac{q_t(s, a)}{t}} \right] + C_2 \min \left\{ \frac{1}{\sqrt{t}}, \mathbb{E} \left[\sqrt{\frac{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a) + \dot{q}(s, a)}{t}} \right] \right\} \right),$$

that is, $\mathcal{O} \left(\sum_{t=1}^T C_1 J_2(t) + C_2 \min \{J_1(t), J_3(t)\} \right)$, where $C_1 = \frac{1+\alpha}{\gamma}$ and $C_2 = \frac{(1+\alpha|A|)\sqrt{|S|L}}{\gamma}$.

On the other hand, by [Lemma 7](#), the stability term is bounded by

$$\mathcal{O} \left(\sum_{t=1}^T \min \left\{ \frac{C_3}{\sqrt{t}}, \mathbb{E} \left[C_4 \sum_{s \neq s_L} \sum_{a \neq \pi(a)} \sqrt{\frac{q_t(s, a)}{t}} \right] \right\} \right),$$

that is, $\mathcal{O} \left(\sum_{t=1}^T \min \{C_3 J_1(t), C_4 J_2(t)\} \right)$, where $C_3 = \gamma\sqrt{L|S||A|}$ and $C_4 = \gamma L^2 (\sqrt{L} + 1/\alpha L)$.

Finally, we plug these bounds into [Lemma 5](#) and show that

$$\begin{aligned} \text{Reg}_T &= \mathcal{O} \left(\sum_{t=1}^T C_1 J_2(t) + C_2 \min \{J_1(t), J_3(t)\} + \min \{C_3 J_1(t), C_4 J_2(t)\} \right) \\ &\quad + \mathcal{O}(L|S||A| \log T). \end{aligned}$$

Noticing that $J_2(t) = \sum_{s \neq s_L} \sum_{a \neq \pi(a)} \sqrt{\frac{q_t(s, a)}{t}} \leq \sqrt{\frac{L|S||A|}{t}} = \sqrt{L|S||A|} J_1(t)$ by Cauchy-Schwarz inequality, we further have

$$\begin{aligned} \text{Reg}_T &= \mathcal{O} \left(\sum_{t=1}^T C_1 J_2(t) + C_2 \min \{J_1(t), J_3(t)\} + \min \{C_3 J_1(t), C_4 J_2(t)\} \right) \\ &\quad + \mathcal{O}(L|S||A| \log T) \\ &= \mathcal{O} \left(\sum_{t=1}^T \min \{C_1 \sqrt{L|S||A|} J_1(t), C_1 J_2(t)\} + \min \{(C_2 + C_3) J_1(t), C_2 J_3(t) + C_4 J_2(t)\} \right) \\ &\quad + \mathcal{O}(L|S||A| \log T) \\ &= \mathcal{O} \left(\sum_{t=1}^T \min \left\{ (C_1 \sqrt{L|S||A|} + C_2 + C_3) J_1(t), C_2 J_3(t) + (C_1 + C_4) J_2(t) \right\} \right) \\ &\quad + \mathcal{O}(L|S||A| \log T). \end{aligned}$$

Therefore, we prove the regret bound stated in [Theorem 4](#) with $X = C_1 + C_4$, $Y = C_2$, and $Z = C_1 \sqrt{L|S||A|} + C_2 + C_3 = \mathcal{O} \left(\sqrt{L|S||A|} \left(\frac{1+\alpha\sqrt{|A|}}{\gamma} + \gamma \right) \right)$. In particular, setting $\gamma = 1$ and $\alpha = 1/\sqrt{|A|}$ exactly leads to [Theorem 1](#). \square

A.1 Proof of [Corollary 3](#)

The proof mostly follows the discussions in [Section 2.2](#), except that we need to deal with the extra term involving \dot{q} . To do that, we first introduce the following important implication of [Condition \(2\)](#).

Lemma 8. *Suppose [Condition \(2\)](#) holds. Then the optimal occupancy measure \dot{q} ensures*

$$T \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \dot{q}(s, a) \Delta(s, a) \leq C.$$

Proof. Simply setting $q_1 = \dots = q_T = \hat{q}$ in Condition (2) gives

$$0 = \mathbb{E} \left[\sum_{t=1}^T \langle \hat{q} - \hat{q}, \ell_t \rangle \right] \geq T \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \hat{q}(s, a) \Delta(s, a) - C,$$

and rearranging finishes the proof. \square

Proof of Corollary 3. By Theorem 4 (setting $\alpha = 1/\sqrt{|A|}$ and $\gamma = 1$ as stated in Corollary 3 and picking $\pi = \pi^*$), Reg_T is bounded by

$$\kappa \left(L|S||A| \log T + \sum_{t=1}^T \mathbb{E} \left[B \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \sqrt{\frac{q_t(s, a)}{t}} + D \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{q_t(s, a) + \hat{q}(s, a)}{t}} \right] \right)$$

where κ is a constant, $B = L^{\frac{5}{2}} + L\sqrt{|A|}$ and $D = \sqrt{L|S||A|}$.

For any $z > 0$, we have

$$\begin{aligned} & \kappa \sum_{t=1}^T \mathbb{E} \left[B \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \sqrt{\frac{q_t(s, a)}{t}} \right] \\ & \leq \sum_{t=1}^T \mathbb{E} \left[\sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \sqrt{\frac{q_t(s, a) \Delta(s, a)}{z} \cdot \frac{z \kappa^2 B^2}{t \Delta(s, a)}} \right] \\ & \leq \mathbb{E} \left[\sum_{t=1}^T \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \left(\frac{q_t(s, a) \Delta(s, a)}{2z} + \frac{z \kappa^2 B^2}{2t \Delta(s, a)} \right) \right] \\ & \leq \frac{\text{Reg}_T + C}{2z} + z \kappa^2 L^2 (L^3 + |A|) \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{\log T}{\Delta(s, a)} \end{aligned}$$

where the third line uses the AM-GM inequality, and the last line uses Eq. (2) and the fact $\sum_{t=1}^T 1/t \leq 2 \log T$.

For the other term, we have for any $z > 0$:

$$\begin{aligned} & \kappa \sum_{t=1}^T \mathbb{E} \left[D \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{q_t(s, a) + \hat{q}(s, a)}{t}} \right] \\ & = \sum_{t=1}^T \mathbb{E} \left[\sqrt{\left(\sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{(q_t(s, a) + \hat{q}(s, a)) \Delta_{\text{MIN}}}{z} \right) \cdot \frac{\kappa^2 D^2}{t \Delta_{\text{MIN}}}} \right] \\ & \leq \mathbb{E} \left[\sum_{t=1}^T \left(\sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{(q_t(s, a) + \hat{q}(s, a)) \Delta_{\text{MIN}}}{2z} \right) + \frac{z \kappa^2 D^2}{2t \Delta_{\text{MIN}}} \right] \\ & \leq \mathbb{E} \left[\sum_{t=1}^T \left(\sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{(q_t(s, a) + \hat{q}(s, a)) \Delta(s, a)}{2z} \right) + \frac{z \kappa^2 D^2}{2t \Delta_{\text{MIN}}} \right] \\ & \leq \frac{\text{Reg}_T + 2C}{2z} + \frac{z \kappa^2 L |S| |A| \log T}{\Delta_{\text{MIN}}} \end{aligned}$$

where the third line uses the AM-GM inequality again, the third line uses the definition of Δ_{MIN} , and the last line uses Eq. (2), Lemma 8, and the fact $\sum_{t=1}^T 1/t \leq 2 \log T$.

Combining the inequalities, we have

$$\text{Reg}_T \leq \frac{\text{Reg}_T}{z} + \frac{2C}{z} + z \kappa^2 U + \kappa V$$

where we use the shorthand U (already defined in the statement of [Corollary 3](#)) and V as

$$U = L^2(L^3 + |A|) \sum_{s \neq s_L} \sum_{a \neq \pi^*(s)} \frac{\log T}{\Delta(s, a)} + \frac{L|S||A| \log T}{\Delta_{\min}}, \quad V = L|S||A| \log T.$$

For any $z > 1$, we can further rearrange and arrive at

$$\begin{aligned} \text{Reg}_T &\leq \frac{2}{(z-1)}C + \frac{z^2}{z-1}\kappa^2U + \frac{z}{z-1}\kappa V \\ &= \frac{2}{x}C + \frac{(x+1)^2}{x}\kappa^2U + \frac{x+1}{x}\kappa V \\ &= \frac{1}{x}(2C + \kappa V + \kappa^2U) + x(\kappa^2U) + (2\kappa^2U + \kappa V) \end{aligned}$$

where we define $x = z - 1 > 0$ and replace all z 's in the second line. Picking the optimal x to balance the first terms gives

$$\begin{aligned} \text{Reg}_T &\leq 2\sqrt{(2C + \kappa V + \kappa^2U)(\kappa^2U)} + 2\kappa^2U + \kappa V \\ &\leq 2\kappa\sqrt{2UC} + 2\sqrt{\kappa^3UV} + 4\kappa^2U + \kappa V \\ &\leq 2\kappa\sqrt{2UC} + \sqrt{\kappa^3}(U + V) + 4\kappa^2U + \kappa V \\ &\leq \mathcal{O}\left(U + V + \sqrt{UC}\right) \end{aligned}$$

where the second line follows from the fact $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, and the third line uses AM-GM inequality. Finally, noticing that $V \leq \frac{L|S||A| \log T}{\Delta_{\min}} \leq U$ finishes the proof. \square

A.2 Different tuning for the stochastic case

Here, we consider the case with stochastic losses (or more generally the case where [Condition \(2\)](#) holds with $C = 0$), and point out what the best bound one can get by tuning α and γ optimally. For simplicity, we consider the worst case when $\Delta(s, a) = \Delta_{\min}$ holds for all state-action pairs (s, a) with $a \neq \pi^*(s)$. Repeating the same argument in the proof of [Corollary 3](#) with the general bound from [Theorem 4](#), one can verify that the final bound is

$$\begin{aligned} \text{Reg}_T &\leq \mathcal{O}\left(\left[\frac{1+\alpha}{\gamma} + \gamma L^2\left(\sqrt{L} + \frac{1}{\alpha L}\right)\right]^2 \frac{|S||A| \log T}{\Delta_{\min}} + \left(\frac{(1+\alpha|A|)\sqrt{|S|L}}{\gamma}\right)^2 \frac{\log T}{\Delta_{\min}}\right) \\ &\quad + \mathcal{O}(L|S||A| \log T). \end{aligned}$$

Picking the optimal parameters leads to

$$\text{Reg}_T \leq \mathcal{O}\left(|S|\sqrt{|A|}L^{3/2}\left(\sqrt{|A|}L + L^{3/2} + |A| + A^{3/4}\sqrt{L}\right)\frac{\log T}{\Delta_{\min}} + L|S||A| \log T\right).$$

This is better than the bound stated in [Corollary 3](#), and could even be better than the bound $\mathcal{O}\left(\frac{L^3|S||A| \log T}{\Delta_{\min}}\right)$ achieved by StrongEuler [\[23\]](#) (although they consider a harder setting where the transition function is unknown). However, this set of parameters leads to a sub-optimal bound for the adversarial case unfortunately.

A.3 The Hessian of ϕ_H

We calculate the Hessian of our hybrid regularizer ϕ_H in this section, which is important for the analysis in [Appendix B](#) and [Appendix D](#). As mentioned in [Section 4](#) (first step), one important trick we do is to use a different but equivalent definition of $q(s)$. Specifically, we analyze the following regularizer:

$$\begin{aligned} \phi_H(q) &= - \sum_{s \neq s_L, a \in A} \left(\sqrt{q(s, a)} + \alpha\sqrt{q(s) - q(s, a)}\right), \\ \text{where } q(s) &= \begin{cases} \sum_{s' \in S_{k(s)-1}} \sum_{a' \in A} q(s', a')P(s|s', a'), & \text{if } k(s) \neq 0, \\ 1, & \text{else.} \end{cases} \end{aligned} \tag{9}$$

We emphasize again that, within the feasible set Ω , this definition is exactly equivalent to Eq. (5), and thus we are not changing the algorithm at all.

Lemma 9. *The Hessian of the regularizer $\phi_H(q)$ defined in Eq. (9) is specified by the following:*

- for any $k = 1, \dots, L - 1$, $s' \in S_{k-1}$, $s \in S_k$, and $a, a' \in A$,

$$\frac{\partial^2 \phi_H}{\partial q(s', a') \partial q(s, a)} = \frac{-\alpha P(s|s', a')}{4(q(s) - q(s, a))^{3/2}}; \quad (10)$$

- for any $k = 1, \dots, L - 1$, $s, s' \in S_{k-1}$, and $(s, a) \neq (s', a')$,

$$\frac{\partial^2 \phi_H}{\partial q(s', a') \partial q(s, a)} = \sum_{s'' \in S_k} \sum_{a'' \in A} \frac{\alpha P(s''|s, a) P(s''|s', a')}{4(q(s'') - q(s'', a''))^{3/2}}; \quad (11)$$

- for any $k = 1, \dots, L$, $s \in S_{k-1}$, and $a \in A$,

$$\frac{\partial^2 \phi_H}{\partial q(s, a)^2} = \frac{1}{4q(s, a)^{3/2}} + \frac{\alpha}{4(q(s) - q(s, a))^{3/2}} + \sum_{s' \in S_k, s' \neq s_L} \sum_{a' \in A} \frac{\alpha P(s'|s, a)^2}{4(q(s') - q(s', a'))^{3/2}}; \quad (12)$$

- all other entries of the Hessian are 0.

Moreover, for any $w : (S \setminus \{s_L\}) \times A \rightarrow \mathbb{R}$, the quadratic form $w^\top \nabla^2 \phi_H(q) w$ can be written as

$$\frac{1}{4} \sum_{s \neq s_L} \sum_a \left(\frac{w(s, a)^2}{q(s, a)^{3/2}} + \frac{\alpha (h(s) - w(s, a))^2}{(q(s) - q(s, a))^{3/2}} \right) \quad (13)$$

where $h(s) = \begin{cases} \sum_{s' \in S_{k(s)-1}} \sum_{a' \in A} P(s|s', a') w(s', a'), & \text{if } k(s) > 0, \\ 0, & \text{else.} \end{cases}$

Consequently, ϕ_H is convex in q .

Proof. Fix a state-action pair (s, a) . By direct calculation, we show that the terms in $\phi_H(q)$ containing the variable $q(s, a)$ are

$$- \left(\sqrt{q(s, a)} + \alpha \sqrt{q(s) - q(s, a)} \right) - \alpha \sum_{s'' \in S_{k(s)+1}} \sum_{a'' \in A} \sqrt{q(s'') - q(s'', a'')} \quad (14)$$

where the last term is zero when $s = s_L$. From Eq. (14), we can infer that the second-order partial derivatives of $q(s, a)$ and $q(s', a')$ are non-zero if and only if $|k(s) - k(s')| = 1$.

We first verify Eq. (10), where $k(s') = k(s) - 1$ and the derivatives are from $-\alpha \sqrt{q(s) - q(s, a)}$. Direct calculations shows

$$\begin{aligned} \frac{\partial^2 \phi_H}{\partial q(s', a') \partial q(s, a)} &= \frac{\partial^2}{\partial q(s', a') \partial q(s, a)} \left(-\alpha \sqrt{q(s) - q(s, a)} \right) \\ &= \frac{\partial}{\partial q(s, a)} \left(\frac{-\alpha P(s|s', a')}{2\sqrt{q(s) - q(s, a)}} \right) \\ &= \frac{-\alpha P(s|s', a')}{4(q(s) - q(s, a))^{3/2}} \end{aligned}$$

where the second step follows because the term $q(s', a') P(s|s', a')$ belongs to $q(s)$ by Eq. (9).

For the second case (Eq. (11)) where $k(s) = k(s')$ and $(s, a) \neq (s', a')$, we have that

$$\frac{\partial^2 \phi_H}{\partial q(s', a') \partial q(s, a)} = \frac{\partial^2}{\partial q(s', a') \partial q(s, a)} \left(-\alpha \sum_{s'' \in S_{k(s)+1}} \sum_{a'' \in A} \sqrt{q(s'') - q(s'', a'')} \right)$$

$$\begin{aligned}
&= \sum_{s'' \in S_{k(s)+1}} \sum_{a'' \in A} \frac{\partial^2}{\partial q(s', a') \partial q(s, a)} \left(-\alpha \sqrt{q(s'') - q(s'', a'')} \right) \\
&= \sum_{s'' \in S_{k(s)+1}} \sum_{a'' \in A} \frac{\partial}{\partial q(s, a)} \left(\frac{-\alpha P(s''|s', a')}{2\sqrt{q(s'') - q(s'', a'')}} \right) \\
&= \sum_{s'' \in S_k} \sum_{a'' \in A} \frac{\alpha P(s''|s, a) P(s''|s', a')}{4(q(s'') - q(s'', a''))^{3/2}}
\end{aligned}$$

where the third step follows from the previous calculation.

Finally, we finish the first part of the proof by verifying [Eq. \(12\)](#) that

$$\begin{aligned}
\frac{\partial^2 \phi_H}{\partial q(s, a)^2} &= -\frac{\partial^2}{\partial q(s, a)^2} \left(\sqrt{q(s, a)} + \alpha \sqrt{q(s) - q(s, a)} \right) \\
&\quad - \sum_{s'' \in S_{k(s)+1}} \sum_{a'' \in A} \frac{\partial^2}{\partial q(s, a)^2} \left(\alpha \sqrt{q(s'') - q(s'', a'')} \right) \\
&= \frac{1}{4q(s, a)^{3/2}} + \frac{\alpha}{4(q(s) - q(s, a))^{3/2}} + \sum_{s' \in S_k, s' \neq s_L} \sum_{a' \in A} \frac{\alpha P(s'|s, a)^2}{4(q(s') - q(s', a'))^{3/2}}.
\end{aligned}$$

Then, we are ready to prove [Eq. \(13\)](#). Indeed, we have

$$\begin{aligned}
w^\top \nabla^2 \phi_H(q) w &= \sum_{s \neq s_L} \sum_a \left(\frac{w(s, a)^2}{4q(s, a)^{3/2}} + \frac{\alpha w(s, a)^2}{4(q(s) - q(s, a))^{3/2}} \right) \\
&\quad + \sum_{0 < k < L} \sum_{s \in S_{k-1}} \sum_a w(s, a)^2 \sum_{s'' \in S_k} \sum_{a'' \in A} \frac{\alpha P(s''|s, a)^2}{4(q(s'') - q(s'', a''))^{3/2}} \\
&\quad + \sum_{0 < k < L} \sum_{s, s' \in S_{k-1}} \sum_{a, a' \in A: (s, a) \neq (s', a')} w(s, a) w(s', a') \sum_{s'' \in S_k} \sum_{a'' \in A} \frac{\alpha P(s''|s, a) P(s''|s', a')}{4(q(s'') - q(s'', a''))^{3/2}} \\
&\quad - 2 \sum_{0 < k < L} \sum_{s \in S_k} \sum_{s' \in S_{k-1}} \sum_{a, a'} \frac{\alpha P(s|s', a') w(s, a) w(s', a')}{4(q(s) - q(s, a))^{3/2}} \\
&= \sum_{s \neq s_L} \sum_a \left(\frac{w(s, a)^2}{4q(s, a)^{3/2}} + \frac{\alpha w(s, a)^2}{4(q(s) - q(s, a))^{3/2}} \right) \\
&\quad + \sum_{0 < k < L} \sum_{s'' \in S_k} \sum_{a'' \in A} \frac{\alpha}{4(q(s'') - q(s'', a''))^{3/2}} \left(\sum_{s \in S_{k-1}} \sum_a w(s, a) P(s''|s, a) \right)^2 \\
&\quad - 2 \sum_{0 < k < L} \sum_{s \in S_k} \sum_a \frac{\alpha w(s, a) \sum_{s' \in S_{k-1}} \sum_{a'} P(s|s', a') w(s', a')}{4(q(s) - q(s, a))^{3/2}} \\
&= \sum_{s \neq s_L} \sum_a \left(\frac{w(s, a)^2}{4q(s, a)^{3/2}} + \frac{\alpha w(s, a)^2}{4(q(s) - q(s, a))^{3/2}} \right) \\
&\quad + \sum_{0 < k < L} \sum_{s \in S_k} \sum_{a \in A} \frac{\alpha h(s)^2}{4(q(s) - q(s, a))^{3/2}} - 2 \sum_{0 < k < L} \sum_{s \in S_k} \sum_a \frac{\alpha w(s, a) h(s)}{4(q(s) - q(s, a))^{3/2}} \\
&= \frac{1}{4} \sum_{s \neq s_L} \sum_a \left(\frac{w(s, a)^2}{q(s, a)^{3/2}} + \frac{\alpha (h(s) - w(s, a))^2}{(q(s) - q(s, a))^{3/2}} \right),
\end{aligned}$$

finishing the proof. \square

Note that the Hessian is clearly non-diagonal. However, by using the alternative definition from [Eq. \(9\)](#), the Hessian has a layered structure which allows an induction-based analysis as we will show.

B Proof of Lemma 5

In this section, we provide the proof for Lemma 5. First, we introduce the following notations.

$$\begin{aligned} F_t(q) &= \langle q, \widehat{L}_{t-1} \rangle + \psi_t(q), & G_t &= \langle q_t, \widehat{L}_t \rangle + \psi_t(q), \\ q_t &= \operatorname{argmin}_{q \in \Omega} F_t(q), & \tilde{q}_t &= \operatorname{argmin}_{q \in \Omega} G_t(q). \end{aligned} \quad (15)$$

Note that the definition of q_t is consistent with Algorithm 1. With these notations, we decompose the regret $\mathbb{E} \left[\sum_{t=1}^T \langle q_t - u, \widehat{\ell}_t \rangle \right]$ against any occupancy measure $u \in \Omega$ into a stability term and a mixed penalty term by adding and subtracting $F_t(q_t)$ and $G_t(\tilde{q}_t)$:

$$= \mathbb{E} \left[\underbrace{\sum_{t=1}^T \left(\langle q_t, \widehat{\ell}_t \rangle + F_t(q_t) - G_t(\tilde{q}_t) \right)}_{\text{stability}} \right] + \mathbb{E} \left[\underbrace{\sum_{t=1}^T \left(G_t(\tilde{q}_t) - F_t(q_t) - \langle u, \widehat{\ell}_t \rangle \right)}_{\text{mixed penalty}} \right]. \quad (16)$$

The rest of the section is organized as follows:

1. With some auxiliary lemmas (Lemma 10 and Lemma 11), we prove in Lemma 12 that the update of the algorithm is smooth in the sense that q_t and \tilde{q}_t are close.
2. With the smoothness guarantee, we bound the stability term by the quadratic norm of $\widehat{\ell}_t$ with respect to the Hessian matrix in Lemma 13, using similar techniques from [16].
3. By standard analysis, we control the mixed penalty term in Lemma 14.
4. Picking a proper u that is closed to \hat{q} and specified in Lemma 15, we finish the proof of Lemma 5 in the end of this section.

We use the notation $M_1 \preceq M_2$ for two matrices M_1 and M_2 to denote the fact that $M_2 - M_1$ is positive semi-definite.

Lemma 10 (Convexity of Ω). *The set of valid occupancy measures Ω is convex.*

Proof. For any $u, v \in \Omega$ and $\lambda \in [0, 1]$, it suffices to verify that $p = \lambda u + (1 - \lambda)v$ satisfies the two constraints described in Section 2. For the first one: we have for any $k = 0, \dots, L - 1$,

$$\begin{aligned} \sum_{s \in S_k} \sum_{a \in A} p(s, a) &= \lambda \sum_{s \in S_k} \sum_{a \in A} u(s, a) + (1 - \lambda) \sum_{s \in S_k} \sum_{a \in A} v(s, a) \\ &= \lambda + (1 - \lambda) = 1. \end{aligned}$$

For the second one (Eq. (1)), we have for any $k = 1, \dots, L - 1$ and every state $s \in S_k$,

$$\begin{aligned} & \sum_{s' \in S_{k-1}, a'} P(s|s', a') p(s', a') \\ &= \sum_{s' \in S_{k-1}, a'} P(s|s', a') (\lambda u(s', a') + (1 - \lambda)v(s', a')) \\ &= \lambda \sum_{s' \in S_{k-1}, a'} P(s|s', a') u(s', a') + (1 - \lambda) \sum_{s' \in S_{k-1}, a'} P(s|s', a') v(s', a') \\ &= \lambda \sum_a u(s, a) + (1 - \lambda) \sum_a v(s, a) \\ &= \sum_a (\lambda u(s, a) + (1 - \lambda)v(s, a)) \\ &= \sum_a p(s, a). \end{aligned}$$

This finishes the proof. □

Lemma 11. For any occupancy measures q and p from Ω satisfying $\frac{1}{2}q(s, a) \leq p(s, a)$ for all state-action pair (s, a) , we have $\frac{1}{4}\nabla^2\psi_t(p) \preceq \nabla^2\psi_t(q)$.

Proof. Let $M_1 = \nabla^2\psi_t(p)$ and $M_2 = \nabla^2\psi_t(q)$. By Eq. (13) in Lemma 9, we have for any w ,

$$w^\top M_2 w = \frac{1}{4\eta_t} \sum_{s \neq s_L, a} \left\{ \frac{w(s, a)^2}{q(s, a)^{3/2}} + \frac{\alpha (h(s) - w(s, a))^2}{(q(s) - q(s, a))^{3/2}} \right\} + \beta \sum_{s \neq s_L, a} \frac{w(s, a)^2}{q(s, a)^2}.$$

According to the condition of the lemma and the fact $q(s) - q(s, a) = \sum_{b \neq a} q(s, b)$ and $p(s) - p(s, a) = \sum_{b \neq a} p(s, b)$, we have $\frac{1}{2}(q(s) - q(s, a)) \leq p(s) - p(s, a)$, and thus

$$\begin{aligned} w^\top M_2 w &\geq \frac{1}{4\eta_t} \sum_{s \neq s_L, a} \left\{ \frac{w(s, a)^2}{(2p(s, a))^{3/2}} + \frac{\alpha (h(s) - w(s, a))^2}{(2(p(s) - p(s, a)))^{3/2}} \right\} + \beta \sum_{s, a} \frac{w(s, a)^2}{(2p(s, a))^2} \\ &\geq \frac{1}{16\eta_t} \sum_{s \neq s_L, a} \left\{ \frac{w(s, a)^2}{p(s, a)^{3/2}} + \frac{\alpha (h(s) - w(s, a))^2}{(p(s) - p(s, a))^{3/2}} \right\} + \frac{\beta}{4} \sum_{s, a} \frac{w(s, a)^2}{p(s, a)^2} \\ &= \frac{1}{4} w^\top M_1 w. \end{aligned}$$

This finishes the proof. \square

With the help of Lemma 10 and Lemma 11, we now prove that q_t to \tilde{q}_t are close.

Lemma 12. With $\beta = 64L$, we have $\frac{1}{2}q_t(s, a) \leq \tilde{q}_t(s, a) \leq 2q_t(s, a)$ for all state-action pairs (s, a) (recall the definitions in Eq. (15)).

Proof. For simplicity, we denote H as the Hessian $\nabla^2\psi_t(q_t)$, and H_L as the Hessian $\nabla^2\phi_L(q_t)$ which is a diagonal matrix with $\frac{\beta}{q_t(s, a)^2}$ on the diagonal. Recalling that $\psi_t = \eta_t^{-1}\phi_H + \phi_L$, by the convexity of ϕ_H (Lemma 9), we have $H_L \preceq H$. Our goal is to prove $\|\tilde{q}_t - q_t\|_H \leq 1$. This is enough to prove the lemma statement because

$$1 \geq \|\tilde{q}_t - q_t\|_H \geq \|\tilde{q}_t - q_t\|_{H_L} = \beta \sum_{s, a} \frac{(\tilde{q}_t(s, a) - q_t(s, a))^2}{q_t(s, a)^2}, \quad (17)$$

and since $\beta \geq 9$, we have $(\tilde{q}_t(s, a) - q_t(s, a))^2 \leq \frac{q_t(s, a)^2}{\beta} \leq \left(\frac{q_t(s, a)}{3}\right)^2$, which indicates $\frac{1}{2}q_t(s, a) \leq \tilde{q}_t(s, a) \leq 2q_t(s, a)$.

To prove $\|\tilde{q}_t - q_t\|_H \leq 1$, it suffices to show that for any occupancy measure q' that satisfies $\|q' - q_t\|_H = 1$, we have $G_t(q') \geq G_t(q_t)$, because this implies that, as the minimizer of the convex function G_t , \tilde{q}_t must be within the convex set $\{q : \|q - q_t\|_H \leq 1\}$. To this end, we bound $G_t(q')$ as follows

$$\begin{aligned} G_t(q') &= G_t(q_t) + \nabla G_t(q_t)^\top (q' - q_t) + \frac{1}{2} \|q' - q_t\|_{\nabla^2\psi_t(\xi)}^2 \\ &= G_t(q_t) + \nabla F_t(q_t)^\top (q' - q_t) + \widehat{\ell}_t^\top (q' - q_t) + \frac{1}{2} \|q' - q_t\|_{\nabla^2\psi_t(\xi)}^2 \\ &\geq G_t(q_t) - \left\| \widehat{\ell}_t \right\|_{H^{-1}} \|q' - q_t\|_H + \frac{1}{2} \|q' - q_t\|_{\nabla^2\psi_t(\xi)}^2 \\ &= G_t(q_t) - \left\| \widehat{\ell}_t \right\|_{H^{-1}} + \frac{1}{2} \|q' - q_t\|_{\nabla^2\psi_t(\xi)}^2 \end{aligned}$$

where in the first step we use Taylor's expansion with ξ being a point between q_t and q' , in the second step we use the definition of G_t and F_t (in Eq. (15)), and in the third step we use Hölder's inequality and the first order optimality condition $\langle \nabla F_t(q_t), q' - q_t \rangle \geq 0$.

Repeating the earlier argument in Eq. (17), $\|q' - q_t\|_H = 1$ also implies $\frac{1}{2}q_t(s, a) \leq q'(s, a) \leq 2q_t(s, a)$ and thus $\frac{1}{2}q_t(s, a) \leq \xi(s, a) \leq 2q_t(s, a)$. Since ξ is in Ω by Lemma 10, we continue to bound the last expression using Lemma 11:

$$G_t(q_t) - \left\| \widehat{\ell}_t \right\|_{H^{-1}} + \frac{1}{2} \|q' - q_t\|_{\nabla^2\psi_t(\xi)}^2$$

$$\begin{aligned}
&\geq G_t(q_t) - \left\| \widehat{\ell}_t \right\|_{H^{-1}} + \frac{1}{8} \|q' - q\|_H^2 \\
&= G_t(q_t) - \left\| \widehat{\ell}_t \right\|_{H^{-1}} + \frac{1}{8}.
\end{aligned} \tag{18}$$

Finally, we bound the term $\left\| \widehat{\ell}_t \right\|_{H^{-1}}$. By the definition of $\widehat{\ell}_t$, with shorthand $\mathbb{I}(s, a) \triangleq \mathbb{I}\{s_{k(s)} = s, a_{k(s)} = a\}$ we show that

$$\begin{aligned}
\left\| \widehat{\ell}_t \right\|_{H^{-1}}^2 &\leq \left\| \widehat{\ell}_t \right\|_{H_L^{-1}}^2 && (H_L \preceq H) \\
&= \sum_{s,a} \frac{\mathbb{I}(s, a) \ell_t(s, a)^2}{q_t(s, a)^2} \frac{q_t(s, a)^2}{\beta} \\
&\leq \sum_{s,a} \frac{\mathbb{I}(s, a)}{\beta} = \frac{L}{\beta} = \frac{1}{64}.
\end{aligned}$$

Plugging it into Eq. (18), we conclude that $G_t(q') \geq G_t(q_t)$, which finishes the proof. \square

We are now ready to bound the first term in Eq. (16).

Lemma 13. *With $\beta = 64L$, we have*

$$\sum_{t=1}^T \left(\left\langle q_t, \widehat{\ell}_t \right\rangle + F_t(q_t) - G_t(\tilde{q}_t) \right) \leq 8 \sum_{t=1}^T \left\| \widehat{\ell}_t \right\|_{\nabla^{-2}\psi_t(q_t)}^2.$$

Proof. We first lower bound the term $\left\langle q_t, \widehat{\ell}_t \right\rangle + F_t(q_t) - G_t(\tilde{q}_t)$ as

$$\begin{aligned}
&\left\langle q_t, \widehat{\ell}_t \right\rangle + F_t(q_t) - G_t(\tilde{q}_t) \\
&= \left\langle q_t, \widehat{\ell}_t + \widehat{L}_{t-1} \right\rangle + \psi_t(q_t) - G_t(\tilde{q}_t) \\
&= G_t(q_t) - G_t(\tilde{q}_t) \\
&= \langle \nabla G_t(\tilde{q}_t), q_t - \tilde{q}_t \rangle + \frac{1}{2} \|q_t - \tilde{q}_t\|_{\nabla^2 \psi_t(\xi)}^2 \\
&\geq \frac{1}{2} \|q_t - \tilde{q}_t\|_{\nabla^2 \psi_t(\xi)}^2,
\end{aligned}$$

where in the second to last step we apply Taylor's expansion with ξ being a point between q_t and \tilde{q}_t , and in the last step we use the first order optimality condition of \tilde{q}_t .

On the other hand, we can upper bound this term as

$$\begin{aligned}
&\left\langle q_t, \widehat{\ell}_t \right\rangle + F_t(q_t) - G_t(\tilde{q}_t) \\
&= \left\langle q_t - \tilde{q}_t, \widehat{\ell}_t \right\rangle + F_t(q_t) - F_t(\tilde{q}_t) \\
&\leq \left\langle q_t - \tilde{q}_t, \widehat{\ell}_t \right\rangle \\
&\leq \|q_t - \tilde{q}_t\|_{\nabla^2 \psi_t(\xi)} \left\| \widehat{\ell}_t \right\|_{\nabla^{-2} \psi_t(\xi)},
\end{aligned}$$

where the second step is by the optimality of q_t and the last step is by Hölder's inequality. Combining the lower and upper bounds we arrive at:

$$\left\langle q_t, \widehat{\ell}_t \right\rangle + F_t(q_t) - G_t(\tilde{q}_t) \leq 2 \left\| \widehat{\ell}_t \right\|_{\nabla^{-2} \psi_t(\xi)}^2.$$

Moreover, by Lemma 12, we know that \tilde{q}_t satisfies $\frac{1}{2}q_t(s, a) \leq \tilde{q}_t(s, a) \leq 2q_t(s, a)$ for all state-action pairs. Since ξ is a middle point between q_t and \tilde{q}_t , it satisfies $\frac{1}{2}q_t(s, a) \leq \xi(s, a) \leq 2q_t(s, a)$ for all state-action pairs as well. According to Lemma 11, we then have $\left\| \widehat{\ell}_t \right\|_{\nabla^{-2} \psi_t(\xi)}^2 \leq 4 \left\| \widehat{\ell}_t \right\|_{\nabla^{-2} \psi_t(q_t)}^2$, which completes the proof. \square

Next, we bound the second term in Eq. (16).

Lemma 14. For any $u \in \Omega$, we have (with $1/\eta_0 \triangleq 0$)

$$\sum_{t=1}^T \left(G_t(\tilde{q}_t) - F_t(q_t) - \langle u, \hat{\ell}_t \rangle \right) = \phi_L(u) - \phi_L(q_1) + \sum_{t=1}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) (\phi_H(u) - \phi_H(q_t)).$$

Proof. Due to the optimality of \tilde{q}_t , we have $G_t(\tilde{q}_t) \leq G_t(q_{t+1})$ and also $G_T(\tilde{q}_T) \leq G_T(u)$. With the help of these inequalities, we proceed as

$$\begin{aligned} & \sum_{t=1}^T \left(G_t(\tilde{q}_t) - F_t(q_t) - \langle u, \hat{\ell}_t \rangle \right) \\ & \leq \sum_{t=1}^T \left(G_t(q_{t+1}) - F_t(q_t) - \langle u, \hat{\ell}_t \rangle \right) \\ & = -F_1(q_1) + \sum_{t=2}^T (G_{t-1}(q_t) - F_t(q_t)) + G_T(u) - \langle u, \hat{L}_T \rangle \\ & = -\psi_1(q_1) - \sum_{t=2}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \phi_H(q_t) + \psi_T(u) \\ & = -\phi_H(q_1) - \phi_L(q_1) - \sum_{t=2}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \phi_H(q_t) + \frac{\phi_H(u)}{\eta_T} + \phi_L(u) \\ & = \phi_L(u) - \phi_L(q_1) + \sum_{t=1}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) (\phi_H(u) - \phi_H(q_t)), \end{aligned}$$

finishing the proof. \square

While it is tempting to set $u = \hat{q}$ to obtain a regret bound, note that $\phi_L(\hat{q})$ can potentially grow to infinity. To this end, we will set u as some point close enough to \hat{q} , specified in the following lemma.

Lemma 15. Suppose $\beta = 64L$. Let $v = (1 - \frac{1}{T})\hat{q} + \frac{q_1}{T}$, where \hat{q} is the optimal occupancy measure and $q_1 = \operatorname{argmin}_{q \in \Omega} \psi_1(q)$ is the initial occupancy measure. Then v satisfies the following:

- $\mathbb{E} \left[\sum_{t=1}^T \langle v - \hat{q}, \hat{\ell}_t \rangle \right] \leq 2L$,
- $\phi_L(v) - \phi_L(q_1) \leq 64L|S||A| \log T$,
- $\phi_H(v) - \phi_H(\hat{q}) \leq \frac{(1+\alpha)|S||A|}{T}$.

Proof. The first statement is by direct calculation:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \langle v - \hat{q}, \hat{\ell}_t \rangle \right] &= \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T \langle q_1 - \hat{q}, \hat{\ell}_t \rangle \right] \\ &= \frac{1}{T} \left\langle q_1 - \hat{q}, \sum_{t=1}^T \ell_t \right\rangle \\ &\leq \frac{1}{T} \|q_1 - \hat{q}\|_1 \left\| \sum_{t=1}^T \ell_t \right\|_\infty \leq \frac{2LT}{T} = 2L. \end{aligned}$$

The second statement directly uses the definition of ϕ_L :

$$\phi_L(v) - \phi_L(q_1) = 64L \sum_{s,a} \log \left(\frac{q_1(s,a)}{v(s,a)} \right) \leq 64L|S||A| \log(T).$$

Finally, we verify the last statement:

$$\begin{aligned}
& \phi_H(v) - \phi_H(\hat{q}) \\
&= \sum_{s,a} \left(\sqrt{\hat{q}(s,a)} + \alpha \sqrt{\hat{q}(s) - \hat{q}(s,a)} - \sqrt{v(s,a)} - \alpha \sqrt{v(s) - v(s,a)} \right) \\
&= \sum_{s,a} \left(\sqrt{\hat{q}(s,a)} - \sqrt{\frac{T-1}{T} \hat{q}(s,a) + \frac{q_1(s,a)}{T}} \right) \\
&\quad + \sum_{s,a} \alpha \left(\sqrt{\hat{q}(s) - \hat{q}(s,a)} - \sqrt{\frac{T-1}{T} (\hat{q}(s) - \hat{q}(s,a)) + \frac{q_1(s) - q_1(s,a)}{T}} \right) \\
&\leq \left(1 - \sqrt{\frac{T-1}{T}} \right) \sum_{s,a} \left(\sqrt{\hat{q}(s,a)} + \alpha \sqrt{\hat{q}(s) - \hat{q}(s,a)} \right) \\
&= \frac{(\sqrt{T} - \sqrt{T-1})(\sqrt{T} + \sqrt{T-1})}{\sqrt{T}(\sqrt{T} + \sqrt{T-1})} ((1+\alpha)|S||A|) \\
&\leq \frac{(1+\alpha)|S||A|}{T}.
\end{aligned}$$

□

Finally, we are ready to prove [Lemma 5](#).

Proof of Lemma 5. By combining the [Lemma 14](#) and [13](#) and taking expectation on both sides, we have that $\mathbb{E} \left[\sum_{t=1}^T \langle q_t - u, \hat{\ell}_t \rangle \right]$ is bounded by

$$\phi_L(u) - \phi_L(q_1) + \sum_{t=1}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \mathbb{E} [\phi_H(u) - \phi_H(q_t)] + 8 \sum_{t=1}^T \mathbb{E} \left[\left\| \hat{\ell}_t \right\|_{\nabla^{-2} \psi_t(q_t)}^2 \right].$$

Picking the intermediate occupancy measure $v = (1 - \frac{1}{T}) \hat{q} + \frac{1}{T} q_1$, by [Lemma 15](#), we have

$$\begin{aligned}
\text{Reg}_T &= \mathbb{E} \left[\sum_{t=1}^T \langle q_t - \hat{q}, \hat{\ell}_t \rangle \right] \\
&= \mathbb{E} \left[\sum_{t=1}^T \langle q_t - v, \hat{\ell}_t \rangle \right] + \mathbb{E} \left[\sum_{t=1}^T \langle v - \hat{q}, \hat{\ell}_t \rangle \right] \\
&\leq 2L + 64L|S||A| \log T + \sum_{t=1}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \mathbb{E} [\phi_H(v) - \phi_H(q_t)] + 8 \sum_{t=1}^T \mathbb{E} \left[\left\| \hat{\ell}_t \right\|_{\nabla^{-2} \psi_t(q_t)}^2 \right] \\
&\leq 2L + 64L|S||A| \log T + 8 \sum_{t=1}^T \mathbb{E} \left[\left\| \hat{\ell}_t \right\|_{\nabla^{-2} \psi_t(q_t)}^2 \right] \\
&\quad + \left(\frac{1}{\eta_T} - \frac{1}{\eta_0} \right) \mathbb{E} [\phi_H(v) - \phi_H(\hat{q})] + \sum_{t=1}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \mathbb{E} [\phi_H(\hat{q}) - \phi_H(q_t)] \\
&\leq \mathcal{O}(L|S||A| \log T) + \sum_{t=1}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \mathbb{E} [\phi_H(\hat{q}) - \phi_H(q_t)] + 8 \sum_{t=1}^T \mathbb{E} \left[\left\| \hat{\ell}_t \right\|_{\nabla^{-2} \psi_t(q_t)}^2 \right]
\end{aligned}$$

where the last line follows from the fact $\frac{(1+\alpha)|S||A|}{\gamma\sqrt{T}} = \mathcal{O}(1)$.

□

C Proof of Lemma 6

Recall the notation defined in Section 4:

$$h_s(\pi) = \sum_{a \in A} \sqrt{\pi(a|s)} + \alpha \sqrt{1 - \pi(a|s)}.$$

Clearly, for any mapping (deterministic policy) $\pi : S \rightarrow A$, we have $h_s(\pi) = 1 + \alpha(|A| - 1)$ for all state s . Therefore, we can decompose $\phi_H(\hat{q}) - \phi_H(q_t)$ as

$$\begin{aligned} \phi_H(\hat{q}) - \phi_H(q_t) &= \sum_{s \neq s_L} \sqrt{q_t(s)} h_s(\pi_t) - \sum_{s \neq s_L} \sqrt{\hat{q}(s)} h_s(\hat{\pi}) \\ &= \sum_{s \neq s_L} \sqrt{q_t(s)} (h_s(\pi_t) - h_s(\hat{\pi})) + (1 + \alpha(|A| - 1)) \sum_{s \neq s_L} (\sqrt{q_t(s)} - \sqrt{\hat{q}(s)}). \end{aligned}$$

In the rest of this section, we first bound the term $\sum_{s \neq s_L} \sqrt{q_t(s)} (h_s(\pi_t) - h_s(\hat{\pi}))$ in Lemma 16, and then show that the term $\sum_{s \neq s_L} (\sqrt{q_t(s)} - \sqrt{\hat{q}(s)})$ can be bounded as

$$\begin{aligned} \sum_{s \neq s_L} (\sqrt{q_t(s)} - \sqrt{\hat{q}(s)}) &= \sum_{s \neq s_L} (\sqrt{q_t(s)} - \sqrt{q^\pi(s)}) + \sum_{s \neq s_L} (\sqrt{q^\pi(s)} - \sqrt{\hat{q}(s)}) \\ &\leq \sqrt{|S|L} \left(\sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a)} + \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} \hat{q}(s, a)} \right) \\ &\leq 2 \sqrt{|S|L \sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a) + \hat{q}(s, a)} \end{aligned}$$

for any mapping π , where the second line follows Lemma 19 and Lemma 20, in which we apply the key induction Lemma 18 based on the *state reach probability* defined in Lemma 17.

Lemma 16. *For any policy π_1 and mapping $\pi_2 : S \rightarrow A$, we have*

$$\sum_{s \neq s_L} \sqrt{q_1(s)} (h_s(\pi_1) - h_s(\pi_2)) \leq (1 + \alpha) \sum_{s \neq s_L} \sum_{a \neq \pi(s)} \sqrt{q_1(s, a)},$$

where π is any mapping from S to A and $q_1 = q^{\pi_1}$.

Proof. Direct calculation shows:

$$\begin{aligned} &h_s(\pi_1) - h_s(\pi_2) \\ &= h_s(\pi_1) - 1 - \alpha(|A| - 1) \\ &= \sqrt{\pi_1(\pi(s)|s)} + \alpha \sqrt{1 - \pi_1(\pi(s)|s)} - 1 + \sum_{a \neq \pi(s)} (\sqrt{\pi_1(a|s)} + \alpha \sqrt{1 - \pi_1(a|s)} - \alpha) \\ &\leq \alpha \sqrt{1 - \pi_1(\pi(s)|s)} + \sum_{a \neq \pi(s)} \sqrt{\pi_1(a|s)} \\ &= \alpha \sqrt{\sum_{a \neq \pi(s)} \pi_1(a|s)} + \sum_{a \neq \pi(s)} \sqrt{\pi_1(a|s)} \\ &\leq (\alpha + 1) \sum_{a \neq \pi(s)} \sqrt{\pi_1(a|s)}. \end{aligned}$$

Multiplying both sides by $\sqrt{q_1(s)}$ and summing over s prove the lemma. \square

Lemma 17. *For any policy π , define its associated reachability probability $p^\pi : S \times A \times S \rightarrow [0, 1]$ as*

$$p^\pi(s'|s, a) = \begin{cases} 0, & \text{if } k(s') \leq k(s), \\ P(s'|s, a), & \text{if } k(s') = k(s) + 1, \\ \sum_{s_m \in S_{k(s')-1}} p^\pi(s_m|s, a) \sum_a \pi(a|s_m) P(s'|s_m, a), & \text{if } k(s') > k(s) + 1, \end{cases} \quad (19)$$

which is simply the probability of reaching state s' after taking action a at state s and then following policy π . For any state-action pair (s, a) , policy π , and $k = 0, \dots, L - 1$, we have

$$\sum_{s' \in S_k} p^\pi(s'|s, a) \leq 1, \quad (20)$$

which implies $\sum_{s \neq s_L} p^\pi(s'|s, a) \leq L$.

Proof. Eq. (20) is clear just based on the definition of p^π . We provide a proof by induction for completeness. Clearly, it holds for all layer l with $l \leq k(s)$ and also $l = k(s) + 1$ where $\sum_{s' \in S_l} p^\pi(s'|s, a) = \sum_{s' \in S_l} P(s'|s, a) = 1$. Now assume that Eq. (20) holds for some layer $l \geq k(s) + 1$. For layer $l + 1$, we have

$$\begin{aligned} & \sum_{s' \in S_{l+1}} p^\pi(s'|s, a) \\ &= \sum_{s' \in S_{l+1}} \left(\sum_{s_m \in S_l} p^\pi(s_m|s, a) \sum_a \pi(a|s_m) P(s'|s_m, a) \right) \\ &= \sum_{s_m \in S_l} p^\pi(s_m|s, a) \left(\sum_{s' \in S_{l+1}} \sum_a \pi(a|s_m) P(s'|s_m, a) \right) \\ &= \sum_{s_m \in S_l} p^\pi(s_m|s, a) \leq 1, \end{aligned}$$

finishing the proof. □

Based on the concept of reachability probability, we prove the following key induction lemma.

Lemma 18. *If a policy π and non-negative functions $f : S \times A \rightarrow \mathbb{R}_+ \cup \{0\}$ and $g : S \times A \rightarrow \mathbb{R}_+ \cup \{0\}$ satisfy*

$$f(s) \leq \sum_{s' \in S_{k(s)-1}} \sum_{a' \in A} g(s', a') P(s|s', a') + \sum_{s' \in S_{k(s)-1}} f(s') \left\{ \sum_{a' \in A} \pi(a'|s') P(s|s', a') \right\}, \forall s \neq s_0$$

and $f(s_0) = 0$, then we have for all $s \neq s_0$,

$$f(s) \leq \sum_{k=0}^{k(s)-1} \sum_{s' \in S_k} \sum_{a'} g(s', a') p^\pi(s|s', a'), \quad (21)$$

where p^π is the reachability probability defined in Eq. (19).

Proof. We prove the statement by induction. First, for $s \in S_1$, using the condition of the lemma we have

$$\begin{aligned} f(s) &\leq \sum_{s' \in S_0} \sum_{a'} g(s', a') P(s|s', a') + \sum_{s' \in S_0} f(s') \left\{ \sum_{a'} \pi(a'|s') P(s|s', a') \right\} \\ &= \sum_{s' \in S_0} \sum_{a'} g(s', a') p^\pi(s|s', a') \end{aligned}$$

where the second line follows from the fact $f(s_0) = 0$ and $p^\pi(s|s', a') = P(s|s', a')$ for $k(s) = k(s') + 1$. This proves the base case.

Now assume that Eq. (21) holds for all states in layer $l > 0$. For $s \in S_{l+1}$, we have

$$\begin{aligned}
f(s) &\leq \sum_{s' \in S_l} \sum_{a'} g(s', a') P(s|s', a') + \sum_{s' \in S_l} f(s') \left\{ \sum_{a'} \pi(a'|s') P(s|s', a') \right\} \\
&\leq \sum_{s' \in S_l} \sum_{a'} g(s', a') P(s|s', a') \\
&\quad + \sum_{s' \in S_l} \left(\sum_{k=0}^{l-1} \sum_{s'' \in S_k} \sum_{a''} g(s'', a'') p^\pi(s'|s'', a'') \right) \left(\sum_{a'} \pi(a'|s') P(s|s', a') \right) \\
&= \sum_{s' \in S_l} \sum_{a'} g(s', a') p^\pi(s|s', a') \\
&\quad + \sum_{k=0}^{l-1} \sum_{s'' \in S_k} \sum_{a''} g(s'', a'') \left(\sum_{s' \in S_l} p^\pi(s'|s'', a'') \sum_{a'} \pi(a'|s') P(s|s', a') \right) \\
&= \sum_{s' \in S_l} \sum_{a'} g(s', a') p^\pi(s|s', a') + \sum_{k=0}^{l-1} \sum_{s'' \in S_k} \sum_{a''} g(s'', a'') p^\pi(s|s'', a'') \\
&= \sum_{k=0}^l \sum_{s' \in S_k} \sum_{a'} g(s', a') p^\pi(s|s', a'),
\end{aligned}$$

where the second step uses the induction hypothesis and the fourth step uses the definition of the reachability probability. This finishes the induction. \square

We now apply the induction lemma to prove the following two key lemmas.

Lemma 19. For any policy π_1 and mapping $\pi_2 : S \rightarrow A$, we have

$$\sum_{s \neq s_L} \sqrt{q_1(s)} - \sqrt{q_2(s)} \leq \sqrt{|S|L} \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi_2(s)} q_1(s, a)} \quad (22)$$

where we denote by $q_1 = q_1^\pi$ and $q_2 = q_2^\pi$ the occupancy measures of π_1 and π_2 respectively.

Proof. We first bound $\sqrt{q_1(s)} - \sqrt{q_2(s)}$ by $\sqrt{\mathbb{I}_s (q_1(s) - q_2(s))}$ where $\mathbb{I}_s \triangleq \mathbb{I}\{q_1(s) \geq q_2(s)\}$. Define $f(s) = \mathbb{I}_s (q_1(s) - q_2(s))$ and $g(s, a) = \mathbb{I}\{a \neq \pi_2(s)\} q_1(s, a)$. Our goal is to prove for any $s \neq s_0$:

$$f(s) \leq \sum_{s' \in S_{k(s)-1}} \sum_{a' \in A} g(s', a') P(s|s', a') + \sum_{s' \in S_{k(s)-1}} f(s') \left(\sum_{a' \in A} \pi_1(a'|s') P(s|s', a') \right) \quad (23)$$

so that we can apply Lemma 18 (clearly we have $f(s_0) = 0$). To prove Eq. (23), consider a fixed state $s \in S_k$ for some $k > 0$. We rewrite $q_1(s) - q_2(s)$ as

$$q_1(s) - q_2(s) = \sum_{s' \in S_{k-1}} \left[q_1(s') \sum_{a' \in A} \pi_1(a'|s') P(s|s', a') - q_2(s') \sum_{a' \in A} \pi_2(s', a') P(s|s', a') \right].$$

For state $s' \in S_{k-1}$ satisfying $q_1(s') \leq q_2(s')$, we have

$$\begin{aligned}
& q_1(s') \sum_{a' \in A} \pi_1(a'|s')P(s|s', a') - q_2(s') \sum_{a' \in A} \pi_2(a'|s')P(s|s', a') \\
& \leq q_1(s') \sum_{a' \in A} \pi_1(a'|s')P(s|s', a') - q_1(s') \sum_{a' \in A} \pi_2(a'|s')P(s|s', a') \\
& = q_1(s') \left(\sum_{a' \in A} (\pi_1(a'|s') - \pi_2(a'|s')) P(s|s', a') \right) \\
& \leq q_1(s') \left(\sum_{a' \neq \pi_2(s')} (\pi_1(a'|s')) P(s|s', a') \right) \\
& = \sum_{a' \neq \pi_2(s')} q_1(s', a')P(s|s', a').
\end{aligned}$$

where the fourth line follows from the fact $\pi_1(a'|s') - \pi_2(a'|s') = \pi_1(a'|s') - 1 \leq 0$ when $a' = \pi_2(s')$.

For the other states with $q_1(s') \geq q_2(s')$, we have

$$\begin{aligned}
& q_1(s') \sum_{a' \in A} \pi_1(a'|s')P(s|s', a') - q_2(s') \sum_{a' \in A} \pi_2(a'|s')P(s|s', a') \\
& = (q_1(s') - q_2(s')) \sum_{a' \in A} \pi_1(a'|s')P(s|s', a') + q_2(s') \sum_{a' \in A} (\pi_1(a'|s') - \pi_2(a'|s')) P(s|s', a') \\
& \leq (q_1(s') - q_2(s')) \sum_{a' \in A} \pi_1(a'|s')P(s|s', a') + q_2(s') \sum_{a' \neq \pi_2(s')} \pi_1(a'|s')P(s|s', a') \\
& \leq \mathbb{I}_{s'} (q_1(s') - q_2(s')) \sum_{a' \in A} \pi_1(a'|s')P(s|s', a') + \sum_{a' \neq \pi_2(s')} q_1(s', a')P(s|s', a')
\end{aligned}$$

where the last line is due to the condition $q_1(s') \geq q_2(s')$.

Combining these two cases together yields [Eq. \(23\)](#). Therefore, applying [Lemma 18](#) gives

$$\begin{aligned}
\mathbb{I}_s (q_1(s) - q_2(s)) & \leq \sum_{k=0}^{k(s)-1} \sum_{s' \in S_k} \sum_{a'} g(s', a') p^{\pi_1}(s|s', a') \\
& = \sum_{k=0}^{k(s)-1} \sum_{s' \in S_k} \sum_{a' \neq \pi_2(s')} q_1(s', a') p^{\pi_1}(s|s', a'),
\end{aligned}$$

for all $s \neq s_0$. Taking square root and summation over all states, we have

$$\sum_{s \neq s_L} \sqrt{\mathbb{I}_s (q_1(s) - q_2(s))} \leq \sum_{s \neq s_L} \sqrt{\sum_{k=0}^{k(s)-1} \sum_{s' \in S_k} \sum_{a' \neq \pi_2(s')} q_1(s', a') p^{\pi_1}(s|s', a')}.$$

Notice that

$$\begin{aligned}
\sum_{s \neq s_L} \sum_{k=0}^{k(s)-1} \sum_{s' \in S_k} \sum_{a' \neq \pi_2(s')} q_1(s', a') p^{\pi_1}(s|s', a') & \leq \sum_{s' \neq s_L} \sum_{a' \neq \pi_2(s')} q_1(s', a') \left(\sum_{s \neq s_L} p^{\pi_1}(s|s', a') \right) \\
& \leq L \sum_{s \neq s_L} \sum_{a' \neq \pi_2(s)} q_1(s, a),
\end{aligned}$$

where the last step uses [Lemma 17](#). Finally by Hölder's inequality, we arrive at

$$\begin{aligned}
\sum_{s \neq s_L} \sqrt{q_1(s)} - \sqrt{q_2(s)} &\leq \sum_{s \neq s_L} \sqrt{\sum_{k=0}^{k(s)-1} \sum_{s' \in S_k} \sum_{a' \neq \pi_2(s')} q_1(s', a') p^{\pi_1}(s|s', a')} \\
&\leq \left(L \sum_{s \neq s_L} \sum_{a \neq \pi_2(s)} q_1(s, a) \right)^{1/2} \left(\sum_{s \neq s_L} 1 \right)^{1/2} \\
&= \sqrt{|S|L} \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi_2(s)} q_1(s, a)},
\end{aligned}$$

which concludes the proof. \square

Lemma 20. For two deterministic policies π_1 and π_2 , we have

$$\sum_{s \neq s_L} \sqrt{q_1(s)} - \sqrt{q_2(s)} \leq \sqrt{|S|L} \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi_1(s)} q_2(s, a)}$$

where we denote by $q_1 = q^{\pi_1}$ and $q_2 = q^{\pi_2}$ the occupancy measures of π_1 and π_2 respectively.

Proof. The proof is similar to that of [Lemma 19](#). First note that $\sqrt{q_1(s)} - \sqrt{q_2(s)} \leq \sqrt{\mathbb{I}_s(q_1(s) - q_2(s))}$ where $\mathbb{I}_s \triangleq \mathbb{I}\{q_1(s) \geq q_2(s)\}$. Define $f(s) = \mathbb{I}_s(q_1(s) - q_2(s))$ and $g(s, a) = \pi_1(a|s) \sum_{a' \neq \pi_1(s)} q_2(s, a')$. Our goal is to prove the following inequality

$$f(s) \leq \sum_{s' \in S_{k(s)-1}} \sum_{a' \in A} g(s', a') P(s|s', a') + \sum_{s' \in S_{k(s)-1}} f(s') \left(\sum_{a' \in A} \pi_1(a'|s') P(s|s', a') \right)$$

for all $s \neq s_L$ so that we can apply [Lemma 18](#) again ($f(s_0) = 0$ holds clearly). To show this, consider a fixed state $s \in S_k$ for some $k > 0$ and rewrite the term $q_1(s) - q_2(s)$ as

$$\begin{aligned}
q_1(s) - q_2(s) &= \sum_{s' \in S_{k-1}} \sum_{a'} (q_1(s', a') P(s|s', a') - q_2(s', a') P(s|s', a')) \\
&= \sum_{s' \in S_{k-1}} (q_1(s') P(s|s', \pi_1(s')) - q_2(s') P(s|s', \pi_2(s')))
\end{aligned}$$

since both π_1 and π_2 are deterministic policies. Then for the states $s' \in S_{k-1}$ with $\pi_1(s') = \pi_2(s')$, we have

$$q_1(s') P(s|s', \pi_1(s')) - q_2(s') P(s|s', \pi_2(s')) \leq \mathbb{I}_{s'} (q_1(s') - q_2(s')) P(s|s', \pi_1(s'))$$

For the other states with $\pi_1(s') \neq \pi_2(s')$, we have

$$\begin{aligned}
&q_1(s') P(s|s', \pi_1(s')) - q_2(s') P(s|s', \pi_2(s')) \\
&\leq (q_1(s') - q_2(s')) P(s|s', \pi_1(s')) + q_2(s') P(s|s', \pi_1(s')) - q_2(s') P(s|s', \pi_2(s')) \\
&\leq (q_1(s') - q_2(s')) P(s|s', \pi_1(s')) + q_2(s') P(s|s', \pi_1(s')) - q_2(s', \pi_1(s')) P(s|s', \pi_1(s')) \\
&= (q_1(s') - q_2(s')) P(s|s', \pi_1(s')) + \left[\sum_{a \neq \pi_1(s')} q_2(s', a) \right] P(s|s', \pi_1(s')) \\
&= (q_1(s') - q_2(s')) \sum_{a'} \pi_1(a'|s') P(s|s', a') + \sum_{a'} \pi_1(a'|s') \left[\sum_{a \neq \pi_1(s')} q_2(s', a) \right] P(s|s', a') \\
&\leq \mathbb{I}_{s'} (q_1(s') - q_2(s')) \sum_{a'} \pi_1(a'|s') P(s|s', a') + \sum_{a'} g(s', a') P(s|s', a'),
\end{aligned}$$

where the third line follows from the fact that $-q_2(s') P(s|s', \pi_2(s')) \leq -q_2(s', \pi_1(s')) P(s|s', \pi_1(s'))$, which holds because the right-hand side is simply zero when $\pi_1(s') \neq \pi_2(s')$.

Combining these two cases proves the desired inequality for all states. Applying [Lemma 18](#), we arrive at

$$f(s) \leq \sum_{k=0}^{k(s)-1} \sum_{s' \in S_k} \sum_{a' \in A} g(s', a') p^{\pi_1}(s|s', a').$$

By the same arguments used in the proof of [Lemma 19](#), we have shown

$$\sum_{s \neq s_L} \sqrt{q_1(s)} - \sqrt{q_2(s)} \leq \sqrt{|S|L} \sqrt{\sum_{s \neq s_L} \sum_{a \in A} g(s, a)}.$$

Noticing that

$$\sum_{s \neq s_L} \sum_{a \in A} g(s, a) = \sum_{s \neq s_L} \left(\sum_{a \in A} \pi_1(a|s) \right) \left[\sum_{a' \neq \pi_1(s)} q_2(s, a') \right] = \sum_{s \neq s_L} \sum_{a' \neq \pi_1(s)} q_2(s, a')$$

finishes the proof. \square

Finally, we are ready to prove [Lemma 6](#).

Proof of Lemma 6. Recall the calculation at the beginning of this section:

$$\begin{aligned} \phi_H(\dot{q}) - \phi_H(q_t) &= \sum_{s \neq s_L} \sqrt{q_t(s)} h_s(\pi_t) - \sum_{s \neq s_L} \sqrt{\dot{q}(s)} h_s(\dot{\pi}) \\ &= \sum_{s \neq s_L} \sqrt{q_t(s)} (h_s(\pi_t) - h_s(\dot{\pi})) + (1 + \alpha(|A| - 1)) \sum_{s \neq s_L} \left(\sqrt{q_t(s)} - \sqrt{\dot{q}(s)} \right). \end{aligned}$$

Using [Lemma 16](#) with $\pi_1 = \pi_t$ and $\pi_2 = \dot{\pi}$, we bound the first term by $(1 + \alpha) \sum_{s \neq s_L} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)}$ for any mapping $\pi : S \rightarrow A$. Then we decompose the term $\sum_{s \neq s_L} \left(\sqrt{q_t(s)} - \sqrt{\dot{q}(s)} \right)$ as

$$\begin{aligned} \sum_{s \neq s_L} \left(\sqrt{q_t(s)} - \sqrt{\dot{q}(s)} \right) &= \sum_{s \neq s_L} \left(\sqrt{q_t(s)} - \sqrt{q^\pi(s)} \right) + \sum_{s \neq s_L} \left(\sqrt{q^\pi(s)} - \sqrt{\dot{q}(s)} \right) \\ &\leq \sqrt{|S|L} \left(\sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a)} + \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} \dot{q}(s, a)} \right) \\ &\leq 2 \sqrt{|S|L \sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a) + \dot{q}(s, a)} \end{aligned}$$

where the second line is by using [Lemma 19](#) with $\pi_1 = \pi_t$ and $\pi_2 = \pi$, and [Lemma 20](#) with $\pi_1 = \pi$ and $\pi_2 = \dot{\pi}$. This provides one upper bound for $\sum_{s \neq s_L} \left(\sqrt{q_t(s)} - \sqrt{\dot{q}(s)} \right)$. On the other hand, it can also be trivially bounded as

$$\sum_{s \neq s_L} \sqrt{q_t(s)} - \sqrt{\dot{q}(s)} \leq \sum_{s \neq s_L} \sqrt{q_t(s)} \leq \sqrt{|S|L}$$

where the second step uses Cauchy-Schwarz inequality and the fact $\sum_{s \neq s_L} q_t(s) = L$. Combining everything shows that $\phi_H(\dot{q}) - \phi_H(q_t)$ is bounded by

$$(1 + \alpha) \sum_{s \neq s_L} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)} + (1 + \alpha|A|) \sqrt{|S|L} \min \left\{ 2 \sqrt{\sum_{s \neq s_L} \sum_{a \neq \pi(s)} q_t(s, a) + \dot{q}(s, a)}, 1 \right\},$$

finishing the proof. \square

D Proof of Lemma 7

In this section, we provide the proof for Lemma 7. Throughout the section, we use the shorthand $H \triangleq \nabla^2 \phi_H(q_t)$. The key is clearly to analyze the Hessian inverse H^{-1} , which is done in Appendix D.1. With a better understanding of the Hessian inverse, we then finish the proof in Appendix D.2.

D.1 Analyzing the Hessian inverse

To facilitate discussions, we first introduce some matrix notation. We see the Hessian H as a matrix in $\mathbb{R}^{(|S||A|) \times (|S||A|)}$ in the natural way.⁵ For subsets $D, E \subseteq S \times A$, we use $\mathbb{R}^{D \times E}$ to represent the set of matrices (in $\mathbb{R}^{|D| \times |E|}$) with the elements in D indexing their rows and the elements in E indexing their columns. The notation $M((s, a), (s', a'))$ represents the entry of a matrix M in the row indexed by (s, a) and in the column indexed by (s', a') . Let $U_k = \{(s, a) : s \in S_k, a \in A\}$ for $k = 0, \dots, L-1$, and $U_{j:k} = U_j \cup \dots \cup U_k$. We use similar notations for vectors, and define $\mathbf{0}_U \in \mathbb{R}^U$ be the all-zero vector with elements in U indexing its coordinates.

Define diagonal matrices $I_k, C_k, D_k \in \mathbb{R}^{U_k \times U_k}$ as

$$\begin{aligned} I_k &= \text{diag} \{1 : (s, a) \in U_k\}, \\ C_k &= \text{diag} \{c(s, a) : (s, a) \in U_k\}, \\ D_k &= \text{diag} \{d(s, a) : (s, a) \in U_k\}, \end{aligned}$$

where

$$c(s, a) = \frac{\alpha}{4(q(s) - q(s, a))^{3/2}} \quad \text{and} \quad d(s, a) = \frac{1}{4q(s, a)^{3/2}}.$$

Also define transition matrices $P_k \in \mathbb{R}^{U_{k-1} \times U_k}$ such that $P_k((s, a), (s', a')) = P(s'|s, a)$, and $\tilde{P}_k \in \mathbb{R}^{U_{0:k-1} \times U_k}$ such that $P_k((s, a), (s', a')) = P(s'|s, a)$ if $s' \in S_{k(s)+1}$ and $P_k((s, a), (s', a')) = 0$ otherwise.

Our first step is to write H in a recursive way with the help of a sequence of matrices:

Lemma 21. Define matrices $M_k \in \mathbb{R}^{U_{0:k} \times U_{0:k}}$ for $k = 0, \dots, L-1$ recursively as

$$M_k = \begin{pmatrix} M_{k-1} + \tilde{P}_k C_k \tilde{P}_k^\top & -\tilde{P}_k C_k \\ -C_k \tilde{P}_k^\top & C_k + D_k \end{pmatrix} = \begin{pmatrix} M_{k-1} & 0 \\ 0 & D_k \end{pmatrix} + \begin{pmatrix} \tilde{P}_k \\ -I_t \end{pmatrix} C_k \begin{pmatrix} \tilde{P}_k^\top & -I_t \end{pmatrix} \quad (24)$$

for $k = 1, \dots, L-1$, and $M_0 = C_0 + D_0$. Then we have $H = M_{L-1}$.

Proof. The proof is by a direct verification based on the calculation of the Hessian done in Lemma 9. The claim is that M_k consists of all the second-order derivatives with respect to $(s, a), (s', a') \in U_{0:k}$, but without the terms involving $P(s''|s, a)$ or $P(s''|s', a')$ for $s'' \in S_{k+1}$. To see this, note that this is clearly true for M_0 based on Lemma 9. Suppose this is true for M_{k-1} and consider M_k .

We first show that the block $\tilde{P}_k C_k \tilde{P}_k^\top$ corresponds to Eq. (11) plus the last term of Eq. (12), that is

$$\begin{aligned} & \left(\tilde{P}_k C_k \tilde{P}_k^\top \right) ((s, a), (s', a')) \\ &= \sum_{(s_1, a_1) \in U_k} \sum_{(s_2, a_2) \in U_k} \tilde{P}_k((s, a); (s_1, a_1)) \cdot C_k((s_1, a_1); (s_2, a_2)) \cdot \tilde{P}_k((s', a'); (s_2, a_2)) \\ &= \sum_{(s_1, a_1) \in U_k} \tilde{P}_k((s, a), (s_1, a_1)) \cdot C_k((s_1, a_1), (s_1, a_1)) \cdot \tilde{P}_k((s', a'), (s_1, a_1)) \\ & \quad + \sum_{(s_1, a_1) \in U_k} \sum_{(s_2, a_2) \in U_k, (s_1, a_1) \neq (s_2, a_2)} \tilde{P}_k((s, a), (s_1, a_1)) \cdot C_k((s_1, a_1), (s_2, a_2)) \cdot \tilde{P}_k((s', a'), (s_1, a_1)) \\ &= \sum_{(s_1, a_1) \in U_k} P(s_1|s, a) \cdot \frac{\alpha}{4(q(s_1) - q(s_1, a_1))^{3/2}} \cdot P(s_1|s', a') \\ &= \sum_{s'' \in S_k} \sum_{a'' \in A} \frac{\alpha P(s''|s, a) P(s''|s', a')}{4(q(s'') - q(s'', a''))^{3/2}} \end{aligned}$$

⁵Technically, S should be $S \setminus \{s_L\}$ instead.

where the third step uses the fact that $C_k((s_1, a_1), (s_2, a_2)) = 0$ when $(s_1, a_1) \neq (s_2, a_2)$.

Then we verify that the blocks $-\tilde{P}_k C_k$ and $-C_k \tilde{P}_k^\top$ correspond to Eq. (10). Direct calculation shows that, for $(s, a) \in U_k$ and $(s', a') \in U_{k-1}$,

$$\begin{aligned} (-\tilde{P}_k C_k)((s', a'), (s, a)) &= - \sum_{(s'', a'') \in U_k} \tilde{P}_k((s', a'), (s'', a'')) C_k((s'', a''), (s, a)) \\ &= - \frac{\alpha}{4(q(s) - q(s, a))^{3/2}} \cdot P(s|s', a'). \end{aligned}$$

Finally, the block $C_k + D_k$ corresponds to the first two terms of Eq. (12). This finishes the proof. \square

To study $H^{-1} = M_{L-1}^{-1}$, we next write M_k^{-1} in terms of M_{k-1}^{-1} .

Lemma 22. *The inverse of M_k (defined in Eq. (24)) is*

$$\begin{aligned} M_k^{-1} &= \begin{pmatrix} M_{k-1}^{-1} & 0 \\ 0 & D_k^{-1} \end{pmatrix} - \begin{pmatrix} M_{k-1}^{-1} \tilde{P}_k \\ -D_k^{-1} \end{pmatrix} \left(C_k^{-1} + D_k^{-1} + \tilde{P}_k^\top M_{k-1}^{-1} \tilde{P}_k \right)^{-1} \begin{pmatrix} \tilde{P}_k^\top M_{k-1}^{-1} & -D_k^{-1} \end{pmatrix} \\ &= \begin{pmatrix} M_{k-1}^{-1} - M_{k-1}^{-1} \tilde{P}_k W_k \tilde{P}_k^\top M_{k-1}^{-1} & M_{k-1}^{-1} \tilde{P}_k W_k D_k^{-1} \\ D_k^{-1} W_k \tilde{P}_k^\top M_{k-1}^{-1} & D_k^{-1} - D_k^{-1} W_k D_k^{-1} \end{pmatrix} \end{aligned} \quad (25)$$

where $W_k = (C_k^{-1} + D_k^{-1} + \tilde{P}_k^\top M_{k-1}^{-1} \tilde{P}_k)^{-1}$.

Proof. The first equality is by the Woodbury matrix identity

$$(A + UBV)^{-1} = A^{-1} - A^{-1}U(B^{-1} + VA^{-1}U)^{-1}VA^{-1}$$

and plugging in the definition of M_k from Eq. (24) with

$$A = \begin{pmatrix} M_{k-1}^{-1} & 0 \\ 0 & D_k^{-1} \end{pmatrix}, \quad U = V^\top = \begin{pmatrix} \tilde{P}_k \\ -I_t \end{pmatrix}, \quad \text{and} \quad B = C_k.$$

The second equality is by direct calculation. \square

The bottom right block of M_k^{-1} plays a key role in the analysis, and is denoted by

$$N_k = D_k^{-1} - D_k^{-1} W_k D_k^{-1} \quad (26)$$

for $k = 1, \dots, L-1$, and $N_0 = M_0^{-1}$. The next lemma shows that we can focus on N_k when analyzing specific quadratic forms of H^{-1} .

Lemma 23. *For any vector $w_k \in \mathbb{R}^{U_k}$, we have*

$$\left(\mathbf{0}_{U_{0:k-1}}^\top, w_k^\top, \mathbf{0}_{U_{k+1:L-1}}^\top \right) H^{-1} \begin{pmatrix} \mathbf{0}_{U_{0:k-1}} \\ w_k \\ \mathbf{0}_{U_{k+1:L-1}} \end{pmatrix} \leq w_k^\top N_k w_k,$$

where N_k is defined in Eq. (26).

Proof. Based on Eq. (24) and the fact that C_k is positive definite, we have

$$M_k \succeq \begin{pmatrix} M_{k-1} & 0 \\ 0 & D_k \end{pmatrix}.$$

Repeatedly using this fact, we can show

$$H = M_{L-1} \succeq \begin{pmatrix} M_k & 0 & \cdots & 0 \\ 0 & D_{k+1} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & D_{L-1} \end{pmatrix},$$

and thus

$$H^{-1} = M_{L-1}^{-1} \preceq \begin{pmatrix} M_k^{-1} & 0 & \cdots & 0 \\ 0 & D_{k+1}^{-1} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & D_{L-1}^{-1} \end{pmatrix}.$$

Note that for the last matrix above, the block with rows and columns indexed by elements in U_k is exactly N_k , based on [Lemma 22](#). Thus, taking the quadratic form on both sides with respect to the vector $(\mathbf{0}_{U_{0:k-1}}^\top, w_k^\top, \mathbf{0}_{U_{k+1:L-1}}^\top)$ finishes the proof. \square

Finally, we point out some important properties of N_k .

Lemma 24. *The matrix N_k (defined in [Eq. \(26\)](#)) is positive definite and satisfies*

$$N_k \preceq D_k^{-1} \quad \text{and} \quad N_k \preceq C_k^{-1} + P_k^\top N_{k-1} P_k.$$

Proof. The fact that N_k is positive definite is directly implied by [Lemma 23](#). To prove the rest of the statement, we first apply Woodbury matrix identity to write W_k as

$$\begin{aligned} W_k &= (D_k^{-1} + C_k^{-1} + \tilde{P}_k^\top M_{k-1}^{-1} \tilde{P}_k)^{-1} = D_k - D_k(D_k + (C_k^{-1} + \tilde{P}_k^\top M_{k-1}^{-1} \tilde{P}_k)^{-1})^{-1} D_k \\ &= D_k - D_k(D_k + (C_k^{-1} + P_k^\top N_{k-1}^{-1} P_k)^{-1})^{-1} D_k. \end{aligned}$$

Plugging this back into the definition of N_k gives:

$$N_k = D_k^{-1} - D_k^{-1} W_k D_k^{-1} = (D_k + (C_k^{-1} + P_k^\top N_{k-1}^{-1} P_k)^{-1})^{-1},$$

which shows $N_k \preceq D_k^{-1}$ and $N_k \preceq C_k^{-1} + P_k^\top N_{k-1} P_k$. \square

D.2 Bounding the stability term

With the tools from the previous section, we are now ready to bound the stability term. We will use the following lemma to relate the quadratic form of H^{-1} to only its diagonal entries.

Lemma 25. *If $M \in \mathbb{R}^{d \times d}$ is a positive semi-definite matrix, then for any $w \in \mathbb{R}^d$ with non-negative coordinates, we have*

$$w^\top M w \leq \left(\sum_{j=1}^d w(j) \right) \sum_{i=1}^d M(i, i) w(i).$$

Proof. Since M is positive semi-definite, we have for any i, j , $(e_i - e_j)^\top M (e_i - e_j) \geq 0$, which implies

$$M(i, j) = M(j, i) \leq \frac{M(i, i) + M(j, j)}{2}.$$

Therefore,

$$w^\top M w = \sum_{i,j} w(i) M(i, j) w(j) \leq \frac{1}{2} \sum_{i,j} w(i) (M(i, i) + M(j, j)) w(j) = \left(\sum_{j=1}^d w(j) \right) \sum_{i=1}^d M(i, i) w(i),$$

where we use the nonnegativity of w in the second step as well. \square

We now bound $\|\hat{\ell}_t\|_{H^{-1}}$ in terms of the diagonal entries of N_k .

Lemma 26. *[Algorithm 1](#) guarantees*

$$\mathbb{E} \left[\|\hat{\ell}_t\|_{H^{-1}} \right] \leq L \mathbb{E} \left[\sum_{k=0}^{L-1} \sum_{(s,a) \in U_k} \frac{N_k((s,a), (s,a))}{q_t(s,a)} \right].$$

Proof. Recall the definition of $\widehat{\ell}_t$: $\widehat{\ell}_t(s, a) = \frac{\ell_t(s, a)}{q_t(s, a)} \mathbb{I}\{s, a\}$ where we use the shorthand $\mathbb{I}\{s, a\} = \mathbb{I}\{s_{k(s)} = s, a_{k(s)} = a\}$. Therefore, we have

$$\begin{aligned} \mathbb{E} \left[\left\| \widehat{\ell}_t \right\|_{H^{-1}} \right] &= \mathbb{E} \left[\left\| \sum_{s, a} \sum_{s', a'} \frac{H^{-1}((s, a), (s', a'))}{q_t(s, a)q_t(s', a')} \mathbb{I}\{s, a\} \ell_t(s, a) \mathbb{I}\{s', a'\} \ell_t(s', a') \right\| \right] \\ &\leq \mathbb{E} \left[\left(\sum_{s', a'} \mathbb{I}\{s', a'\} \ell_t(s', a') \right) \sum_{s, a} \left(\frac{H^{-1}((s, a), (s, a))}{q_t(s, a)^2} \right) \mathbb{I}\{s, a\} \ell_t(s, a) \right], \end{aligned} \quad (27)$$

where in the last step we use [Lemma 25](#) with M being a matrix in the same shape of H and with entry $M((s, a), (s', a')) = \frac{H^{-1}((s, a), (s', a'))}{q_t(s, a)q_t(s', a')}$ (which is clearly positive definite). Using the fact $\sum_{s', a'} \mathbb{I}\{s', a'\} \ell_t(s', a') \leq L$, $H^{-1}((s, a), (s, a)) \geq 0$, and $\ell_t(s, a) \in [0, 1]$, we continue with

$$\begin{aligned} \mathbb{E} \left[\left\| \widehat{\ell}_t \right\|_{H^{-1}} \right] &\leq L \mathbb{E} \left[\sum_{s, a} \left(\frac{H^{-1}((s, a), (s, a))}{q_t(s, a)^2} \right) \mathbb{I}\{s, a\} \right] \\ &= L \mathbb{E} \left[\sum_{s, a} \frac{H^{-1}((s, a), (s, a))}{q_t(s, a)} \right] \\ &\leq L \mathbb{E} \left[\sum_{k=0}^{L-1} \sum_{(s, a) \in U_k} \frac{N_k((s, a), (s, a))}{q_t(s, a)} \right], \end{aligned}$$

where in the last step we use [Lemma 23](#). □

Next, we continue to bound the term involving N_k using the following lemma.

Lemma 27. *Algorithm 1 guarantees*

$$\sum_{k=0}^{L-1} \sum_{(s, a) \in U_k} \frac{N_k((s, a), (s, a))}{q_t(s, a)} \leq 8eL \left(\sqrt{L} + \frac{1}{\alpha L} \right) \sum_{s \neq s_L} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)}$$

for any mapping π from S to A .

Proof. For notational convenience, define $R(s, a) = N_{k(s)}((s, a), (s, a))$. We first prove that for any $k = 1, \dots, L-1$,

$$\sum_{(s, a) \in U_k} \frac{R(s, a)}{q_t(s, a)} \leq 8 \left(\sqrt{L} + \frac{1}{\alpha L} \right) \sum_{s \in S_k} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)} + \left(1 + \frac{1}{L} \right) \sum_{(s', a') \in U_{k-1}} \frac{R(s', a')}{q_t(s', a')}, \quad (28)$$

and for $k = 0$,

$$\sum_{a \in A} \frac{R(s_0, a)}{q_t(s_0, a)} \leq 8 \left(\sqrt{L} + \frac{1}{\alpha L} \right) \sum_{a \neq \pi(s)} \sqrt{q(s_0, a)}, \quad (29)$$

where π is any mapping from S to A . Indeed, repeatedly applying (28) and using the fact $(1 + 1/L)^L \leq e$ show

$$\sum_{(s, a) \in U_k} \frac{R(s, a)}{q_t(s, a)} \leq 8e \left(\sqrt{L} + \frac{1}{\alpha L} \right) \sum_{l=0}^k \sum_{s \in S_l} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)},$$

and summing over k finishes the proof.

To prove [Eq. \(29\)](#), note that by definition, when $s = s_0$ we have

$$R(s, a) = \frac{4}{q_t(s, a)^{-3/2} + \alpha(q_t(s) - q_t(s, a))^{-3/2}} \leq 4 \min \left\{ q_t(s, a)^{3/2}, \frac{(q_t(s) - q_t(s, a))^{3/2}}{\alpha} \right\}.$$

Now consider two cases, if $\frac{q_t(s) - q_t(s, \pi(s))}{q_t(s, \pi(s))} \leq \frac{1}{L}$, then

$$\frac{R(s, \pi(s))}{q_t(s, \pi(s))} \leq \frac{4}{\alpha L} \sqrt{\sum_{a \neq \pi(s)} q_t(s, a)} \leq \frac{4}{\alpha L} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)}.$$

On the other hand, if $\frac{q_t(s) - q_t(s, \pi(s))}{q_t(s, \pi(s))} > \frac{1}{L}$, then $q_t(s, \pi(s)) \leq L(q_t(s) - q_t(s, \pi(s)))$ and

$$\frac{R(s, \pi(s))}{q_t(s, \pi(s))} \leq 4\sqrt{q_t(s, \pi(s))} \leq 4\sqrt{L} \sqrt{q_t(s) - q_t(s, \pi(s))} \leq 4\sqrt{L} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)}.$$

Combining the two cases and also the fact $\sum_{a \neq \pi(s)} \frac{R(s, a)}{q_t(s, a)} \leq 4 \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)}$ proves Eq. (29).

It remains to prove Eq. (28). First, using the fact $N_k \preceq D_k^{-1}$ from Lemma 24, we again have

$$R(s, a) \leq 4q_t(s, a)^{3/2}. \quad (30)$$

At the same time, using another fact $N_k \preceq C_k^{-1} + P_k^\top N_{k-1} P_k$ from Lemma 24 and shorthand $R(s, a, s', a') \triangleq N_{k(s)}((s, a), (s', a'))$, we have

$$\begin{aligned} R(s, a) &\leq \frac{4(q(s) - q(s, a))^{3/2}}{\alpha} + \sum_{(s_1, a_1), (s_2, a_2) \in U_{k-1}} P(s|s_1, a_1) P(s|s_2, a_2) R(s_1, a_1, s_2, a_2) \\ &= \frac{4(q(s) - q(s, a))^{3/2}}{\alpha} \\ &\quad + \sum_{(s_1, a_1), (s_2, a_2) \in U_{k-1}} P(s|s_1, a_1) q_t(s_1, a_1) P(s|s_2, a_2) q_t(s_2, a_2) \frac{R(s_1, a_1, s_2, a_2)}{q_t(s_1, a_1) q_t(s_2, a_2)} \\ &\leq \frac{4(q(s) - q(s, a))^{3/2}}{\alpha} \\ &\quad + \left(\sum_{(s_2, a_2) \in U_{k-1}} P(s|s_2, a_2) q_t(s_2, a_2) \right) \sum_{(s_1, a_1) \in U_{k-1}} P(s|s_1, a_1) q_t(s_1, a_1) \frac{R(s_1, a_1)}{q(s_1, a_1)^2} \\ &= \frac{4(q(s) - q(s, a))^{3/2}}{\alpha} + q_t(s) \sum_{(s_1, a_1) \in U_{k-1}} P(s|s_1, a_1) \frac{R(s_1, a_1)}{q_t(s_1, a_1)}, \end{aligned}$$

where the second inequality is by applying Lemma 25 again, with $M \in \mathbb{R}^{U_{k-1} \times U_{k-1}}$ such that $M((s_1, a_1), (s_2, a_2)) = \frac{R(s_1, a_1, s_2, a_2)}{q_t(s_1, a_1) q_t(s_2, a_2)}$ (which is positive definite by Lemma 24). Again, we fix s and consider two cases. First, if $\frac{q_t(s) - q_t(s, \pi(s))}{q_t(s, \pi(s))} \leq \frac{1}{L}$, then

$$\begin{aligned} \frac{R(s, \pi(s))}{q_t(s, \pi(s))} &\leq \frac{4(q_t(s) - q_t(s, \pi(s)))^{3/2}}{\alpha q_t(s, \pi(s))} + \frac{q_t(s)}{q_t(s, \pi(s))} \sum_{(s_1, a_1) \in U_{k-1}} P(s|s_1, a_1) \frac{R(s_1, a_1)}{q_t(s_1, a_1)} \\ &\leq \frac{4(q_t(s) - q_t(s, \pi(s)))}{\alpha q_t(s, \pi(s))} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)} + \left(1 + \frac{1}{L}\right) \sum_{(s_1, a_1) \in U_{k-1}} P(s|s_1, a_1) \frac{R(s_1, a_1)}{q_t(s_1, a_1)} \\ &\leq \frac{4}{\alpha L} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)} + \left(1 + \frac{1}{L}\right) \sum_{(s_1, a_1) \in U_{k-1}} P(s|s_1, a_1) \frac{R(s_1, a_1)}{q_t(s_1, a_1)}. \end{aligned}$$

On the other hand, if $\frac{q_t(s) - q_t(s, \pi(s))}{q_t(s, \pi(s))} > \frac{1}{L}$ and thus $q_t(s, \pi(s)) \leq L(q_t(s) - q_t(s, \pi(s)))$, then using Eq. (30) we have

$$\frac{R(s, \pi(s))}{q_t(s, \pi(s))} \leq 4\sqrt{q_t(s, \pi(s))} \leq 4\sqrt{L} \sum_{a \neq \pi(s)} q_t(s, a) \leq 4\sqrt{L} \sum_{a \neq \pi(s)} \sqrt{q_t(s, a)}.$$

Combining the two cases and also $\sum_{a \neq \pi(s)} \frac{R(s,a)}{q_t(s,a)} \leq 4 \sum_{a \neq \pi(s)} \sqrt{q_t(s,a)}$ (using Eq. (30) again) leads to

$$\sum_a \frac{R(s,a)}{q_t(s,a)} \leq 8 \left(\sqrt{L} + \frac{1}{\alpha L} \right) \sum_{a \neq \pi(s)} \sqrt{q_t(s,a)} + \left(1 + \frac{1}{L} \right) \sum_{(s_1, a_1) \in U_{k-1}} P(s|s_1, a_1) \frac{R(s_1, a_1)}{q_t(s_1, a_1)}.$$

Finally, we sum over all $s \in S_k$ and obtain

$$\begin{aligned} & \sum_{(s,a) \in U_k} \frac{R(s,a)}{q_t(s,a)} \\ & \leq 8 \left(\sqrt{L} + \frac{1}{\alpha L} \right) \sum_{s \in S_k} \sum_{a \neq \pi(s)} \sqrt{q_t(s,a)} + \left(1 + \frac{1}{L} \right) \sum_{s \in S_k} \sum_{(s', a') \in U_{k-1}} P(s|s', a') \frac{R(s', a')}{q_t(s', a')} \\ & = 8 \left(\sqrt{L} + \frac{1}{\alpha L} \right) \sum_{s \in S_k} \sum_{a \neq \pi(s)} \sqrt{q_t(s,a)} + \left(1 + \frac{1}{L} \right) \sum_{(s', a') \in U_{k-1}} \frac{R(s', a')}{q_t(s', a')}. \end{aligned}$$

This proves Eq. (28) and thus finishes the proof. \square

We are now ready to finish the proof for Lemma 7.

Proof of Lemma 7. Combining Lemma 26 and Lemma 27, we prove

$$\mathbb{E} \left[\|\widehat{\ell}_t\|_{\nabla^{-2} \phi_H(q_t)}^2 \right] \leq \mathbb{E} \left[8eL^2 \left(\sqrt{L} + \frac{1}{\alpha \cdot L} \right) \sum_{s \neq s_L} \sum_{a \neq \pi(a)} \sqrt{q_t(s,a)} \right].$$

It thus remains to prove the other bound

$$\mathbb{E} \left[\|\widehat{\ell}_t\|_{\nabla^{-2} \phi_H(q_t)}^2 \right] \leq 4\sqrt{L|S||A|}.$$

This is simply by considering only the regular $1/2$ -Tsallis entropy part of the regularizer: $\phi_D(q) = -\sum_{s,a} \sqrt{q(s,a)}$. Specifically, by Lemma 9 we have $\nabla^2 \phi_H(q_t) \succeq \nabla^2 \phi_D(q_t)$, and thus

$$\begin{aligned} \mathbb{E} \left[\|\widehat{\ell}_t\|_{\nabla^{-2} \phi_H(q_t)}^2 \right] & \leq \mathbb{E} \left[\|\widehat{\ell}_t\|_{\nabla^{-2} \phi_D(q_t)}^2 \right] \\ & = \mathbb{E} \left[\sum_{s,a} \frac{4q_t(s,a)^{3/2}}{q_t(s,a)^2} \mathbb{I}\{s,a\} \ell_t(s,a)^2 \right] \\ & \leq 4\mathbb{E} \left[\sum_{s,a} \sqrt{q_t(s,a)} \right] \\ & \leq 4\sqrt{L|S||A|}, \end{aligned}$$

where the last step uses the Cauchy-Schwarz inequality. \square