From Dueling Bandits to Episodic MDPs: Learning with Adversarial Preference Feedback

Anonymous authors
Paper under double-blind review

Abstract

In dueling bandits, the learner compares two arms in each round and observes only which one is preferred. This paper asks what changes when the compared pair also affects the next state, as in a recommendation interface where the pair shown to a user affects the user's next context. We formulate this problem as a state-wise preference-based MDP, an episodic tabular MDP in which the learner selects a pair of arms at each visited state, observes only a noisy comparison between them, and moves to a next state that may depend on the pair. The preferences may change adversarially across episodes, and the loss of a pair is defined through Borda scores. We show that the minimax regret of this problem under known transitions is $\widetilde{\Theta}(\min\{HT, (H^3SK)^{1/3}T^{2/3}\})$, where H, S, K, and T are the numbers of steps per episode, states, arms, and episodes, and that this rate is achieved by an algorithm based on online linear optimization over the set of occupancy measures. If the pair did not affect the next state, each state would be a separate dueling bandit, and the $T^{2/3}$ term of the minimax regret would be smaller by a factor of $H^{1/3}$, demonstrating the cost of having action-dependent state transitions. We also give a policy optimization algorithm with closed-form updates, whose regret is $\widetilde{O}((H^4SK^3)^{1/3}T^{2/3})$ up to lower-order terms, and we extend both algorithms to unknown transitions.

1 Introduction

Reinforcement learning from preference feedback instead of numerical rewards has attracted increasing attention (Wirth et al., 2017; Christiano et al., 2017; Novoseller et al., 2020; Hejna & Sadigh, 2023; Saha et al., 2023; Wu & Sun, 2024). A basic model in this line of research is the dueling bandit problem (Yue & Joachims, 2009), where the learner compares two arms and observes only a binary outcome indicating which one is preferred. Using only this feedback, the learner aims to minimize the regret, that is, the cumulative difference between the losses of the compared arms and the loss of the best arm. Most dueling bandit algorithms assume the existence of a Condorcet winner (Urvoy et al., 2013; Zoghi et al., 2014; Komiyama et al., 2015), that is, an arm that wins in comparisons against every other arm, or of an underlying utility value for each arm (Ailon et al., 2014; Gajane et al., 2015).

Both assumptions fail when preferences are inconsistent or cyclic, since a Condorcet winner or an underlying utility need not exist. Recent work has thus turned to an alternative notion of optimality, the *Borda winner* (Urvoy et al., 2013; Jamieson et al., 2015; Ramamohan et al., 2016; Falahatgar et al., 2017; Heckel et al., 2018; Saha et al., 2021; Wu et al., 2024; Suk & Agarwal, 2024). The Borda winner is the arm with the highest win probability against a uniformly random opponent and always exists. Dueling bandits with a Borda winner can therefore handle arbitrary preferences, but they do not model how the compared arms affect future states. For example, a recommendation interface shows a user a pair of items and observes which one is clicked, and the pair shown affects the user's next context. This paper asks how to learn from comparisons when the compared pair also affects the future states.

1.1 Contributions of this paper

To answer this question, we propose a new online learning model of episodic MDPs with adversarial state-wise pairwise preference feedback, which we call *state-wise preference-based MDPs (state-wise PbMDPs for short)*. State-wise PbMDPs extend episodic MDPs with adversarial losses (Neu

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Table 1: Regret bounds established in this paper for state-wise PbMDPs with Borda scores. The notation $\widetilde{O}$ hides logarithmic factors. Each bound is of the form $\min\{HT, \cdot\}$, and the trivial term HT is omitted. OLO stands for online linear optimization. †Up to additive lower-order terms.

Setting	Result	Regret	Algorithm
Lower bound	Theorem 3	$\Omega\big((H^3SK)^{1/3}T^{2/3}\big)$	–
Known transitions	Theorem 1	$\widetilde{O}\big((H^3SK)^{1/3}T^{2/3}\big)$	OLO over polytope
	Theorem 2	$\widetilde{O}\big((H^4SK^3)^{1/3}T^{2/3}\big)^\dagger$	policy optimization
Unknown transitions	Theorem 8	$\widetilde{O}\big((H^3SK)^{1/3}T^{2/3} + HSK\sqrt{T}\big)^\dagger$	OLO over polytope
	Theorem 17	$\widetilde{O}\big((H^4SK^3)^{1/3}T^{2/3} + H^2SK\sqrt{T}\big)^\dagger$	policy optimization

et al., 2010; Zimin & Neu, 2013; Jin et al., 2020a; 2021) in the same way as dueling bandits extend multi-armed bandits. At each visited state, the learner selects a pair of arms as its action and observes only the outcome of a comparison between them, rather than a numerical loss. The preference probabilities may vary arbitrarily across episodes. We measure the loss of a pair by the Borda scores of its arms, which are well defined for arbitrary preferences. The model then contains adversarial dueling bandits with a Borda winner as a special case. To our knowledge, this is the first regret analysis for episodic MDPs with adversarial state-wise preference feedback.

Our first algorithm runs follow-the-regularized-leader (FTRL) with the negative Shannon entropy over the set of occupancy measures, following the online linear optimization approach to adversarial MDPs (Zimin & Neu, 2013; Jin et al., 2020a), and achieves a regret bound of $O((H^3SK\log(SK))^{1/3}T^{2/3})$ under known transitions, where H, S, K, and T denote the numbers of steps per episode, states, arms, and episodes, respectively (Theorem 1). The main difficulty is that the Borda scores at a state are estimated from a uniformly random comparison at that state, whereas the probability of visiting the state can be arbitrarily small, so that an estimator that divides by this probability has a large second moment. We resolve this by restricting the decision set of FTRL to the occupancy measures of policies that make a uniformly random comparison with a fixed probability at every state. The executed policy then makes these comparisons itself, and the resulting regret bound does not depend on the state-visit probabilities.

The first algorithm has to solve a convex program in every episode, which is computationally demanding, and this approach is difficult to extend to structured MDPs. These drawbacks have motivated policy optimization methods for adversarial MDPs (Shani et al., 2020; Luo et al., 2021; Dann et al., 2023; Li et al., 2026b). Our second algorithm follows this policy optimization framework, which maintains a distribution over pairs at each state and updates it by a multi-armed bandit algorithm whose losses correspond to an underlying Q-function. The main new ingredient is an estimator of the Q-function from preference feedback. It combines the losses of all pairs at a visited state, estimated from a single uniform comparison there, with the losses at later steps, estimated from the executed trajectory. Under known transitions, the algorithm achieves a regret bound of $\widetilde{O}((H^4SK^3)^{1/3}T^{2/3})$ up to lower-order terms (Theorem 2), which has the same dependence on T and S as our first bound and is worse by a factor of $H^{1/3}K^{2/3}$.

We also extend both algorithms to unknown transitions using confidence sets for the transition kernel. In this setting, we obtain a regret bound of $\widetilde{O}((H^3SK)^{1/3}T^{2/3} + HSK\sqrt{T})$ for global optimization and a bound of $\widetilde{O}((H^4SK^3)^{1/3}T^{2/3} + H^2SK\sqrt{T})$ for policy optimization, in both cases up to an additive term polylogarithmic in T (Theorems 8 and 17), and estimating the transitions thus adds only lower-order terms. The main body presents the algorithms and analyses for known transitions, and the unknown-transition settings are deferred to Appendices B.5 and D.

Finally, we prove that the regret bound of our first algorithm is optimal up to the logarithmic factor. We show that, for any learning algorithm, there exists a state-wise PbMDP with Borda scores on which the expected regret is $\Omega(\min\{HT, (H^3SK)^{1/3}T^{2/3}\})$ (Theorem 3), that is, $\Omega((H^3SK)^{1/3}T^{2/3})$ for $T \geq SK$. Hence, under the assumptions of Theorem 3, the minimax regret of state-wise PbMDPs with Borda scores under known transitions is $\widetilde{\Theta}(\min\{HT, (H^3SK)^{1/3}T^{2/3}\})$, and the leading term of the policy optimization bound is optimal in T and S. The proof requires a new instance in which the comparisons that reveal the Borda winner also move the learner to states with larger losses for the rest of the episode. When the pair

does not affect the next state, each state is a separate dueling bandit, and the minimax regret is $\widetilde{\Theta}(\min\{HT, (H^2SK)^{1/3}T^{2/3}\})$ (Appendix F), which is smaller by a factor of $H^{1/3}$ in the $T^{2/3}$ term, demonstrating the cost of having action-dependent state transitions. Table 1 summarizes our results. Experiments on small instances (Appendix G) show that the regret grows at a rate consistent with $T^{2/3}$ for global optimization and at a rate approaching $T^{2/3}$ for policy optimization as T increases, and that policy optimization has a shorter running time.

1.2 Related work

Our work relates to preference-based reinforcement learning (PbRL). We refer to the survey by Wirth et al. (2017) and, for recent developments, to Christiano et al. (2017); Novoseller et al. (2020); Xu et al. (2020); Hejna & Sadigh (2023); Saha et al. (2023); Wu & Sun (2024).

Many PbRL methods learn from comparisons of trajectories or trajectory segments (Christiano et al., 2017; Xu et al., 2020; Saha et al., 2023; Wu & Sun, 2024). In contrast, state-wise PbMDPs adopt *state-wise* comparisons. At each visited state, the learner selects a *pair* of arms and observes a preference outcome between them, and the dynamics may depend on the displayed pair of arms. The two formulations use different feedback and are not directly reducible to each other. Table 2 in Appendix A compares our interaction protocol with the trajectory-comparison formulation.

A natural example is an interactive ranking or recommendation interface (Joachims, 2002; Radlinski et al., 2008; Hofmann et al., 2016), where the state is the observed user context, the displayed pair is compared through the user's choice, and the pair shown can affect the subsequent context. This example motivates the state-wise formulation, which models the comparison of the displayed pair at the current state and its effect on the next state, and allows the preferences to vary across time and states in an inconsistent or nonstationary manner. Comparing actions at the same state has also been considered in reinforcement and imitation learning (Fürnkranz et al., 2012; Myers et al., 2023; Yang et al., 2023; Sekhari et al., 2023).

2 Preliminaries

Notation. For $n \in \mathbb{N}$, let $[n] = \{1, \dots, n\}$. For a vector x, x_i denotes its i-th element and $\|x\|_p$ its ℓ_p-norm for $p \in [1, \infty]$. For a finite set $\mathcal{K}$, $\Delta(\mathcal{K})$ denotes the set of all probability distributions over $\mathcal{K}$. We write $\mathbb{1}[\cdot]$ for the indicator function, $\mathcal{A}^{\mathcal{B}}$ for the set of all functions from $\mathcal{B}$ to $\mathcal{A}$, $\mathsf{Ber}(p)$ for the Bernoulli distribution with mean p, and $\mathsf{Unif}(\mathcal{A})$ for the uniform distribution over $\mathcal{A}$.

Tabular MDPs. We study a tabular MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, H, s_0)$, where $\mathcal{S}$ is a finite state space with $S = |\mathcal{S}|$, $\mathcal{A}$ is a finite action space with $A = |\mathcal{A}|$, $P\colon \mathcal{S} \times \mathcal{A} \to \Delta(\mathcal{S})$ is the transition kernel, H is the horizon, that is, the number of steps per episode, and $s_0 \in \mathcal{S}$ is the initial state. Here $P(s' \mid s, a)$ denotes the probability of transitioning to state s' when taking action a in state s.

Following existing studies (Zimin & Neu, 2013; Du et al., 2019; Jin et al., 2020a; 2021), we assume that $\mathcal{S}$ is partitioned into $(H+1)$ disjoint layers, $\mathcal{S} = \bigcup_{h=0}^{H} \mathcal{S}_h$ with $S_h = |\mathcal{S}_h|$, where $\mathcal{S}_0 = \{s_0\}$ and $\mathcal{S}_H = \{s_H\}$ are singletons and all intermediate layers $\mathcal{S}_1, \dots, \mathcal{S}_{H-1}$ are nonempty. Transitions proceed forward by one layer: for $h \in \{0, \dots, H-1\}$ and $(s, a) \in \mathcal{S}_h \times \mathcal{A}$, $\sum_{s' \in \mathcal{S}_{h+1}} P(s' \mid s, a) = 1$ and $P(s' \mid s, a) = 0$ for $s' \notin \mathcal{S}_{h+1}$. A general finite-horizon MDP can be unfolded into a layered MDP by augmenting the state with the time index, at the cost of increasing the state space by a factor of H. All bounds in this paper are stated in terms of the layered state space. For simplicity, we sometimes write $\mathcal{S}$ for $\mathcal{S} \setminus \{s_H\}$.

Policies and value functions. An episode starts at s_0 and lasts H steps. At step h, the learner observes the state s_h, selects an action a_h, and moves to $s_{h+1} \sim P(\cdot \mid s_h, a_h)$. A stochastic Markov policy $\pi \in \Pi = \Delta(\mathcal{A})^{\mathcal{S}}$ selects $a_h \sim \pi(\cdot \mid s_h)$. Let $h(s)$ denote the layer of s, that is, $h(s) \in \{0, \dots, H\}$ with $s \in \mathcal{S}_{h(s)}$. Given P, a function $\ell\colon \mathcal{S} \times \mathcal{A} \to [0, 1]$, and $\pi \in \Pi$, the value function $V^\pi(\cdot; \ell)$ and the Q-function $Q^\pi(\cdot, \cdot; \ell)$ are defined by $V^\pi(s; \ell) = \mathbb{E}\big[\sum_{k=h(s)}^{H-1} \ell(s_k, a_k) \;\big|\; s_{h(s)} = s, \pi, P\big]$ and $Q^\pi(s, a; \ell) = \ell(s, a) + \mathbb{E}\big[\sum_{k=h(s)+1}^{H-1} \ell(s_k, a_k) \;\big|\; s_{h(s)} = s,\, a_{h(s)} = a,\, \pi,\, P\big]$, where the expectation is over $a_k \sim \pi(\cdot \mid s_k)$ and $s_{k+1} \sim P(\cdot \mid s_k, a_k)$ for $k \geq h(s)$. We omit the dependence on P when it is clear from the context. The occupancy measure $q^\pi\colon \mathcal{S} \times \mathcal{A} \to [0, 1]$

of π is the probability of visiting (s, a) under π and P: $q^\pi(s,a) = \Pr[\exists h \in \{0,\dots,H-1\}, s_h = s, a_h = a \mid \pi, P]$, with $q^\pi(s) = \sum_{a\in\mathcal{A}} q^\pi(s,a)$. Let $\Omega = \{q^\pi : \pi \in \Pi\}$ denote the set of all occupancy measures, which forms a polytope with $\Theta(S)$ linear constraints (Puterman, 2014), and let $\pi^q \in \Pi$ be the Markov policy induced by $q \in \Omega$, given by $\pi^q(a \mid s) = q(s,a)/q(s)$.

3 State-wise preference-based MDPs and the Borda score

This section introduces state-wise preference-based MDPs (state-wise PbMDPs) and then focuses on the case where the loss functions are determined by Borda scores.

State-wise preference-based MDPs. Formally, a state-wise PbMDP is specified by a tabular MDP $\mathcal{M}$ with action space $\mathcal{A} = [K] \times [K]$ for the set of arms $[K] = \{1,\dots,K\}$, where $K \geq 2$, and an adversarial sequence $\{\mathsf{P}_t\}_{t=1}^T$ of preference functions $\mathsf{P}_t : \mathcal{S}\times\mathcal{A} \to [0,1]$ satisfying $\mathsf{P}_t(s,(a^{\mathrm{L}},a^{\mathrm{R}})) = 1 - \mathsf{P}_t(s,(a^{\mathrm{R}},a^{\mathrm{L}}))$ for every $s \in \mathcal{S}$ and every pair $a = (a^{\mathrm{L}},a^{\mathrm{R}}) \in \mathcal{A}$. Here $\mathsf{P}_t(s,(i,j))$ is the probability that arm i is preferred to arm j at state s in episode t. The incurred loss $\ell_t : \mathcal{S}\times\mathcal{A} \to [0,1]$ is determined by P_t, and in all main results below it is the Borda-induced loss defined in (1).

The interaction protocol proceeds as follows:

Protocol. For each episode $t = 1,\dots,T$:

1. Environment chooses a preference function $\mathsf{P}_t : \mathcal{S}\times\mathcal{A} \to [0,1]$;
2. Learner selects a policy π_t based on past observations without knowing P_t;
3. Learner executes policy π_t to obtain a trajectory $\{(s_{t,h}, a_{t,h}, o_t(s_{t,h}, a^{\mathrm{L}}_{t,h}, a^{\mathrm{R}}_{t,h}))\}_{h=0}^{H-1} \in (\mathcal{S}\times\mathcal{A}\times\{0,1\})^H$. Specifically, for each step $h = 0,\dots,H-1$, the learner selects an arm pair $a_{t,h} = (a^{\mathrm{L}}_{t,h}, a^{\mathrm{R}}_{t,h}) \sim \pi_t(\cdot \mid s_{t,h})$, moves to $s_{t,h+1} \sim P(\cdot \mid s_{t,h}, a_{t,h})$, incurs the loss $\ell_t(s_{t,h}, a_{t,h})$, and observes the preference feedback $o_t(s_{t,h}, a^{\mathrm{L}}_{t,h}, a^{\mathrm{R}}_{t,h}) \sim \mathsf{Ber}(\mathsf{P}_t(s_{t,h}, a_{t,h}))$.

We assume an oblivious adversary, that is, the entire sequence $\{\mathsf{P}_t\}_{t=1}^T$ is fixed before the first episode and does not depend on the learner's randomization, the transition randomness, or the preference-feedback noise. The learner aims to minimize the *regret* $\mathsf{Reg}_T = \sum_{t=1}^T V^{\pi_t}(s_0;\ell_t) - \sum_{t=1}^T V^{\pi^*}(s_0;\ell_t)$, where $\pi^* \in \arg\min_{\pi\in\Pi} \sum_{t=1}^T V^{\pi}(s_0;\ell_t)$. We study both the known-transition setting, in which the learner has full access to P, and the unknown-transition setting, in which the learner infers P from interactions.

Borda score. Given P_t, the (shifted) Borda score $b_t : \mathcal{S}\times[K] \to [-1,0]$ and the loss function $\ell_t : \mathcal{S}\times\mathcal{A} \to [0,1]$ are defined as

$$b_t(s,i) = \tfrac{1}{K}\sum_{j=1}^K (\mathsf{P}_t(s,(i,j)) - 1)\,, \quad \ell_t(s,a) = -\tfrac{1}{2}\big(b_t(s,a^{\mathrm{L}}) + b_t(s,a^{\mathrm{R}})\big)\,. \tag{1}$$

The shifted Borda score is the probability that arm i is preferred to a uniformly random opponent at state s, shifted by -1 for analytical convenience. Note that state-wise PbMDPs with Borda scores reduce to dueling bandits with a Borda winner when $H = 1$, that is, when the MDP has a single nonterminal state and a single step. Even in this much simpler multi-armed bandit setting, different notions of a winner lead to different problems that require different algorithms (Zoghi et al., 2014; Ramamohan et al., 2016; Dudík et al., 2015; Saha et al., 2021). We therefore focus on the Borda winner, which is always defined, and leave other preference-based objectives, such as those based on the Condorcet winner, for future work.

Additional notation. Let $\mathbb{I}_t(s,a) := \mathbb{1}[\exists h \in \{0,\dots,H-1\}, (s_{t,h}, a_{t,h}) = (s,a)]$ indicate whether (s,a) is visited in episode t, and $\mathbb{I}_t(s) = \sum_{a\in\mathcal{A}} \mathbb{I}_t(s,a)$. We use $\mathbb{E}_t[\cdot]$ for the expectation conditioned on all observations before episode t.

4 Global optimization

This section presents Algorithm 1, a global optimization algorithm for state-wise PbMDPs with Borda scores. We first describe the algorithm and its regret bound for known transitions, and then state the regret bound for unknown transitions.

Algorithm 1: Global optimization under known transitions

Input: Learning rate $\eta > 0$ and exploration rate $\rho \in (0, 1/2]$

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

2 Compute $\widetilde{q}_t$ by FTRL over Ω_ρ and set $\pi_t = \pi^{\widetilde{q}_t}$:

$$\widetilde{q}_t \in \arg\min_{q \in \Omega_\rho} \left\{ \left\langle q, \sum_{\tau=1}^{t-1} \widehat{\ell}_\tau \right\rangle - \frac{1}{\eta} \mathrm{H}(q) \right\}, \quad \mathrm{H}(q) = \sum_{s,a} q(s,a) \log\left(\frac{1}{q(s,a)}\right). \tag{2}$$

3 Compute the residual policy: $\pi_t^{\mathrm{r}}(a \mid s) = (\pi_t(a \mid s) - \rho/A)/(1-\rho)$. (3)

4 **for** $h = 0, \dots, H-1$ **do**

5 At $s = s_{t,h}$, sample $Z_t(s) \sim \mathsf{Ber}(\rho)$, and select $a_{t,h}$ uniformly from $\mathcal{A}$ if $Z_t(s) = 1$ and from $\pi_t^{\mathrm{r}}(\cdot \mid s)$ otherwise;

6 Observe $o_t(s, a_{t,h}^{\mathrm{L}}, a_{t,h}^{\mathrm{R}})$ and $s_{t,h+1}$;

7 Set $Z_t(s) = 0$ at unvisited states, compute $\widehat{b}_t$ by (5), and set

$$\widetilde{b}_t(s,i) = \frac{\mathbb{I}_t(s) Z_t(s) \widehat{b}_t(s,i)}{\rho \widetilde{q}_t(s)}, \qquad \widehat{\ell}_t(s,a) = -\frac{\widetilde{b}_t(s, a^{\mathrm{L}}) + \widetilde{b}_t(s, a^{\mathrm{R}})}{2}. \tag{4}$$

4.1 Algorithm design under known transitions

Reduction to online linear optimization. The regret can be written as $\mathbb{E}[\mathsf{Reg}_T] = \mathbb{E}\big[\sum_{t=1}^T \langle q^{\pi_t} - q^{\pi^*}, \ell_t \rangle\big]$, where ℓ_t is the Borda-induced loss in (1). Hence, it suffices to run FTRL over the set of occupancy measures Ω, following the online linear optimization approach to adversarial MDPs (Zimin & Neu, 2013; Jin et al., 2020a). Unlike in adversarial MDPs with loss feedback, the learner observes no entry of ℓ_t, and ℓ_t has to be estimated from the preference feedback.

Borda score estimation. We first describe how to estimate the Borda scores from the preference feedback. For a visited state s, let $a_{t,s}^{\mathrm{L}}$ and $a_{t,s}^{\mathrm{R}}$ be the arms compared at s in episode t, and define

$$\widehat{b}_t(s,i) = K \mathbb{1}\big[a_{t,s}^{\mathrm{L}} = i\big]\big(o_t(s, a_{t,s}^{\mathrm{L}}, a_{t,s}^{\mathrm{R}}) - 1\big) \qquad \text{for each } i \in [K]. \tag{5}$$

We set $\widehat{b}_t(s, \cdot) = 0$ when s is not visited. If the compared pair is drawn uniformly from $[K] \times [K]$, then $\widehat{b}_t(s,i)$ is an unbiased estimator of $b_t(s,i)$ (Lemma 4). Since the estimates in (5) for all $i \in [K]$ are computed from the same comparison, a single uniform comparison at s yields unbiased estimates of the Borda scores of all K arms, and hence of the losses $\ell_t(s,a)$ of all A pairs at s.

The question is how to make such uniform comparisons without losing control of the state-visit probabilities. In adversarial dueling bandits, Saha et al. (2021) draw the two arms independently from a distribution over arms and estimate the Borda scores by inverse weighting with respect to that distribution. In a state-wise PbMDP, the transition depends on the ordered pair, and drawing the two arms independently from the marginals of a policy would in general change the next-state distribution. For example, a policy that selects $(1,1)$ and $(2,2)$ with probability $1/2$ each never selects $(1,2)$, whereas independent draws from its marginals select it with probability $1/4$. We instead keep the decision variable as an occupancy measure over ordered pairs and build the uniform comparisons into the policy returned by the optimization problem.

FTRL with uniform comparisons. For $\rho \in (0, 1/2]$, we define the restricted set of occupancy measures $\Omega_\rho = \{q \in \Omega : q(s,a) \geq \rho q(s)/A \text{ for all } (s,a) \in \mathcal{S} \times \mathcal{A}\}$. The constraints are linear and require only that, conditional on a visit to s, every pair is selected with probability at least ρ/A. The set Ω_ρ is a convex polytope, and it is nonempty since it contains the occupancy measure of the mixture $(1-\rho)\pi + (\rho/A)\mathbf{1}$ of any policy π with the uniform policy.

In each episode t, the algorithm computes $\widetilde{q}_t$ by FTRL over Ω_ρ with the negative Shannon entropy regularizer, as in (2), and executes the induced policy $\pi_t = \pi^{\widetilde{q}_t}$ (Line 2). To make the uniform comparisons explicit, it also computes the residual policy π_t^{r} in (3), which satisfies $\pi_t(a \mid s) = (1-\rho)\pi_t^{\mathrm{r}}(a \mid s) + \rho/A$ (Line 3). The residual probabilities are nonnegative and sum to one since $\pi_t(a \mid s) \geq \rho/A$, and at a state with $\widetilde{q}_t(s) = 0$ both policies are defined to be uniform. At each visited state s, the algorithm samples $Z_t(s) \sim \mathsf{Ber}(\rho)$ and selects the pair uniformly from $\mathcal{A}$ if $Z_t(s) = 1$ and from $\pi_t^{\mathrm{r}}(\cdot \mid s)$ otherwise (Line 5). This implements π_t exactly, and $Z_t(s) = 1$ indicates that a uniform comparison is made at s.

After the episode, the algorithm sets $Z_t(s) = 0$ at unvisited states and computes the estimates $\widetilde{b}_t$ and $\widehat{\ell}_t$ in (4), where the first estimate is defined as zero if its denominator is zero (Line 7). Since $\rho\widetilde{q}_t(s)$ is exactly the probability that a uniform comparison is made at s in episode t, both estimates are unbiased at every reachable state (Lemma 5).

4.2 Regret upper bounds

The above algorithm achieves the following bound, whose proof is given in Appendix B.

Theorem 1. *Suppose that the transition kernel is known. With* $\eta = (\log(SA)/(T\sqrt{SK}))^{2/3}$ *and* $\rho = \min\{1/2, \sqrt{\eta SK}/H\}$, *Algorithm 1 satisfies* $\mathbb{E}[\mathsf{Reg}_T] = O(\min\{HT, (H^3SK\log(SK))^{1/3}T^{2/3}\})$.

For $H = 1$, state-wise PbMDPs with Borda scores reduce to adversarial dueling bandits with a Borda winner, and the bound in Theorem 1 becomes $O((K\log K)^{1/3}T^{2/3})$, which recovers the minimax regret of Saha et al. (2021, Theorems 2 and 16) up to the logarithmic factor. Section 6 shows that the bound is minimax optimal up to the logarithmic factor for general H as well.

Technical and proof overview. With loss feedback, the losses of the visited pairs are observed, and FTRL over the set of occupancy measures needs no explicit exploration. With preference feedback, the Borda scores at s are estimated from a uniform comparison at s, which the executed policy makes with probability $\rho\widetilde{q}_t(s)$, and the estimator in (4) divides by this probability. Since the estimate in (5) has second moment at most K under a uniform comparison (Lemma 4), we have $\mathbb{E}_t[\widetilde{b}_t(s,i)^2] \le K/(\rho\widetilde{q}_t(s))$ (Lemma 5), which is large at rarely visited states. The restriction to Ω_ρ lets the algorithm execute $\pi^{\widetilde{q}_t}$ without further mixing it with the uniform policy, so that the occupancy measure of the executed policy is the FTRL iterate $\widetilde{q}_t$ and the factor $1/\widetilde{q}_t(s)$ is cancelled by the weight $\widetilde{q}_t(s,a)$ in the FTRL bound, as in the loss-feedback case. This can be confirmed in the following proof sketch.

Since q^{π^*} need not lie in Ω_ρ, we compare with the mixture $(1-\rho)\pi^* + (\rho/A)\mathbf{1}$, whose occupancy measure q^*_ρ lies in Ω_ρ, at a cost of $\rho H^2 T$ (Lemma 7). Since $\widehat{\ell}_t$ is unbiased, the regret against q^*_ρ equals $\mathbb{E}[\sum_t \langle \widetilde{q}_t - q^*_\rho, \widehat{\ell}_t\rangle]$, and the standard analysis of FTRL over Ω_ρ gives $\sum_t \langle \widetilde{q}_t - q^*_\rho, \widehat{\ell}_t\rangle \le H\log(SA)/\eta + (\eta/2)\sum_t\sum_{s,a}\widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2$ (Lemma 6). The weight $\widetilde{q}_t(s,a)$ in the second term is the probability of visiting (s,a) under the executed policy, and hence

$$\mathbb{E}_t\left[\sum_{s,a}\widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2\right] \le \sum_{s,a}\widetilde{q}_t(s,a)\,(K/\rho\widetilde{q}_t(s)) \le SK/\rho,$$

where the first inequality follows from $\mathbb{E}_t[\widetilde{b}_t(s,i)^2] \le K/(\rho\widetilde{q}_t(s))$ and the second inequality from $\sum_a \widetilde{q}_t(s,a) = \widetilde{q}_t(s)$. The regret is thus at most $H\log(SA)/\eta + \eta SKT/(2\rho) + \rho H^2 T$, and choosing η and ρ to balance the three terms gives the bound in Theorem 1.

Unknown transitions. When the transition kernel is unknown, we combine Algorithm 1 with confidence sets for P as in Jin et al. (2020a). The resulting algorithm runs the FTRL update over the set of occupancy measures that are consistent with the current confidence set, and divides the Borda estimates by an upper confidence bound on the state-visit probability instead of by $\widetilde{q}_t(s)$. The algorithm and its proof are given in Appendix B.5, and it achieves $\mathbb{E}[\mathsf{Reg}_T] = \widetilde{O}(\min\{HT, (H^3SK)^{1/3}T^{2/3} + HSK\sqrt{T} + HS^4K^2\})$, where the last two terms are the cost of estimating the transitions (Theorem 8).

5 Policy optimization

The convex program over Ω_ρ in the last section can be solved in polynomial time by standard convex optimization methods, but it is a global optimization problem over all states and has to be solved in every episode. This section presents Algorithm 2, a policy optimization algorithm that avoids it, updating a distribution over pairs at each state in closed form and using only dynamic programming.

5.1 Algorithm design

Reduction to online linear optimization at each state. By the performance difference lemma (Kakade & Langford, 2002), we can decompose the regret as

$$\mathsf{Reg}_T = \sum_{s\in\mathcal{S}} q^{\pi^*}(s)\sum_{t=1}^{T}\langle \pi_t(\cdot \mid s) - \pi^*(\cdot \mid s), Q^{\pi_t}(s,\cdot;\ell_t)\rangle. \tag{6}$$

Hence it suffices to solve an online linear optimization problem over $\Delta(\mathcal{A})$ with loss vectors $Q^{\pi_t}(s,\cdot;\ell_t)$ for each $s \in \mathcal{S}$. As in the previous section, these loss vectors must be estimated from a single trajectory per episode. We next describe the exploration scheme and the Q-function estimate.

Algorithm 2: Policy optimization under known transitions

Input: Learning rate $\eta > 0$, exploration rate $\gamma \in (0, 1/2]$, and $\delta = 2\eta(K/\gamma + H)$

1 Set $\widetilde{\pi}_1(a \mid s) = 1/A$ for every nonterminal (s, a);

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

3 Set π_t and compute q_t under P as follows:

$$\pi_t(a \mid s) = (1-\gamma)\widetilde{\pi}_t(a \mid s) + \gamma/A, \qquad q_t = q^{\pi_t}. \tag{7}$$

4 Compute m_t and M_t, interpreting $m_t(s, a) = m_t(s)$:

$$m_t(s) = 2\delta H \sum_{a \in \mathcal{A}} \frac{\widetilde{\pi}_t(a \mid s)}{q_t(s, a) + \delta}, \qquad M_t(s, a) = Q^{\pi_t}(s, a; m_t). \tag{8}$$

5 **for** $h = 0, \dots, H-1$ **do**

6 At $s = s_{t,h}$, draw $Z_t(s) \sim \mathsf{Ber}(\gamma)$, and select $a_{t,h}$ uniformly from $\mathcal{A}$ if $Z_t(s) = 1$ and from $\widetilde{\pi}_t(\cdot \mid s)$ otherwise;

7 Observe $o_t(s, a^{\mathrm{L}}_{t,h}, a^{\mathrm{R}}_{t,h})$ and $s_{t,h+1}$;

8 Set $Z_t(s) = 0$ at unvisited states, compute $\widehat{b}_t$ by (5), and set

$$\widehat{\ell}_t(s, a) = -\frac{Z_t(s)}{2\gamma}\big(\widehat{b}_t(s, a^{\mathrm{L}}) + \widehat{b}_t(s, a^{\mathrm{R}})\big). \tag{9}$$

9 Compute $\widehat{L}_{t,h}$ and $\widehat{Q}_t(s, a)$ for all h and nonterminal (s, a):

$$\begin{aligned} \widehat{L}_{t,h} &= \sum_{j=h}^{H-1} \sum_{a' \in \mathcal{A}} \pi_t(a' \mid s_{t,j})\widehat{\ell}_t(s_{t,j}, a'), \\ \widehat{Q}_t(s, a) &= \frac{\mathbb{I}_t(s)\pi_t(a \mid s)\widehat{\ell}_t(s, a) + \mathbb{I}_t(s, a)\widehat{L}_{t,h(s)+1}}{q_t(s, a) + \delta}. \end{aligned} \tag{10}$$

10 Update each distribution over $\mathcal{A}$ by FTRL:

$$\widetilde{\pi}_{t+1}(a \mid s) \propto \widetilde{\pi}_t(a \mid s) \exp\big(-\eta(\widehat{Q}_t(s, a) - M_t(s, a))\big). \tag{11}$$

Exploration and loss estimation. At each state s, the algorithm maintains a distribution $\widetilde{\pi}_t(\cdot \mid s)$ over pairs, whose update is described at the end of this section. As in Section 4, uniform comparisons are used for estimation. At each visited state s, the algorithm samples $Z_t(s) \sim \mathsf{Ber}(\gamma)$ and selects the pair uniformly from $\mathcal{A}$ if $Z_t(s) = 1$ and from $\widetilde{\pi}_t(\cdot \mid s)$ otherwise (Line 6), so that the executed policy is the mixture π_t in (7).

The loss estimate $\widehat{\ell}_t$ in (9) is built from $\widehat{b}_t$ in (5). Since $Z_t(s)$ is drawn after reaching s, this estimate satisfies $\mathbb{E}_t[\widehat{\ell}_t(s, a) \mid \mathbb{I}_t(s) = 1] = \ell_t(s, a)$ for every pair a (Lemma 13). In contrast to the estimator in (4) for global optimization, $\widehat{\ell}_t$ is divided only by the probability γ and not by a state-visit probability, since the visit probabilities are accounted for in the Q-function estimate below.

Q-function estimation. We next describe the Q-function estimate $\widehat{Q}_t$ in (10). With loss feedback, the standard estimate of $Q^{\pi_t}(s, a; \ell_t)$ is the sum of the losses observed at and after (s, a) in the episode, divided by $q_t(s, a)$ when (s, a) is visited (Luo et al., 2021). With preference feedback, the losses are not observed, and the two parts of $Q^{\pi_t}(s, a; \ell_t)$ have to be estimated differently. The immediate loss $\ell_t(s, a)$ can be estimated from any visit to s, since a uniform comparison at s yields the estimates $\widehat{\ell}_t(s, \cdot)$ of the losses of all pairs at s. The losses at later steps can be estimated only at the visited states, and hence only from the trajectory after (s, a), which is available when a is executed at s. Accordingly, the first term of $\widehat{Q}_t(s, a)$ in (10) is weighted by $\mathbb{I}_t(s)$ and the second by $\mathbb{I}_t(s, a)$.

In the second term, the loss at a later visited state $s_{t,j}$ is estimated by the average $\sum_{a'} \pi_t(a' \mid s_{t,j})\widehat{\ell}_t(s_{t,j}, a')$ over pairs, since $\widehat{\ell}_t(s_{t,j}, a')$ is unbiased for every fixed pair a' but not for the executed pair $a_{t,j}$, from which $\widehat{b}_t$ is computed. This average is unbiased for the expected loss of π_t at $s_{t,j}$, as required by the Bellman recursion for Q^{π_t} (Lemma 14). The parameter $\delta > 0$ keeps the denominators in (10) positive.

Bonus term. We next describe the bonus term M_t in (8). It is needed because the decomposition (6) weights the estimation error of $\widehat{Q}_t(s, \cdot)$ by the comparator's visit probability $q^{\pi^*}(s)$, whereas the error grows as the learner's visit probability $q_t(s)$ decreases (Lemma 15), and $q^{\pi^*}(s)$ can be much larger than $q_t(s)$. To correct for this discrepancy, following Luo et al. (2021), the algorithm subtracts from

$\widehat{Q}_t$ the bonus $M_t(s,a) = Q^{\pi_t}(s,a;m_t)$, where $m_t(s)$ accounts for the bias and the second moment of $\widehat{Q}_t(s,\cdot)$. The bonus makes FTRL favor pairs leading to rarely visited states.

Policy update. Finally, at each state s, the algorithm updates $\widetilde{\pi}_t(\cdot \mid s)$ by FTRL with the negative Shannon entropy applied to the loss vectors $\widehat{Q}_\tau(s,\cdot) - M_\tau(s,\cdot)$, which has the closed form $\widetilde{\pi}_t(a \mid s) \propto \exp(-\eta \sum_{\tau<t}(\widehat{Q}_\tau(s,a) - M_\tau(s,a)))$ (Lines 1 and 10).

5.2 Regret upper bounds

The above algorithm achieves the following bound, whose proof is given in Appendix C.

Theorem 2. *Under known transitions, with the parameter choices specified in the proof, Algorithm 2 satisfies* $\mathbb{E}[\mathsf{Reg}_T] = O(\min\{HT, (H^4SK^3 \log K)^{1/3}T^{2/3} + (H^3 + H^{3/2}K\sqrt{S})\sqrt{T \log K}\})$.

Compared with Theorem 1, the leading term of the bound in Theorem 2 has the same dependence on T and S and is larger by the factor $H^{1/3}K^{2/3}$, and the second term is of lower order in T.

Technical and proof overview. As in the standard analysis for policy optimization (Luo et al., 2021), subtracting the bonus M_t from $\widehat{Q}_t$ lets the analysis weight the estimation error at each state by the learner's visit probability $q_t(s)$ instead of $q^{\pi^\star}(s)$ in (6). Lemma 16 in Appendix C states this in a form that applies to any nonnegative estimate $\widehat{Q}_t$ with the bonus in (8). For $\eta \le 1/(9H^2)$, suppose that $\mathbb{E}_t[\widehat{Q}_t(s,a)] \le Q^{\pi_t}(s,a;\ell_t)$ for every (s,a) and $\sum_a \widetilde{\pi}_t(a \mid s)\big(Q^{\pi_t}(s,a;\ell_t) - \mathbb{E}_t[\widehat{Q}_t(s,a)] + \eta\mathbb{E}_t[\widehat{Q}_t(s,a)^2]\big) \le m_t(s)$ for every s. Then $\mathbb{E}[\mathsf{Reg}_T] \le \gamma H^2 T + H\log A/\eta + 81\eta H^5 T + \mathbb{E}[\sum_t \sum_s q_t(s) m_t(s)]$. We state Lemma 16 in this form since, with loss feedback, $Q^{\pi_t}(s,a;\ell_t)$ is estimated without bias from the observed losses, whereas the estimate in (10) from preference feedback, as well as its unknown-transition version in Appendix D, has a different mean and a second moment that depends on K and γ.

It remains to evaluate $\mathbb{E}_t[\widehat{Q}_t]$ and $\mathbb{E}_t[\widehat{Q}_t^2]$ in the conditions of Lemma 16 for the estimate in (10). By (9) and $\mathbb{E}_t[\widehat{b}_t(s,i)^2 \mid \mathbb{I}_t(s) = 1, Z_t(s) = 1] \le K$ (Lemma 4), we have $\mathbb{E}_t[\widehat{\ell}_t(s,a)^2 \mid \mathbb{I}_t(s) = 1] \le K/\gamma$ (Lemma 13). Since $\widehat{L}_{t,h(s)+1}$ in (10) has at most H terms, each of which is nonnegative and has conditional mean at most one, we can show that $\mathbb{E}_t[\widehat{L}^2_{t,h(s)+1} \mid \mathbb{I}_t(s,a) = 1] \le HK/\gamma + H^2$ (Lemma 14), where the first term comes from the squared terms and the second from the cross terms. Lemma 15 then gives

$$\mathbb{E}_t\left[\widehat{Q}_t(s,a)\right] = \frac{q_t(s,a)}{q_t(s,a)+\delta} Q^{\pi_t}(s,a;\ell_t), \qquad \mathbb{E}_t\left[\widehat{Q}_t(s,a)^2\right] \le \frac{2H(K/\gamma + H)}{q_t(s,a)+\delta},$$

so that the conditions of Lemma 16 hold with $m_t(s)$ in (8) for $\delta = 2\eta(K/\gamma + H)$. The last term in the bound of Lemma 16 is thus at most $4\delta HSAT = 8\eta H(K/\gamma + H)SK^2T$, and choosing η and γ to balance the terms gives the bound in Theorem 2.

Unknown transitions. When the transition kernel is unknown, we combine Algorithm 2 with the same confidence sets for P as in Section 4, which changes only the denominator of $\widehat{Q}_t$ and the bonus. The unknown $q_t(s,a)$ in (10) is replaced by its upper confidence bound, and the bonus M_t takes the maximum over the confidence set at each step as in Luo et al. (2021), with an additional term in m_t for the bias caused by this replacement. The algorithm and its proof are given in Appendix D, and it achieves $\mathbb{E}[\mathsf{Reg}_T] = \widetilde{O}(\min\{HT, (H^4SK^3)^{1/3}T^{2/3} + H^2SK\sqrt{T} + H^2S^4K^2\})$, where the last two terms are again the cost of estimating the transitions (Theorem 17).

6 Regret lower bound

This section gives a regret lower bound for state-wise PbMDPs with Borda scores.

Theorem 3. *Suppose that $H \ge 3$, $K \ge 5$, and $S \ge 2H$. For any learning algorithm and any $T \ge 1$, there exists an episodic state-wise PbMDP with Borda scores and a known transition kernel such that* $\mathbb{E}[\mathsf{Reg}_T] = \Omega(\min\{HT, (H^3SK)^{1/3}T^{2/3}\})$.

Theorem 3, whose proof is given in Appendix E, matches the upper bound of Theorem 1 up to the factor $(\log(SK))^{1/3}$. Hence, under the assumptions of Theorem 3, the minimax regret of state-wise

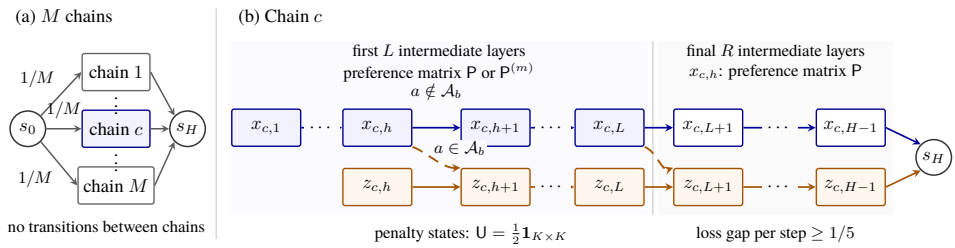


Figure 1: Lower-bound instance used to prove Theorem 3. Each chain is selected with probability $1/M$. At $x_{c,h}$ with $h \leq L$, a pair in $\mathcal{A}_b = \mathcal{A} \setminus (\mathcal{K}_g \times \mathcal{K}_g)$, where $\mathcal{K}_g = [\lfloor K/2 \rfloor]$, leads to the penalty state $z_{c,h+1}$, and any other pair leads to $x_{c,h+1}$. All other transitions do not depend on the pair.

PbMDPs with Borda scores under known transitions is $\widetilde{\Theta}(\min\{HT, (H^3SK)^{1/3}T^{2/3}\})$, and the leading term of the policy optimization bound of Theorem 2 is optimal in T and S. If the transition kernel does not depend on the selected pair, the learner faces a separate dueling bandit problem at each state, and the minimax regret is $\widetilde{\Theta}(\min\{HT, (H^2SK)^{1/3}T^{2/3}\})$ (Appendix F). The factor $H^{1/3}$ in the $T^{2/3}$ term is thus the cost of transitions that depend on the compared pair.

Proof overview. The instance used in the proof is shown in Figure 1, where $M = \Theta(S/H)$, $L = \Theta(H)$, and $R = \Theta(H)$. It is built so that every comparison that helps to identify the Borda winner moves the learner to the penalty states. At each state $x_{c,h}$ with $h \leq L$, the preference matrix is either P in (86), under which all arms in $\mathcal{K}_g = [\lfloor K/2 \rfloor]$ have the largest Borda score, or $\mathsf{P}^{(m)}$ with $m \in \mathcal{K}_g$, which differs from P only in that arm m beats each arm in $\mathcal{K}_b = [K] \setminus \mathcal{K}_g$ with probability larger by 2ε, so that (m, m) is the unique optimal pair with a loss gap of at least $\varepsilon/2$ (see (90)). The two matrices thus differ only on the pairs in $\mathcal{A}_b = \mathcal{A} \setminus (\mathcal{K}_g \times \mathcal{K}_g)$, which lead to the penalty states, where every pair has loss $1/2$, exceeding the loss of the optimal policy by at least $1/5$ at each of the final R layers (see (91)).

This instance gives the desired lower bound. Identifying the optimal index m among $\Theta(K)$ candidates requires $\Omega(K/\varepsilon^2)$ comparisons in $\mathcal{A}_b$, each costing $\Omega(H)$ instead of a constant, whereas a learner that does not identify m incurs regret $\Omega(\varepsilon T')$ at $x_{c,h}$ or $z_{c,h}$ in layer h, where $T' = T/M$ is the expected number of visits to the chain. Hence, each of the $LM = \Theta(S)$ states $x_{c,h}$ with $h \leq L$ contributes $\Omega(\min\{\varepsilon T', HK/\varepsilon^2\}) = \Omega((HK)^{1/3}(T')^{2/3})$ to the regret, where we choose $\varepsilon \asymp (HK/T')^{1/3}$ as in (92), and summing over these states gives $\Omega((H^3SK)^{1/3}T^{2/3})$.

7 Conclusion and Future Work

This paper studied online learning in episodic tabular MDPs with adversarial state-wise preference feedback. For losses induced by Borda scores, we gave a global optimization algorithm and a policy optimization algorithm under known and unknown transitions, together with a lower bound showing that the first algorithm is minimax optimal up to a logarithmic factor.

Our results have several limitations. First, the policy optimization bound is worse than the lower bound by a factor of $H^{1/3}K^{2/3}$ in the leading term, and under unknown transitions both of our bounds contain additional lower-order terms. Whether these terms can be removed is left open. Second, our guarantees are specific to Borda scores, since the estimators use that the Borda score is linear in the preference probabilities and that the loss of a pair is the negative average of the scores of its two arms. Other objectives such as Condorcet winners (Urvoy et al., 2013; Zoghi et al., 2014; Komiyama et al., 2015) and von Neumann winners (Dudík et al., 2015) deserve study. Third, the stochastic setting and best-of-both-worlds guarantees remain to be studied, and it is unclear whether the techniques for online tabular MDPs (Jin & Luo, 2020; Dann et al., 2023; Li et al., 2026a;b) can be combined with learning rates tailored to minimax regret of order $T^{2/3}$ (Tsuchiya & Ito, 2024). Finally, extensions to structured MDPs such as linear MDPs (Jin et al., 2020b; Cai et al., 2020; Neu & Olkhovskaya, 2021) and to adversarial trajectory-comparison feedback are left for future work.

AI USE STATEMENT

The problem formulation, the algorithms, and the original analysis in this paper were developed without generative AI tools. For the original manuscript, such tools were used only to proofread the English. In revising the paper, we used generative AI tools (GPT-5.6-Sol, GPT-6-Astra, and Fable-5) to improve the regret upper bound for the global optimization algorithm and the regret lower bound. We also used them to check the proofs in the appendix for errors, to implement the algorithms for the experiments, to create the figures, and to proofread and edit the text for readability. We have reviewed all AI-assisted work and we take responsibility for the final content of this work.

ETHICS STATEMENT

This work is theoretical and does not involve human subjects, personal data, or applications that raise ethical concerns.

REPRODUCIBILITY STATEMENT

The assumptions of all theoretical results are stated in the theorems, the algorithms are given as pseudocode (Algorithms 1–4), and complete proofs are given in Appendices B–F. The experimental setup, including the instances, the parameter selection, and the software versions, is described in Appendix G, and the code for the experiments is included in the supplementary material.

REFERENCES

Nir Ailon, Zohar Karnin, and Thorsten Joachims. Reducing dueling bandits to cardinal bandits. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pp. 856–864. PMLR, 2014.

Noga Alon, Nicolò Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback graphs: Beyond bandits. In *Proceedings of The 28th Conference on Learning Theory*, volume 40, pp. 23–35, 2015.

Gábor Bartók, Dávid Pál, and Csaba Szepesvári. Minimax regret of finite partial-monitoring games in stochastic environments. In *Proceedings of the 24th Annual Conference on Learning Theory*, volume 19, pp. 133–154, 2011.

Viktor Bengs, Róbert Busa-Fekete, Adil El Mesaoudi-Paul, and Eyke Hüllermeier. Preference-based online learning with dueling bandits: A survey. *Journal of Machine Learning Research*, 22(7): 1–108, 2021.

Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 1283–1294. PMLR, 2020.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, volume 30, pp. 4299–4307. Curran Associates, Inc., 2017.

Jonathan Colaço Carr, Prakash Panangaden, Doina Precup, and Benjamin Van Roy. Reinforcement learning with pairwise preferences in long-term decision problems. In *Proceedings of the 43rd International Conference on Machine Learning*. PMLR, 2026.

Christoph Dann, Chen-Yu Wei, and Julian Zimmert. Best of both worlds policy optimization. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 6968–7008. PMLR, 2023.

Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Episodic reinforcement learning in finite MDPs: Minimax lower bounds revisited. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, volume 132, pp. 578–598. PMLR, 2021.

Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudik, and John Langford. Provably efficient RL with rich observations via latent state decoding. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 1665–1674. PMLR, 2019.

Miroslav Dudík, Katja Hofmann, Robert E. Schapire, Aleksandrs Slivkins, and Masrour Zoghi. Contextual dueling bandits. In *Proceedings of The 28th Conference on Learning Theory*, volume 40, pp. 563–587. PMLR, 2015.

Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Online Markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009.

Moein Falahatgar, Yi Hao, Alon Orlitsky, Venkatadheeraj Pichapati, and Vaishakh Ravindrakumar. Maxing and ranking with few assumptions. In *Advances in Neural Information Processing Systems*, volume 30, pp. 7060–7070. Curran Associates, Inc., 2017.

Johannes Fürnkranz, Eyke Hüllermeier, Weiwei Cheng, and Sang-Hyeun Park. Preference-based reinforcement learning: A formal framework and a policy iteration algorithm. *Machine Learning*, 89:123–156, 2012.

Pratik Gajane, Tanguy Urvoy, and Fabrice Clérot. A relative exponential weighing algorithm for adversarial utility-based dueling bandits. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pp. 218–227. PMLR, 2015.

Reinhard Heckel, Max Simchowitz, Kannan Ramchandran, and Martin Wainwright. Approximate ranking from pairwise comparisons. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84, pp. 1057–1066. PMLR, 2018.

Joey Hejna and Dorsa Sadigh. Inverse preference learning: Preference-based RL without a reward function. In *Advances in Neural Information Processing Systems*, volume 36, pp. 18806–18827. Curran Associates, Inc., 2023.

Katja Hofmann, Lihong Li, and Filip Radlinski. Online evaluation for information retrieval. *Foundations and Trends in Information Retrieval*, 10(1):1–117, 2016.

Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(51):1563–1600, 2010.

Kevin Jamieson, Sumeet Katariya, Atul Deshpande, and Robert Nowak. Sparse dueling bandits. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38, pp. 416–424. PMLR, 2015.

Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, volume 31, pp. 4863–4873. Curran Associates, Inc., 2018.

Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu. Learning adversarial Markov decision processes with bandit feedback and unknown transition. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 4860–4869. PMLR, 2020a.

Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Proceedings of Thirty Third Conference on Learning Theory*, volume 125, pp. 2137–2143. PMLR, 2020b.

Tiancheng Jin and Haipeng Luo. Simultaneously learning stochastic and adversarial episodic MDPs with known transition. In *Advances in Neural Information Processing Systems*, volume 33, pp. 16557–16566. Curran Associates, Inc., 2020.

Tiancheng Jin, Longbo Huang, and Haipeng Luo. The best of both worlds: stochastic and adversarial episodic MDPs with unknown transition. In *Advances in Neural Information Processing Systems*, volume 34, pp. 20491–20502. Curran Associates, Inc., 2021.

Thorsten Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 133–142. Association for Computing Machinery, 2002.

Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning*, pp. 267–274, 2002.

Junpei Komiyama, Junya Honda, Hisashi Kashima, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in dueling bandit problem. In *Proceedings of The 28th Conference on Learning Theory*, volume 40, pp. 1141–1154. PMLR, 2015.

Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.

Mingyi Li, Taira Tsuchiya, and Kenji Yamanishi. Data- and variance-dependent regret bounds for online tabular MDPs. In *Proceedings of the 43rd International Conference on Machine Learning*. PMLR, 2026a.

Mingyi Li, Taira Tsuchiya, and Kenji Yamanishi. Policy optimization achieves data-dependent regret bounds in MDPs with unknown transitions. *arXiv preprint arXiv:2606.31769*, 2026b.

Haipeng Luo, Chen-Yu Wei, and Chung-Wei Lee. Policy optimization in adversarial MDPs: Improved exploration via dilated bonuses. In *Advances in Neural Information Processing Systems*, volume 34, pp. 22931–22942. Curran Associates, Inc., 2021.

Vivek Myers, Erdem Bıyık, and Dorsa Sadigh. Active reward learning from online preferences. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 7511–7518, 2023.

Gergely Neu and Julia Olkhovskaya. Online learning in MDPs with linear function approximation and bandit feedback. In *Advances in Neural Information Processing Systems*, volume 34, pp. 10407–10417. Curran Associates, Inc., 2021.

Gergely Neu, András György, and Csaba Szepesvári. The online loop-free stochastic shortest-path problem. In *Proceedings of the 23rd Conference on Learning Theory*, pp. 231–243, 2010.

Gergely Neu, András György, Csaba Szepesvári, and András Antos. Online Markov decision processes under bandit feedback. *IEEE Transactions on Automatic Control*, 59(3):676–691, 2014.

Ellen Novoseller, Yibing Wei, Yanan Sui, Yisong Yue, and Joel Burdick. Dueling posterior sampling for preference-based reinforcement learning. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence*, volume 124, pp. 1029–1038. PMLR, 2020.

Martin L Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.

Filip Radlinski, Madhu Kurup, and Thorsten Joachims. How does clickthrough data reflect retrieval quality? In *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, pp. 43–52. Association for Computing Machinery, 2008.

Siddartha Y. Ramamohan, Arun Rajkumar, and Shivani Agarwal. Dueling bandits: Beyond Condorcet winners to general tournament solutions. In *Advances in Neural Information Processing Systems*, volume 29, pp. 1253–1261. Curran Associates, Inc., 2016.

Aadirupa Saha, Tomer Koren, and Yishay Mansour. Adversarial dueling bandits. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 9235–9244, 2021.

Aadirupa Saha, Aldo Pacchiano, and Jonathan Lee. Dueling RL: Reinforcement learning with trajectory preferences. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 6263–6289. PMLR, 25–27 Apr 2023. URL https://proceedings.mlr.press/v206/saha23a.html.

Ayush Sekhari, Karthik Sridharan, Wen Sun, and Runzhe Wu. Contextual bandits and imitation learning with preference-based active queries. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Lior Shani, Yonathan Efroni, Aviv Rosenberg, and Shie Mannor. Optimistic policy optimization with bandit feedback. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 8604–8613. PMLR, 2020.

Joe Suk and Arpit Agarwal. Non-stationary dueling bandits under a weighted Borda criterion. *Transactions on Machine Learning Research*, 2024.

Taira Tsuchiya and Shinji Ito. A simple and adaptive learning rate for FTRL in online learning with minimax regret of $\Theta(T^{2/3})$ and its application to best-of-both-worlds. In *Advances in Neural Information Processing Systems*, volume 37, pp. 8477–8514. Curran Associates, Inc., 2024.

Tanguy Urvoy, Fabrice Clerot, Raphael Féraud, and Sami Naamane. Generic exploration and K-armed voting bandits. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, pp. 91–99. PMLR, 2013.

Yuanhao Wang, Qinghua Liu, and Chi Jin. Is RLHF more difficult than standard RL? A theoretical perspective. In *Advances in Neural Information Processing Systems*, volume 36. Curran Associates, Inc., 2023.

Christian Wirth, Riad Akrour, Gerhard Neumann, and Johannes Fürnkranz. A survey of preference-based reinforcement learning methods. *Journal of Machine Learning Research*, 18(136):1–46, 2017.

Runzhe Wu and Wen Sun. Making RL with preference-based feedback efficient via randomization. In *The Twelfth International Conference on Learning Representations*, 2024.

Yongtao Wu, Luca Viano, Kimon Antonakopoulos, Yihang Chen, Zhenyu Zhu, Quanquan Gu, and Volkan Cevher. Multi-step alignment as Markov games: An optimistic online mirror descent approach with convergence guarantees. *Transactions on Machine Learning Research*, 2026.

Yue Wu, Tao Jin, Qiwei Di, Hao Lou, Farzad Farnoud, and Quanquan Gu. Borda regret minimization for generalized linear dueling bandits. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pp. 53571–53596. PMLR, 2024.

Yichong Xu, Ruosong Wang, Lin Yang, Aarti Singh, and Artur Dubrawski. Preference-based reinforcement learning with finite-time guarantees. In *Advances in Neural Information Processing Systems*, volume 33, pp. 18784–18794. Curran Associates, Inc., 2020.

Qisen Yang, Shenzhi Wang, Matthieu Gaetan Lin, Shiji Song, and Gao Huang. Boosting offline reinforcement learning with action preference query. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 39509–39523. PMLR, 2023.

Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th International Conference on Machine Learning*, pp. 1201–1208, 2009.

Wenhao Zhan, Masatoshi Uehara, Nathan Kallus, Jason D. Lee, and Wen Sun. Provable offline preference-based reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024a.

Wenhao Zhan, Masatoshi Uehara, Wen Sun, and Jason D. Lee. Provable reward-agnostic preference-based reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024b.

Banghua Zhu, Michael Jordan, and Jiantao Jiao. Principled reinforcement learning with human feedback from pairwise or K-wise comparisons. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 43037–43067. PMLR, 2023.

Alexander Zimin and Gergely Neu. Online learning in episodic Markovian decision processes by relative entropy policy search. In *Advances in Neural Information Processing Systems*, volume 26, pp. 1583–1591. Curran Associates, Inc., 2013.

Masrour Zoghi, Shimon Whiteson, Remi Munos, and Maarten de Rijke. Relative upper confidence bound for the K-armed dueling bandit problem. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pp. 10–18. PMLR, 2014.

Table 2: Comparison of state-wise PbMDPs with the trajectory-comparison formulation.

	Trajectory-comparison formulation	State-wise PbMDPs (this paper)
Feedback	Preference over two trajectories	Preference over two arms per state
Decision unit	One arm per state (per policy)	One ordered pair of arms per state
Transition kernel	Depends on the chosen arm	May depend on the compared pair

A Additional related work

Several PbRL models assume that the preferences are generated from an underlying reward function or from the Q-values of an expert policy (Zhu et al., 2023; Zhan et al., 2024a;b; Wang et al., 2023; Sekhari et al., 2023). General preferences in sequential decision problems have also been studied through Nash equilibria of the preference game between policies (Wu et al., 2026; Colaço Carr et al., 2026), which generalize the von Neumann winner of dueling bandits (Dudík et al., 2015). We instead study the regret for the losses induced by Borda scores, with preference probabilities that may vary arbitrarily across episodes.

Learning in state-wise PbMDPs extends adversarial dueling bandits to settings with state transitions. For a survey of dueling bandits, we refer to Bengs et al. (2021). Adversarial dueling bandits with a Borda winner were studied by Saha et al. (2021), who showed that the minimax regret is of order $K^{1/3}T^{2/3}$ up to logarithmic factors. Extensions to generalized linear preference models (Wu et al., 2024) and to nonstationary preferences (Suk & Agarwal, 2024) have also been studied. Our estimator of the Borda scores differs from that of Saha et al. (2021), who draw the two arms independently and reweight the outcome, in that the uniform comparisons are built into the policy chosen by FTRL (Section 4). Regret of order $T^{2/3}$ also arises in other problems with limited feedback, for instance in hard partial monitoring games (Bartók et al., 2011) and under weakly observable feedback graphs (Alon et al., 2015).

Online learning in MDPs with adversarial losses has been studied extensively since the work of Even-Dar et al. (2009) and Neu et al. (2010; 2014). For episodic tabular MDPs with bandit feedback, the two main approaches are online linear optimization over the set of occupancy measures (Zimin & Neu, 2013; Jin et al., 2020a; Li et al., 2026a) and policy optimization (Shani et al., 2020; Luo et al., 2021; Dann et al., 2023; Li et al., 2026a;b), and best-of-both-worlds guarantees are also available (Jin & Luo, 2020; Jin et al., 2021; Dann et al., 2023; Li et al., 2026a;b). Our two algorithms follow these two approaches, respectively. Existing regret lower bounds for tabular MDPs (Jaksch et al., 2010; Jin et al., 2018; Domingues et al., 2021) concern loss feedback, whereas our lower bound concerns preference feedback and uses transitions that depend on the compared pair (Section 6).

B Deferred proofs of the global optimization approach

Here we provide the proofs of the global optimization approach. Lemmas 4–7 below are used in both the known- and unknown-transition proofs. We use $\mathbb{E}_t$ as in Section 3, with the conditioning including the learner's past random choices, and write $\mathbb{P}_t(E) = \mathbb{E}_t[\mathbb{1}[E]]$ for conditional probability. Sums over $s \in \mathcal{S}$ exclude the terminal state, and conditional expectations given a visit are used only when the conditioning event has positive conditional probability.

B.1 Borda-score estimates

The first lemma concerns the estimate in (5) when the pair compared at s is drawn uniformly from $\mathcal{A}$, that is, when $Z_t(s) = 1$.

Lemma 4. *For the estimate in (5),*

$$\mathbb{E}_t\left[\widehat{b}_t(s,i) \;\middle|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1\right] = b_t(s,i),$$
$$\mathbb{E}_t\left[\widehat{b}_t(s,i)^2 \;\middle|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1\right] \le K.$$

Proof. When $Z_t(s) = 1$, the pair compared at s is drawn uniformly from $\mathcal{A}$, so that each ordered pair (u, v) has conditional probability $1/K^2$. Moreover, by the feedback protocol in Section 3,

$$\mathbb{E}_t\big[o_t(s,u,v) \;\big|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1, a^{\mathrm{L}}_{t,s} = u, a^{\mathrm{R}}_{t,s} = v\big] = \mathsf{P}_t(s,(u,v)).$$

Consequently,

$$\begin{aligned}\mathbb{E}_t\Big[\widehat{b}_t(s,i) \;\Big|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1\Big] &= K \sum_{u,v\in[K]} \frac{\mathbb{1}[i=u]}{K^2}\big(\mathsf{P}_t(s,(u,v)) - 1\big) \\ &= \frac{1}{K}\sum_{v\in[K]}\big(\mathsf{P}_t(s,(i,v)) - 1\big) = b_t(s,i).\end{aligned}$$

Since $(o-1)^2 = 1 - o$ for $o \in \{0,1\}$, we similarly obtain

$$\begin{aligned}\mathbb{E}_t\Big[\widehat{b}_t(s,i)^2 \;\Big|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1\Big] &= K^2 \sum_{u,v\in[K]} \frac{\mathbb{1}[i=u]}{K^2}\big(1 - \mathsf{P}_t(s,(u,v))\big) \\ &= -K b_t(s,i) \le K,\end{aligned}$$

which completes the proof. □

For known transitions, we omit states unreachable under every policy. The remaining states have positive visit probability under every policy satisfying $\pi(a \mid s) \ge \rho/A$, since every path with positive transition probabilities also has positive probability under such a policy. Thus the denominators in (4) are positive on these states.

Lemma 5. *The estimates in (4) satisfy*

$$\mathbb{E}_t\Big[\widetilde{b}_t(s,i)\Big] = b_t(s,i), \qquad \mathbb{E}_t\Big[\widehat{\ell}_t(s,a)\Big] = \ell_t(s,a), \qquad \mathbb{E}_t\Big[\widetilde{b}_t(s,i)^2\Big] \le \frac{K}{\rho\widetilde{q}_t(s)}.$$

Moreover,

$$\mathbb{E}_t\left[\sum_{s,a} \widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2\right] \le \frac{SK}{\rho}. \tag{12}$$

Proof. Since the sampling rule in Line 5 of Algorithm 1 selects each pair with probability $\pi_t(a \mid s)$, the occupancy measure of the executed policy is $\widetilde{q}_t$. Since $Z_t(s) \sim \mathsf{Ber}(\rho)$ is drawn when s is visited, independently of the trajectory up to s, we have

$$\mathbb{P}_t(\mathbb{I}_t(s) = 1, Z_t(s) = 1) = \rho\widetilde{q}_t(s).$$

Hence, by the definition of $\widetilde{b}_t$ in (4) and Lemma 4, we have

$$\begin{aligned}\mathbb{E}_t\Big[\widetilde{b}_t(s,i)\Big] &= \frac{\rho\widetilde{q}_t(s)}{\rho\widetilde{q}_t(s)}\mathbb{E}_t\Big[\widehat{b}_t(s,i) \;\Big|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1\Big] = b_t(s,i), \\ \mathbb{E}_t\Big[\widetilde{b}_t(s,i)^2\Big] &= \frac{\rho\widetilde{q}_t(s)}{\rho^2\widetilde{q}_t(s)^2}\mathbb{E}_t\Big[\widehat{b}_t(s,i)^2 \;\Big|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1\Big] \le \frac{K}{\rho\widetilde{q}_t(s)}.\end{aligned}$$

By the definition of $\widehat{\ell}_t$ in (4) and the first identity,

$$\mathbb{E}_t\Big[\widehat{\ell}_t(s,a)\Big] = -\frac{\mathbb{E}_t[\widetilde{b}_t(s,a^{\mathrm{L}})] + \mathbb{E}_t[\widetilde{b}_t(s,a^{\mathrm{R}})]}{2} = -\frac{b_t(s,a^{\mathrm{L}}) + b_t(s,a^{\mathrm{R}})}{2} = \ell_t(s,a).$$

For each state s, we have

$$\begin{aligned}\mathbb{E}_t\left[\sum_{i,j} \widetilde{q}_t(s,(i,j))\widehat{\ell}_t(s,(i,j))^2\right] &\le \frac{1}{2}\sum_{i,j}\widetilde{q}_t(s,(i,j))\Big(\mathbb{E}_t\big[\widetilde{b}_t(s,i)^2\big] + \mathbb{E}_t\big[\widetilde{b}_t(s,j)^2\big]\Big) \\ &\le \frac{K}{\rho\widetilde{q}_t(s)}\sum_{i,j}\widetilde{q}_t(s,(i,j)) = \frac{K}{\rho},\end{aligned}$$

where the first inequality follows from the definition of $\widehat{\ell}_t$ and $(x+y)^2 \le 2x^2 + 2y^2$, and the second inequality from the bound on $\mathbb{E}_t[\widetilde{b}_t(s,i)^2]$ shown above. Summing over the states completes the proof. □

B.2 FTRL WITH DECREASING FEASIBLE SETS

In the unknown-transition algorithm of Appendix B.5, the feasible set is determined by transition confidence sets that can only shrink over time. We state the FTRL bound for decreasing feasible sets, which includes the case of a fixed feasible set.

Lemma 6. *Fix $\eta > 0$, and let $\mathcal{Q}_1 \supseteq \cdots \supseteq \mathcal{Q}_n$ be nonempty compact convex subsets of $\mathbb{R}_{\geq 0}^{\mathcal{S}\times\mathcal{A}}$ satisfying $\sum_{s\in\mathcal{S}_h, a\in\mathcal{A}} q(s,a) = 1$ for every $q \in \mathcal{Q}_1$ and every $h \in \{0,\dots,H-1\}$. For nonnegative loss vectors $\widehat{\ell}_1,\dots,\widehat{\ell}_n$, let*

$$\widetilde{q}_t \in \operatorname*{arg\,min}_{q\in\mathcal{Q}_t}\left\{\left\langle q, \sum_{\tau=1}^{t-1}\widehat{\ell}_\tau\right\rangle - \frac{1}{\eta}\,\mathrm{H}(q)\right\}.$$

Then, for every $q \in \mathcal{Q}_n$,

$$\sum_{t=1}^{n}\langle \widetilde{q}_t - q, \widehat{\ell}_t\rangle \leq \frac{H\log(SA)}{\eta} + \frac{\eta}{2}\sum_{t=1}^{n}\sum_{s,a}\widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2.$$

Proof. Set $\mathcal{Q}_{n+1} = \mathcal{Q}_n$ and define the minimum objective value

$$\Phi_t = \min_{u\in\mathcal{Q}_t}\left\{\left\langle u, \sum_{\tau=1}^{t-1}\widehat{\ell}_\tau\right\rangle - \frac{1}{\eta}\,\mathrm{H}(u)\right\}.$$

We use the unnormalized relative entropy

$$D(u,v) = \sum_{s,a}\left(u(s,a)\log\frac{u(s,a)}{v(s,a)} - u(s,a) + v(s,a)\right) = -\,\mathrm{H}(u) + \mathrm{H}(v) + \langle\nabla\,\mathrm{H}(v), u - v\rangle.$$

Coordinates that vanish on all of $\mathcal{Q}_t$ can be dropped. On every other coordinate, $\widetilde{q}_t$ is positive, since the derivative of $y\log y$ tends to $-\infty$ as $y \to 0$, and hence the first-order optimality condition gives

$$\left\langle \sum_{\tau<t}\widehat{\ell}_\tau - \frac{1}{\eta}\nabla\,\mathrm{H}(\widetilde{q}_t), u - \widetilde{q}_t\right\rangle \geq 0$$

for every $u \in \mathcal{Q}_t$. Substituting the identity for D gives

$$\left\langle u, \sum_{\tau<t}\widehat{\ell}_\tau\right\rangle - \frac{1}{\eta}\,\mathrm{H}(u) \geq \Phi_t + \frac{1}{\eta}D(u,\widetilde{q}_t)$$

for every $u \in \mathcal{Q}_t$. Since $\mathcal{Q}_{t+1} \subseteq \mathcal{Q}_t$, this lower bound holds for every candidate in the definition of Φ_{t+1}. Dropping all constraints except $u \geq 0$ then yields

$$\begin{aligned}
\Phi_{t+1} - \Phi_t &\geq \min_{u\in\mathcal{Q}_{t+1}}\left\{\langle u, \widehat{\ell}_t\rangle + \frac{1}{\eta}D(u,\widetilde{q}_t)\right\} \\
&\geq \sum_{s,a}\min_{y\geq 0}\left\{y\widehat{\ell}_t(s,a) + \frac{1}{\eta}\left(y\log\frac{y}{\widetilde{q}_t(s,a)} - y + \widetilde{q}_t(s,a)\right)\right\} \\
&= \frac{1}{\eta}\sum_{s,a}\widetilde{q}_t(s,a)\left(1 - \mathrm{e}^{-\eta\widehat{\ell}_t(s,a)}\right) \\
&\geq \langle\widetilde{q}_t, \widehat{\ell}_t\rangle - \frac{\eta}{2}\sum_{s,a}\widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2, \qquad (13)
\end{aligned}$$

where the equality follows from the fact that the minimum over $y \geq 0$ is attained at $y = \widetilde{q}_t(s,a)\mathrm{e}^{-\eta\widehat{\ell}_t(s,a)}$, and the last inequality from $\mathrm{e}^{-z} \leq 1 - z + z^2/2$ for $z \geq 0$.

For $q \in \mathcal{Q}_n = \mathcal{Q}_{n+1}$, we have $\Phi_{n+1} \leq \langle q, \sum_{t=1}^{n}\widehat{\ell}_t\rangle - \mathrm{H}(q)/\eta$ and $\Phi_1 = -\max_{u\in\mathcal{Q}_1}\mathrm{H}(u)/\eta$. Summing (13) gives

$$\begin{aligned}
\sum_{t=1}^{n}\langle\widetilde{q}_t - q, \widehat{\ell}_t\rangle &\leq \Phi_{n+1} - \Phi_1 - \left\langle q, \sum_{t=1}^{n}\widehat{\ell}_t\right\rangle + \frac{\eta}{2}\sum_{t=1}^{n}\sum_{s,a}\widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2 \\
&\leq \frac{\max_{u\in\mathcal{Q}_1}\mathrm{H}(u) - \mathrm{H}(q)}{\eta} + \frac{\eta}{2}\sum_{t=1}^{n}\sum_{s,a}\widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2.
\end{aligned}$$

By the assumption on $\mathcal{Q}_1$, the restriction of q to each layer $\mathcal{S}_h \times \mathcal{A}$ is a probability distribution, whose entropy lies in $[0, \log(S_h A)]$. Summing over the layers gives $0 \le \mathrm{H}(q) \le \sum_{h=0}^{H-1} \log(S_h A) \le H \log(SA)$, which completes the proof. □

B.3 Cost of the uniform component

The occupancy measure of the best policy need not belong to Ω_ρ. The following lemma bounds the cost of adding the required uniform component.

Lemma 7. *For a Markov policy π and $\rho \in [0,1]$, define*

$$\pi_\rho(a \mid s) = (1-\rho)\pi(a \mid s) + \rho/A. \tag{14}$$

For every $\ell \colon \mathcal{S} \times \mathcal{A} \to [0,1]$,

$$V^{\pi_\rho}(s_0;\ell) - V^{\pi}(s_0;\ell) \le \rho H^2.$$

Proof. By the performance difference lemma (Kakade & Langford, 2002), we have

$$\begin{aligned} V^{\pi_\rho}(s_0;\ell) - V^{\pi}(s_0;\ell) &= \sum_s q^{\pi_\rho}(s) \langle \pi_\rho(\cdot \mid s) - \pi(\cdot \mid s), Q^{\pi}(s,\cdot;\ell)\rangle \\ &= \rho \sum_s q^{\pi_\rho}(s) \left(\frac{1}{A} \sum_a Q^{\pi}(s,a;\ell) - \sum_a \pi(a \mid s) Q^{\pi}(s,a;\ell) \right) \\ &\le \rho H \sum_s q^{\pi_\rho}(s) = \rho H^2, \end{aligned}$$

where the inequality follows from $Q^{\pi}(s,a;\ell) \in [0,H]$ for every (s,a) and the last equality from $\sum_s q^{\pi_\rho}(s) = H$. This completes the proof. □

B.4 Known transition kernel (Proof of Theorem 1)

Proof of Theorem 1. Let π^* be the best fixed Markov policy in hindsight and set $q^*_\rho = q^{\pi^*_\rho}$, with π^*_ρ defined by (14) in Lemma 7. Then $q^*_\rho \in \Omega_\rho$. By Lemma 7,

$$\mathbb{E}[\mathsf{Reg}_T] \le \rho H^2 T + \mathbb{E}\left[\sum_{t=1}^T \langle \widetilde{q}_t - q^*_\rho, \ell_t \rangle \right].$$

Since $\widetilde{q}_t$ and q^*_ρ are determined before episode t, Lemma 5 gives

$$\mathbb{E}\left[\sum_{t=1}^T \langle \widetilde{q}_t - q^*_\rho, \ell_t \rangle \right] = \mathbb{E}\left[\sum_{t=1}^T \langle \widetilde{q}_t - q^*_\rho, \mathbb{E}_t[\widehat{\ell}_t] \rangle \right] = \mathbb{E}\left[\sum_{t=1}^T \langle \widetilde{q}_t - q^*_\rho, \widehat{\ell}_t \rangle \right].$$

Since $\widehat{\ell}_t \ge 0$ by $o_t \le 1$, this is further upper bounded as

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^T \langle \widetilde{q}_t - q^*_\rho, \widehat{\ell}_t \rangle \right] &\le \frac{H \log(SA)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \mathbb{E}\left[\sum_{s,a} \widetilde{q}_t(s,a) \widehat{\ell}_t(s,a)^2 \right] \\ &\le \frac{H \log(SA)}{\eta} + \frac{\eta S K T}{2\rho}, \end{aligned}$$

where the first inequality follows from Lemma 6 with $\mathcal{Q}_t = \Omega_\rho$ for every t and the second inequality from (12) in Lemma 5. Combining these inequalities yields

$$\mathbb{E}[\mathsf{Reg}_T] \le \frac{H \log(SA)}{\eta} + \frac{\eta S K T}{2\rho} + \rho H^2 T. \tag{15}$$

Choose

$$\eta = \left(\frac{\log(SA)}{T\sqrt{SK}} \right)^{2/3}, \qquad \rho = \min\left\{ \frac{1}{2}, \frac{\sqrt{\eta SK}}{H} \right\}.$$

If $\sqrt{\eta SK}/H \le 1/2$, then

$$\frac{\eta SKT}{2\rho} + \rho H^2 T = \frac{3}{2} HT\sqrt{\eta SK},$$
$$\frac{H\log(SA)}{\eta} = HT\sqrt{\eta SK} = \big(H^3 SK\log(SA)\big)^{1/3} T^{2/3},$$

and hence the right-hand side of (15) is at most $\frac{5}{2}(H^3 SK\log(SA))^{1/3}T^{2/3}$. Otherwise,

$$\frac{(H^3 SK\log(SA))^{1/3}T^{2/3}}{HT} = \left(\frac{SK\log(SA)}{T}\right)^{1/3} = \sqrt{\eta SK} > \frac{H}{2} \ge \frac{1}{2},$$

and hence $HT \le 2(H^3 SK\log(SA))^{1/3}T^{2/3}$. Combining the two cases with the trivial bound $\mathsf{Reg}_T \le HT$, we obtain $\mathbb{E}[\mathsf{Reg}_T] \le \frac{5}{2}\min\{HT, (H^3 SK\log(SA))^{1/3}T^{2/3}\}$ for every $T \ge 1$. Using $A = K^2$ and $\log(SA) \le 2\log(SK)$ completes the proof. □

B.5 Unknown transition kernel

We now turn to the case where the transition kernel P is fixed but unknown. The algorithm keeps the FTRL update and the uniform-comparison rule of Algorithm 1, but the feasible set of FTRL is now determined by a confidence set for P, and the estimator divides by an upper confidence bound on the state-visit probability instead of by $\widetilde{q}_t(s)$.

Theorem 8. *Suppose that the transition kernel is fixed but unknown. With the parameter choices specified in the proof, Algorithm 3 satisfies, for every $T \ge 1$,*

$$\mathbb{E}[\mathsf{Reg}_T] = \widetilde{O}\Big(\min\Big\{HT, (H^3 SK)^{1/3}T^{2/3} + HSK\sqrt{T} + HS^4K^2\Big\}\Big). \tag{16}$$

The last two terms in (16) are the cost of estimating the transitions. The term $HSK\sqrt{T}$ comes from the width of the confidence sets, and HS^4K^2 comes from the lower-order term of the bound of Dann et al. (2023) on the difference between the occupancy measures under the true and the estimated transition kernels (Lemma 10).

B.5.1 Transition confidence sets

We construct the confidence sets as in Jin et al. (2020a). The episodes are divided into epochs, and the empirical transition kernel and its entrywise Bernstein confidence intervals are updated only at the beginning of each epoch. Let $N_k(s,a)$ and $M_k(s,a,s')$ be the visit counts before epoch k, and let $N(s,a)$ and $M(s,a,s')$ denote the current counts. An epoch ends when $N(s,a) \ge \max\{1, 2N_k(s,a)\}$ for some visited pair. For $k \ge 2$, define

$$\bar{P}_k(s' \mid s,a) = \frac{M_k(s,a,s')}{\max\{1, N_k(s,a)\}}, \tag{17}$$

$$\beta_k(s,a,s') = 2\sqrt{\frac{\bar{P}_k(s' \mid s,a)\iota}{\max\{1, N_k(s,a)-1\}}} + \frac{14\iota}{3\max\{1, N_k(s,a)-1\}}, \qquad \iota = \log\frac{TSA}{\delta}. \tag{18}$$

Let $\mathcal{P}_1$ be the set of all layered transition kernels and set

$$\mathcal{C}_k = \big\{P' \in \mathcal{P}_1 : |P'(s' \mid s,a) - \bar{P}_k(s' \mid s,a)| \le \beta_k(s,a,s') \text{ for all } (s,a,s') \text{ with } s' \in \mathcal{S}_{h(s)+1}\big\}. \tag{19}$$

The remaining entries of a layered kernel are zero. With $\mathcal{C}_1 = \mathcal{P}_1$, update

$$\mathcal{P}_k = \begin{cases} \mathcal{P}_{k-1} \cap \mathcal{C}_k, & \text{if } \mathcal{P}_{k-1} \cap \mathcal{C}_k \ne \varnothing, \\ \mathcal{P}_{k-1}, & \text{otherwise.} \end{cases} \tag{20}$$

Under this rule, $\mathcal{P}_k$ is nonempty and $\mathcal{P}_k \subseteq \mathcal{P}_{k-1}$ for every history. If $P \in \mathcal{C}_k$ for every epoch, the first case always applies and $P \in \mathcal{P}_k \subseteq \mathcal{C}_k$.

The constraints defining $\mathcal{P}_k$ are intervals for the individual entries $P'(s' \mid s,a)$. For every (s,a,s') with $s' \in \mathcal{S}_{h(s)+1}$, define

$$P_k^-(s' \mid s,a) = \min_{P' \in \mathcal{P}_k} P'(s' \mid s,a), \qquad P_k^+(s' \mid s,a) = \max_{P' \in \mathcal{P}_k} P'(s' \mid s,a).$$

Thus

$$\mathcal{P}_k = \{P' \in \mathcal{P}_1 : P_k^- \le P' \le P_k^+\},$$

where both inequalities hold for every (s, a, s') with $s' \in \mathcal{S}_{h(s)+1}$.

B.5.2 Extended occupancy measures

Since P is unknown, the algorithm replaces Ω_ρ by the set of occupancy measures that are consistent with some kernel in $\mathcal{P}_k$. To describe this set by linear constraints, we use occupancy measures over triples (s, a, s'), as in Jin et al. (2020a). For a kernel P' and a policy π, write $q^{P',\pi}$ for their occupancy measure. For (s, a, s') with $s' \in \mathcal{S}_{h(s)+1}$, define the extended occupancy measure

$$x(s, a, s') = q^{P',\pi}(s, a)P'(s' \mid s, a), \qquad q_x(s, a) = \sum_{s'} x(s, a, s'), \quad q_x(s) = \sum_a q_x(s, a).$$

Let $\mathcal{X}_\rho(\mathcal{P}_k)$ be the set of all $x(s, a, s') \ge 0$ satisfying

$$\sum_{a\in\mathcal{A},s'\in\mathcal{S}_1} x(s_0, a, s') = 1,$$

$$\sum_{s^-\in\mathcal{S}_{h-1},a} x(s^-, a, s) = q_x(s) \qquad \text{for every } s \in \mathcal{S}_h \text{ with } 1 \le h < H,$$

$$P_k^-(s' \mid s, a)q_x(s, a) \le x(s, a, s') \le P_k^+(s' \mid s, a)q_x(s, a),$$

$$q_x(s, a) \ge \rho q_x(s)/A.$$

The last two constraints hold for all (s, a, s') with $s' \in \mathcal{S}_{h(s)+1}$ and for all nonterminal (s, a), respectively. All the constraints above are linear. Since the extended occupancy measure of a uniform policy under any $P' \in \mathcal{P}_k$ is feasible, $\mathcal{X}_\rho(\mathcal{P}_k)$ is a nonempty convex polytope. Its projection

$$\Omega_\rho(\mathcal{P}_k) = \{q_x : x \in \mathcal{X}_\rho(\mathcal{P}_k)\}$$

is also convex and satisfies $\Omega_\rho(\mathcal{P}_{k+1}) \subseteq \Omega_\rho(\mathcal{P}_k)$.

We now show that $\Omega_\rho(\mathcal{P}_k)$ consists exactly of the occupancy measures $q^{P',\pi}$ with $P' \in \mathcal{P}_k$ and $\pi(a \mid s) \ge \rho/A$ for all (s, a). For $x \in \mathcal{X}_\rho(\mathcal{P}_k)$, define

$$P_x(s' \mid s, a) = \frac{x(s, a, s')}{q_x(s, a)}, \qquad \pi_x(a \mid s) = \frac{q_x(s, a)}{q_x(s)}$$

whenever the denominators are positive. If $q_x(s, a) = 0$, let $P_x(\cdot \mid s, a)$ be $P'(\cdot \mid s, a)$ for an arbitrary fixed $P' \in \mathcal{P}_k$, and if $q_x(s) = 0$, let $\pi_x(\cdot \mid s)$ be uniform. Since the constraints defining $\mathcal{P}_k$ are imposed separately on each $P'(\cdot \mid s, a)$, we have $P_x \in \mathcal{P}_k$, and the last constraint in the definition of $\mathcal{X}_\rho(\mathcal{P}_k)$ gives $\pi_x(a \mid s) \ge \rho/A$. By induction over the layers starting from s_0, the second constraint gives $q_x = q^{P_x,\pi_x}$. Conversely, for any $P' \in \mathcal{P}_k$ and any policy π with $\pi(a \mid s) \ge \rho/A$, the extended occupancy measure $x(s, a, s') = q^{P',\pi}(s, a)P'(s' \mid s, a)$ satisfies all constraints. In particular, $\Omega_\rho \subseteq \Omega_\rho(\mathcal{P}_k)$ whenever $P \in \mathcal{P}_k$.

B.5.3 Algorithm under unknown transitions

Let k_t be the epoch containing episode t. The algorithm replaces the update (2) of Algorithm 1 by

$$\widetilde{x}_t \in \operatorname*{arg\,min}_{x\in\mathcal{X}_\rho(\mathcal{P}_{k_t})} \left\{ \left\langle q_x, \sum_{\tau<t} \widehat{\ell}_\tau \right\rangle - \frac{1}{\eta} \mathrm{H}(q_x) \right\}, \qquad \widetilde{q}_t = q_{\widetilde{x}_t}. \tag{21}$$

Since the objective depends only on q_x, $\widetilde{q}_t$ is the FTRL minimizer over $\Omega_\rho(\mathcal{P}_{k_t})$. The algorithm sets $\pi_t = \pi_{\widetilde{x}_t}$, defines π_t^{r} by (3), and computes

$$u_t(s) = \max_{P'\in\mathcal{P}_{k_t}} q^{P',\pi_t}(s). \tag{22}$$

The maximization is over the kernel only, and u_t can be computed by dynamic programming over the intervals defining $\mathcal{P}_{k_t}$ (Jin et al., 2020a, Section 3.3). The algorithm executes π_t with the uniform-comparison rule of Algorithm 1 and replaces the estimator in (4) by

$$\widetilde{b}_t(s, i) = \frac{\mathbb{I}_t(s)Z_t(s)\widehat{b}_t(s, i)}{\rho u_t(s)}, \qquad \widehat{\ell}_t(s, a) = -\frac{\widetilde{b}_t(s, a^{\mathrm{L}}) + \widetilde{b}_t(s, a^{\mathrm{R}})}{2}, \tag{23}$$

Algorithm 3: Global optimization under unknown transitions

Input: Learning rate $\eta > 0$, exploration rate $\rho \in (0, 1/2]$, and $\delta \in (0, 1)$

1 Set $k = 1$, let $\mathcal{P}_1 = \mathcal{C}_1$ be the set of all layered transition kernels, and initialize all counts N, M, N_1, M_1 to zero;
2 **for** $t = 1, \dots, T$ **do**
3 Set $k_t = k$ and compute $\widetilde{x}_t, \widetilde{q}_t$ by (21);
4 Set $\pi_t = \pi_{\widetilde{x}_t}$, and compute π_t^{r} and u_t by (3) and (22);
5 Execute π_t as in Lines 4–6 of Algorithm 1;
6 Set $Z_t(s) = 0$ at unvisited states and compute the estimates by (5) and (23);
7 Increment $N(s, a)$ and $M(s, a, s')$ for every observed transition (s, a, s');
8 **if** $t < T$ *and* $N(s, a) \geq \max\{1, 2N_k(s, a)\}$ *for some visited* (s, a) **then**
9 Set $k \leftarrow k + 1$, $N_k \leftarrow N$, and $M_k \leftarrow M$;
10 Compute $\bar{P}_k, \beta_k, \mathcal{C}_k$ by (17)–(19), and update $\mathcal{P}_k$ by (20);

with the first estimate defined as zero when $u_t(s) = 0$.

Three properties of Algorithm 3 are used in the analysis. The quantities $\widetilde{q}_t$, π_t, and u_t are determined by the observations before episode t (Lines 3–4), the transition counts are generated by the executed policies π_t (Line 7), and the feasible set changes between epochs without restarting FTRL (Lines 9–10). If $\mathcal{P}_k = \{P\}$ for every k, then $\widetilde{x}_t(s, a, s') = \widetilde{q}_t(s, a)P(s' \mid s, a)$ and $u_t(s) = \widetilde{q}_t(s)$, and the algorithm reduces to Algorithm 1.

B.5.4 Borda-score estimates under unknown transitions

In what follows, let $q_t = q^{P,\pi_t}$ be the occupancy measure of π_t under the true transition kernel. In general, q_t differs from the selected $\widetilde{q}_t$.

Lemma 9. *Suppose that $P \in \mathcal{P}_{k_t}$. Then, for every state s with $u_t(s) > 0$,*

$$\begin{aligned} \mathbb{E}_t\left[\widetilde{b}_t(s, i)\right] &= \frac{q_t(s)}{u_t(s)} b_t(s, i), \\ \mathbb{E}_t\left[\widehat{\ell}_t(s, a)\right] &= \frac{q_t(s)}{u_t(s)} \ell_t(s, a) \leq \ell_t(s, a), \\ \mathbb{E}_t\left[\widetilde{b}_t(s, i)^2\right] &\leq \frac{K}{\rho u_t(s)}. \end{aligned} \tag{24}$$

For every state s with $u_t(s) = 0$, both $q_t(s)$ and $\widetilde{q}_t(s)$ are zero and the estimates at s are zero. Moreover, still under $P \in \mathcal{P}_{k_t}$,

$$\mathbb{E}_t\left[\sum_{s,a} \widetilde{q}_t(s, a)\widehat{\ell}_t(s, a)^2\right] \leq \frac{SK}{\rho}. \tag{25}$$

Proof. Since both P and $P_{\widetilde{x}_t}$ belong to $\mathcal{P}_{k_t}$, the definition of u_t in (22) gives

$$q_t(s) \leq u_t(s), \qquad \widetilde{q}_t(s) \leq u_t(s).$$

If $q_t(s) = 0$, then s is not visited, and the claims hold since the estimates at s are zero. Suppose $q_t(s) > 0$. Then $\mathbb{P}_t(\mathbb{I}_t(s) = 1, Z_t(s) = 1) = \rho q_t(s) > 0$, and Lemma 4 gives

$$\begin{aligned} \mathbb{E}_t\left[\widetilde{b}_t(s, i)\right] &= \frac{\rho q_t(s)}{\rho u_t(s)} \mathbb{E}_t\left[\widehat{b}_t(s, i) \;\middle|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1\right] = \frac{q_t(s)}{u_t(s)} b_t(s, i), \\ \mathbb{E}_t\left[\widetilde{b}_t(s, i)^2\right] &= \frac{\rho q_t(s)}{\rho^2 u_t(s)^2} \mathbb{E}_t\left[\widehat{b}_t(s, i)^2 \;\middle|\; \mathbb{I}_t(s) = 1, Z_t(s) = 1\right] \\ &\leq \frac{K q_t(s)}{\rho u_t(s)^2} \leq \frac{K}{\rho u_t(s)}. \end{aligned}$$

By the definition of $\widehat{\ell}_t$ in (23) and the first identity,

$$\mathbb{E}_t\left[\widehat{\ell}_t(s, a)\right] = -\frac{\mathbb{E}_t[\widetilde{b}_t(s, a^{\mathrm{L}})] + \mathbb{E}_t[\widetilde{b}_t(s, a^{\mathrm{R}})]}{2} = \frac{q_t(s)}{u_t(s)} \ell_t(s, a) \leq \ell_t(s, a),$$

where the inequality follows from $q_t(s) \le u_t(s)$ and $\ell_t(s,a) \ge 0$.

As in the proof of Lemma 5, for every state s with $u_t(s) > 0$ we have

$$\mathbb{E}_t\left[\sum_{i,j} \widetilde{q}_t(s,(i,j))\widehat{\ell}_t(s,(i,j))^2\right] \le \frac{1}{2}\sum_{i,j} \widetilde{q}_t(s,(i,j))\Big(\mathbb{E}_t\big[\widetilde{b}_t(s,i)^2\big] + \mathbb{E}_t\big[\widetilde{b}_t(s,j)^2\big]\Big) \le \frac{K\widetilde{q}_t(s)}{\rho u_t(s)} \le \frac{K}{\rho},$$

where the last inequality follows from $\widetilde{q}_t(s) \le u_t(s)$. Since states with $u_t(s) = 0$ contribute zero, summing over the at most S nonterminal states completes the proof. □

B.5.5 Transition estimation error

In this section, π_t may be any Markov policy determined by the observations before episode t, and $q_t = q^{P,\pi_t}$. The counts and the confidence sets $\mathcal{C}_k$ are those defined above, computed from the trajectories of these policies. By Jin et al. (2020a, Lemma 2),

$$\mathbb{P}(P \in \mathcal{C}_k \text{ for every epoch } k) \ge 1 - 4\delta. \tag{26}$$

The following bound on the occupancy measures under kernels in the confidence sets is used for both global and policy optimization.

Lemma 10. *Suppose that $T \ge A$, and let $\delta \in (0,1)$ be the confidence parameter in (18). There are universal constants $C, C_0 > 0$ such that, with probability at least $1 - C_0 H\delta$, $P \in \mathcal{C}_k$ for every epoch k, and*

$$\sum_{t=1}^{T}\sum_{s} |q^{P_t^s,\pi_t}(s) - q_t(s)| \le C\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big) \tag{27}$$

holds simultaneously for every choice of kernels $P_t^s \in \mathcal{C}_{k_t}$, one for each pair (t,s), where $\iota = \log(TSA/\delta)$.

Lemma 10 follows from Dann et al. (2023, Lemma C.4), whose proof uses the counts and δ only through two inequalities. We state their lemma with these inequalities as conditions, writing $N^{<t}(s,a) = \sum_{u<t} \mathbb{I}_u(s,a)$ for the number of visits to (s,a) before episode t and $n_t(s,a) = \max\{1, N^{<t}(s,a)\}$.

Lemma 11 (adapted from Dann et al., 2023, Lemma C.4)**.** *Let $\iota \ge 1$ and $C > 0$. Suppose that*

$$|P'(s' \mid s,a) - P(s' \mid s,a)| \le C\min\left\{1, \sqrt{\frac{P(s' \mid s,a)\iota}{n_t(s,a)}} + \frac{\iota}{n_t(s,a)}\right\} \tag{28}$$

for every $t \in [T]$, every (s,a,s') with $s' \in \mathcal{S}_{h(s)+1}$, and every $P' \in \mathcal{C}_{k_t}$, and that

$$\sum_{t=1}^{T}\sum_{s\in\mathcal{S}_h, a\in\mathcal{A}} \frac{q_t(s,a)}{n_t(s,a)} \le CS_hA\iota \tag{29}$$

for every $h \in \{0,\dots,H-1\}$. Then, for every choice of kernels $P_t^s \in \mathcal{C}_{k_t}$, one for each pair (t,s), and every $g_t\colon \mathcal{S} \to [0,G]$,

$$\sum_{t=1}^{T}\sum_{s} |q^{P_t^s,\pi_t}(s) - q_t(s)|\, g_t(s) \le C'\left(\sqrt{HS^2A\iota^2\sum_{t=1}^{T}\sum_{s} q_t(s)g_t(s)^2} + HS^4AG\iota^2\right),$$

where $C' > 0$ depends only on C.

Proof. Dann et al. (2023) prove this bound with $\iota \log T$ in place of ι^2, since they set $\delta = T^{-3}$. Their proof uses only the identity of their Lemma C.3, their Lemma B.1, which is (28), and their Lemma B.2, which is (29) with $S_hA\log T + \log(1/\delta)$ in place of $S_hA\iota$. The same proof therefore applies under the conditions above. □

Proof of Lemma 10. We verify the conditions of Lemma 11 with $\iota = \log(TSA/\delta)$. Since an epoch ends as soon as a visited pair satisfies $N(s,a) \geq \max\{1, 2N_k(s,a)\}$ and the count of a pair changes only when the pair is visited, every pair satisfies $N^{<t}(s,a) < \max\{1, 2N_{k_t}(s,a)\}$ during epoch k_t, and hence

$$\max\{1, N_{k_t}(s,a)\} \leq n_t(s,a) \leq 2\max\{1, N_{k_t}(s,a)\} \leq 4\max\{1, N_{k_t}(s,a) - 1\}. \tag{30}$$

Suppose that $P \in \mathcal{C}_k$ for every epoch k, which holds with probability at least $1 - 4\delta$ by (26), and fix t, (s,a,s') with $s' \in \mathcal{S}_{h(s)+1}$, and $P' \in \mathcal{C}_{k_t}$. If $k_t = 1$, then $n_t(s,a) \leq 2$ by (30), and (28) holds for $C \geq 2$ since its left-hand side is at most $1 \leq 2\iota/n_t(s,a)$. If $k_t \geq 2$, write $\beta = \beta_{k_t}(s,a,s')$, $\bar{P} = \bar{P}_{k_t}(s' \mid s,a)$, $p = P(s' \mid s,a)$, and $m = \max\{1, N_{k_t}(s,a) - 1\}$. Since $P'(s' \mid s,a)$ and p both lie within β of $\bar{P}$,

$$|P'(s' \mid s,a) - p| \leq 2\beta, \qquad \bar{P} \leq p + \beta = p + 2\sqrt{\frac{\bar{P}\iota}{m}} + \frac{14\iota}{3m}.$$

Solving the second inequality for $\sqrt{\bar{P}}$ and substituting the result into the definition of β give

$$\sqrt{\bar{P}} \leq \sqrt{\frac{\iota}{m}} + \sqrt{p + \frac{17\iota}{3m}} \leq \sqrt{p} + 4\sqrt{\frac{\iota}{m}}, \qquad \beta \leq 2\sqrt{\frac{p\iota}{m}} + \frac{13\iota}{m}.$$

Since $m \geq n_t(s,a)/4$ by (30) and $|P'(s' \mid s,a) - p| \leq 1$, condition (28) holds for a universal constant C.

By Dann et al. (2023, Lemma B.2), which is Lemma D.3.7 of Jin et al. (2021), for each layer h, with probability at least $1 - \delta$,

$$\sum_{t=1}^{T} \sum_{s \in \mathcal{S}_h, a \in \mathcal{A}} \frac{q_t(s,a)}{n_t(s,a)} \leq C_2\big(S_h A \log(2T) + \log(1/\delta)\big) \leq 2C_2 S_h A\iota$$

for a universal constant $C_2 > 0$, where the second inequality follows from $\log(2T) \leq \iota$ and $\log(1/\delta) \leq \iota$. The proof of the first inequality uses only $\mathbb{E}_t[\mathbb{I}_t(s,a)] = q_t(s,a)$ and therefore applies even though the policies depend on the preference feedback. By a union bound over the H layers, condition (29) holds with probability at least $1 - H\delta$.

On the intersection of the two events, which has probability at least $1 - (4+H)\delta \geq 1 - 5H\delta$, Lemma 11 with $g_t \equiv 1$ gives, for every choice of the kernels $P_t^s \in \mathcal{C}_{k_t}$,

$$\begin{aligned}\sum_{t=1}^{T} \sum_{s} |q^{P_t^s, \pi_t}(s) - q_t(s)| &\leq C'\left(\sqrt{HS^2A\iota^2 \sum_{t=1}^{T}\sum_{s} q_t(s)} + HS^4A\iota^2\right) \\ &= C'\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big),\end{aligned}$$

where the equality follows from $\sum_s q_t(s) = H$. This completes the proof. □

For global optimization, the kernels defining $\widetilde{q}_t$ and u_t lie in $\mathcal{P}_{k_t}$, which is contained in $\mathcal{C}_{k_t}$ whenever $P \in \mathcal{C}_k$ for every epoch k. Lemma 10 thus gives the following two bounds.

Lemma 12. *Suppose that $T \geq A$. For a sufficiently large universal constant C_{tr}, define*

$$\mathcal{E} = \{P \in \mathcal{C}_k \text{ for every epoch } k\} \cap \left\{\max\left\{\sum_{t,s,a} |\widetilde{q}_t(s,a) - q_t(s,a)|, \sum_{t,s} |u_t(s) - q_t(s)|\right\} \leq C_{\mathrm{tr}}\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big)\right\}. \tag{31}$$

Then $\mathbb{P}(\mathcal{E}) \geq 1 - C_0 H\delta$ for a universal constant $C_0 > 0$. In particular, $P \in \mathcal{P}_k$ for every epoch on $\mathcal{E}$.

Proof. On the event of Lemma 10, we have $P \in \mathcal{C}_k$ for every epoch k, and hence $P \in \mathcal{P}_k \subseteq \mathcal{C}_k$ by (20). Let $\widetilde{P}_t = P_{\widetilde{x}_t}$. For any choice of kernels $P_t^s \in \mathcal{P}_{k_t}$,

$$\begin{aligned}\sum_{t,s,a} |q^{P_t^s, \pi_t}(s,a) - q_t(s,a)| &= \sum_{t,s} |q^{P_t^s, \pi_t}(s) - q_t(s)| \\ &\leq C_{\mathrm{tr}}\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big),\end{aligned}$$

where the equality follows from $q^{P',\pi_t}(s,a) = \pi_t(a \mid s) q^{P',\pi_t}(s)$ and $\sum_a \pi_t(a \mid s) = 1$, and the inequality from (27) in Lemma 10. Taking $P_t^s = \widetilde{P}_t$ for every s gives the bound on the first sum in (31).

For the second sum, choose

$$P_t^s \in \arg\max_{P' \in \mathcal{P}_{k_t}} q^{P',\pi_t}(s).$$

These maxima exist by compactness, and

$$|u_t(s) - q_t(s)| = |q^{P_t^s,\pi_t}(s) - q_t(s)|.$$

Thus, on the same event, the second sum in (31) is at most $C_{\mathrm{tr}}(HS\sqrt{AT}\,\iota + HS^4A\iota^2)$ by (27) in Lemma 10. Since both bounds are obtained by applying Lemma 10 to the same executed policies and counts, the lemma gives $\mathbb{P}(\mathcal{E}) \geq 1 - C_0 H\delta$, which completes the proof. □

B.5.6 Regret bound under unknown transitions (Proof of Theorem 8)

Proof of Theorem 8. If $T < A = K^2$, then $HSK\sqrt{T} \geq HT$, and the trivial bound $\mathsf{Reg}_T \leq HT$ proves the theorem. In the rest of the proof, we assume $T \geq A$. Define $q_\rho^* = q^{P,\pi_\rho^*}$ as in the known-transition proof. For each episode, let

$$\mathcal{E}_t = \{P \in \mathcal{P}_{k_t}\}.$$

These events are determined by the observations before episode t and satisfy $\mathcal{E}_1 \supseteq \mathcal{E}_2 \supseteq \cdots \supseteq \mathcal{E}_T$, whereas the event $\mathcal{E}$ in Lemma 12 depends on the entire interaction. We have

$$\begin{aligned}\mathbb{E}[\mathsf{Reg}_T] &\leq \rho H^2 T + \mathbb{E}\left[\sum_{t=1}^T \langle q_t - q_\rho^*, \ell_t\rangle\right] \\ &\leq \rho H^2 T + \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\mathcal{E}_t]\langle q_t - q_\rho^*, \ell_t\rangle\right] + HT\mathbb{P}(\mathcal{E}_T^c), \end{aligned} \tag{32}$$

where the first inequality follows from Lemma 7 and the second inequality from $V^\pi(s_0;\ell_t) \in [0,H]$ and $\mathcal{E}_t^c \subseteq \mathcal{E}_T^c$. For each episode t, we decompose the second term in (32) as

$$\langle q_t - q_\rho^*, \ell_t\rangle = \langle q_t - \widetilde{q}_t, \ell_t\rangle + \langle \widetilde{q}_t, \ell_t - \widehat{\ell}_t\rangle + \langle \widetilde{q}_t - q_\rho^*, \widehat{\ell}_t\rangle + \langle q_\rho^*, \widehat{\ell}_t - \ell_t\rangle. \tag{33}$$

We bound the four terms separately, each multiplied by $\mathbb{1}[\mathcal{E}_t]$ and summed over $t \in [T]$.

The first term of (33) satisfies

$$\langle q_t - \widetilde{q}_t, \ell_t\rangle \leq \sum_{s,a} |q_t(s,a) - \widetilde{q}_t(s,a)|. \tag{34}$$

For the second term of (33), since $\widetilde{q}_t$ and $\mathbb{1}[\mathcal{E}_t]$ are determined before episode t and states with $u_t(s) = 0$ contribute zero, Lemma 9 gives

$$\begin{aligned}\mathbb{E}_t\left[\mathbb{1}[\mathcal{E}_t]\langle \widetilde{q}_t, \ell_t - \widehat{\ell}_t\rangle\right] &= \mathbb{1}[\mathcal{E}_t] \sum_{s:u_t(s)>0} \sum_a \widetilde{q}_t(s,a)\left(1 - \frac{q_t(s)}{u_t(s)}\right)\ell_t(s,a) \\ &\leq \mathbb{1}[\mathcal{E}_t] \sum_{s:u_t(s)>0} \frac{\widetilde{q}_t(s)}{u_t(s)}(u_t(s) - q_t(s)) \\ &\leq \mathbb{1}[\mathcal{E}_t] \sum_s (u_t(s) - q_t(s)), \end{aligned} \tag{35}$$

where the first inequality follows from $\ell_t \leq 1$ and $q_t(s) \leq u_t(s)$ on $\mathcal{E}_t$, and the second inequality from $\widetilde{q}_t(s) \leq u_t(s)$ on $\mathcal{E}_t$.

For the third term of (33), define $n = \sum_{t=1}^T \mathbb{1}[\mathcal{E}_t]$. Since the events decrease, $\mathbb{1}[\mathcal{E}_t] = \mathbb{1}[t \leq n]$ for every $t \in [T]$ and every realization. For $t \leq n$, we have $P \in \mathcal{P}_{k_t}$, and hence $q_\rho^* \in \Omega_\rho \subseteq \Omega_\rho(\mathcal{P}_{k_t})$

by Appendix B.5.2. Since $\widehat{\ell}_t \geq 0$ as in the known-transition case, Lemma 6 applies to the first n episodes with $\mathcal{Q}_t = \Omega_\rho(\mathcal{P}_{k_t})$ and gives

$$\begin{aligned}\sum_{t=1}^{T} \mathbb{1}[\mathcal{E}_t]\langle \widetilde{q}_t - q_\rho^*, \widehat{\ell}_t\rangle &= \sum_{t=1}^{n} \langle \widetilde{q}_t - q_\rho^*, \widehat{\ell}_t\rangle \\ &\leq \frac{H\log(SA)}{\eta} + \frac{\eta}{2}\sum_{t=1}^{n}\sum_{s,a} \widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2 \\ &= \frac{H\log(SA)}{\eta} + \frac{\eta}{2}\sum_{t=1}^{T}\mathbb{1}[\mathcal{E}_t]\sum_{s,a} \widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2.\end{aligned}$$

This inequality holds for every realization of n. For each fixed t, since $\mathbb{1}[\mathcal{E}_t]$ is determined before episode t, we have

$$\begin{aligned}\mathbb{E}\left[\mathbb{1}[\mathcal{E}_t]\sum_{s,a} \widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2\right] &= \mathbb{E}\left[\mathbb{1}[\mathcal{E}_t]\mathbb{E}_t\left[\sum_{s,a} \widetilde{q}_t(s,a)\widehat{\ell}_t(s,a)^2\right]\right] \\ &\leq \mathbb{E}\left[\mathbb{1}[\mathcal{E}_t]\frac{SK}{\rho}\right] \leq \frac{SK}{\rho},\end{aligned}$$

where the first inequality follows from (25) in Lemma 9, which holds on $\mathcal{E}_t$. Consequently,

$$\mathbb{E}\left[\sum_{t=1}^{T} \mathbb{1}[\mathcal{E}_t]\langle \widetilde{q}_t - q_\rho^*, \widehat{\ell}_t\rangle\right] \leq \frac{H\log(SA)}{\eta} + \frac{\eta SKT}{2\rho}.$$

Finally, for the fourth term of (33), since the comparator q_ρ^* is fixed independently of the interaction and the estimate is optimistic on $\mathcal{E}_t$, that is, $\mathbb{E}_t[\widehat{\ell}_t(s,a)] \leq \ell_t(s,a)$ by (24) in Lemma 9, we have

$$\mathbb{E}\Big[\mathbb{1}[\mathcal{E}_t]\langle q_\rho^*, \widehat{\ell}_t - \ell_t\rangle\Big] = \mathbb{E}\left[\mathbb{1}[\mathcal{E}_t]\sum_{s,a} q_\rho^*(s,a)\Big(\mathbb{E}_t\big[\widehat{\ell}_t(s,a)\big] - \ell_t(s,a)\Big)\right] \leq 0.$$

We now take expectations of the first two terms. By Lemma 12, on the event $\mathcal{E}$,

$$\begin{aligned}\sum_{t=1}^{T}\mathbb{1}[\mathcal{E}_t]\sum_{s,a}|q_t(s,a) - \widetilde{q}_t(s,a)| &\leq C_{\mathrm{tr}}\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big), \\ \sum_{t=1}^{T}\mathbb{1}[\mathcal{E}_t]\sum_{s}(u_t(s) - q_t(s)) &\leq C_{\mathrm{tr}}\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big),\end{aligned}$$

whereas the two left-hand sides are at most $2HT$ and ST, respectively, on $\mathcal{E}^c$, since $\sum_{s,a}(q_t(s,a) + \widetilde{q}_t(s,a)) = 2H$ and $u_t(s) - q_t(s) \leq 1$. Since $\mathbb{P}(\mathcal{E}^c) \leq C_0H\delta$, taking expectations in (34) and (35) gives

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{T}\mathbb{1}[\mathcal{E}_t]\langle q_t - \widetilde{q}_t, \ell_t\rangle\right] &\leq C_{\mathrm{tr}}\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big) + 2C_0H^2T\delta, \\ \mathbb{E}\left[\sum_{t=1}^{T}\mathbb{1}[\mathcal{E}_t]\langle \widetilde{q}_t, \ell_t - \widehat{\ell}_t\rangle\right] &\leq C_{\mathrm{tr}}\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big) + C_0HST\delta.\end{aligned}$$

Substituting these bounds and those on the third and fourth terms into (32), and using $\mathbb{P}(\mathcal{E}_T^c) \leq 4\delta$ by (26), we obtain

$$\begin{aligned}\mathbb{E}[\mathsf{Reg}_T] \leq \frac{H\log(SA)}{\eta} + \frac{\eta SKT}{2\rho} + \rho H^2T \\ + 2C_{\mathrm{tr}}\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2\Big) + C_1H(3H+S)T\delta,\end{aligned}$$

where $C_1 = \max\{C_0, 4\}$. Choose $\delta = (HSA(T+1))^{-2}$ and choose η and ρ as in the proof of Theorem 1. Then the last term is bounded by a universal constant. As shown in the proof of Theorem 1, if $\sqrt{\eta SK}/H \le 1/2$, then the sum of the first three terms is at most $\frac{5}{2}(H^3SK\log(SA))^{1/3}T^{2/3}$, and otherwise $HT \le 2(H^3SK\log(SA))^{1/3}T^{2/3}$, so that the trivial bound $\mathrm{Reg}_T \le HT$ is of the claimed order.

Substituting $A = K^2$ gives $HS\sqrt{AT}\,\iota = HSK\sqrt{T}\,\iota$ and $HS^4A\iota^2 = HS^4K^2\iota^2$, and combining these bounds with the trivial bound $\mathrm{Reg}_T \le HT$ proves (16). This completes the proof. □

C Deferred proofs of the policy optimization approach

We first prove the additional estimates needed for policy optimization, then derive the regret bound under known transitions, using Lemmas 4 and 7. The unknown-transition algorithm in Appendix D uses the same regret decomposition.

Throughout, $\widehat{\ell}_t$ denotes (9), not (4). Recall that we use $\mathbb{E}_t$ and $\mathbb{P}_t(E) = \mathbb{E}_t[\mathbb{1}[E]]$ as in Appendix B, including the learner's past random choices in the conditioning.

C.1 Loss estimates

The sampling rule in Lines 6–7 does not require knowledge of P, and hence the conditional moment bounds below apply in both settings. Write $q_t = q^{P,\pi_t}$, $Q_t(s,a) = Q^{\pi_t}(s,a;\ell_t)$, and $V_t(s) = V^{\pi_t}(s;\ell_t)$. The true occupancy measure q_t is used only for analysis when P is unknown. Conditional expectations given a visit are used only when that event has positive conditional probability. For the analysis, define

$$\widehat{g}_t(s) = \sum_{a\in\mathcal{A}} \pi_t(a \mid s)\widehat{\ell}_t(s,a), \tag{36}$$

so that $\widehat{L}_{t,h} = \sum_{j=h}^{H-1} \widehat{g}_t(s_{t,j})$ in (10).

Lemma 13. *For every state s with positive visit probability and every $a \in \mathcal{A}$,*

$$\mathbb{E}_t\Big[\widehat{\ell}_t(s,a) \;\Big|\; \mathbb{I}_t(s)=1\Big] = \ell_t(s,a), \qquad \mathbb{E}_t\Big[\widehat{\ell}_t(s,a)^2 \;\Big|\; \mathbb{I}_t(s)=1\Big] \le K/\gamma, \tag{37}$$

$$\mathbb{E}_t[\widehat{g}_t(s) \mid \mathbb{I}_t(s)=1] = \sum_a \pi_t(a \mid s)\ell_t(s,a), \qquad \mathbb{E}_t\big[\widehat{g}_t(s)^2 \;\big|\; \mathbb{I}_t(s)=1\big] \le K/\gamma. \tag{38}$$

Proof. The sampling rule in Line 6 gives $\mathbb{P}_t(Z_t(s)=1 \mid \mathbb{I}_t(s)=1) = \gamma$. Using the conditional moments of $\widehat{b}_t$ from Lemma 4 in (9), we obtain

$$\begin{aligned}\mathbb{E}_t\Big[\widehat{\ell}_t(s,a) \;\Big|\; \mathbb{I}_t(s)=1\Big] &= -\frac{\gamma}{2\gamma}\mathbb{E}_t\Big[\widehat{b}_t(s,a^{\mathrm{L}}) + \widehat{b}_t(s,a^{\mathrm{R}}) \;\Big|\; \mathbb{I}_t(s)=1, Z_t(s)=1\Big] \\ &= -\frac{1}{2}\big(b_t(s,a^{\mathrm{L}}) + b_t(s,a^{\mathrm{R}})\big) = \ell_t(s,a), \end{aligned} \tag{39}$$

$$\mathbb{E}_t\Big[\widehat{\ell}_t(s,a)^2 \;\Big|\; \mathbb{I}_t(s)=1\Big] \le \frac{1}{2\gamma}\mathbb{E}_t\Big[\widehat{b}_t(s,a^{\mathrm{L}})^2 + \widehat{b}_t(s,a^{\mathrm{R}})^2 \;\Big|\; \mathbb{I}_t(s)=1, Z_t(s)=1\Big] \le \frac{K}{\gamma}, \tag{40}$$

where the first inequality follows from $(x+y)^2 \le 2x^2+2y^2$ and the last inequality from the second-moment bound in Lemma 4. For the average in (36),

$$\begin{aligned}\mathbb{E}_t[\widehat{g}_t(s) \mid \mathbb{I}_t(s)=1] &= \sum_a \pi_t(a \mid s)\mathbb{E}_t\Big[\widehat{\ell}_t(s,a) \;\Big|\; \mathbb{I}_t(s)=1\Big] = \sum_a \pi_t(a \mid s)\ell_t(s,a), \\ \mathbb{E}_t\big[\widehat{g}_t(s)^2 \;\big|\; \mathbb{I}_t(s)=1\big] &\le \sum_a \pi_t(a \mid s)\mathbb{E}_t\Big[\widehat{\ell}_t(s,a)^2 \;\Big|\; \mathbb{I}_t(s)=1\Big] \le K/\gamma, \end{aligned}$$

where the first equality follows from the fact that π_t is determined before episode t, the second equality from (39), the first inequality from Jensen's inequality, and the second inequality from (40), which completes the proof. □

C.2 Estimated losses at later steps

The variable $\widehat{L}_{t,h}$ in (10) sums the estimated losses from step h to the end of episode t. The next lemma computes its conditional expectation and bounds its second moment, separating the diagonal and cross terms.

Lemma 14. *For $s \in \mathcal{S}_h$ with $0 \le h < H$, let $n = H - h - 1$. Whenever the conditioning events have positive probability,*

$$\mathbb{E}_t\Big[\widehat{L}_{t,h} \;\Big|\; \mathbb{I}_t(s) = 1\Big] = V_t(s),$$

$$\mathbb{E}_t\Big[\widehat{L}_{t,h+1} \;\Big|\; \mathbb{I}_t(s,a) = 1\Big] = Q_t(s,a) - \ell_t(s,a), \tag{41}$$

$$\mathbb{E}_t\Big[\widehat{L}^2_{t,h+1} \;\Big|\; \mathbb{I}_t(s,a) = 1\Big] \le nK/\gamma + n(n-1). \tag{42}$$

Proof. Let $\mathcal{F}_{t,0}$ be the history before episode t, so that $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot \mid \mathcal{F}_{t,0}]$. For $0 \le j \le H$, let $\mathcal{F}_{t,j}$ additionally contain the variables $Z_t(s_{t,k})$, actions, and observations at steps $0, \dots, j-1$ and the states $s_{t,0}, \dots, s_{t,j}$. In particular, $\sigma(s_{t,j}) \subseteq \mathcal{F}_{t,j}$, whereas the variable $Z_t(s_{t,j})$, action, and feedback at step j are not included. We first prove $\mathbb{E}[\widehat{L}_{t,j} \mid \mathcal{F}_{t,j}] = V_t(s_{t,j})$ by backward induction. Both sides are zero at $j = H$. If the identity holds at $j+1$, then

$$\begin{aligned}
\mathbb{E}\Big[\widehat{L}_{t,j} \;\Big|\; \mathcal{F}_{t,j}\Big] &= \mathbb{E}[\widehat{g}_t(s_{t,j}) \mid \mathcal{F}_{t,j}] + \mathbb{E}\Big[\mathbb{E}\big[\widehat{L}_{t,j+1} \;\big|\; \mathcal{F}_{t,j+1}\big] \;\Big|\; \mathcal{F}_{t,j}\Big] \\
&= \sum_a \pi_t(a \mid s_{t,j})\ell_t(s_{t,j},a) + \sum_{a,s'} \pi_t(a \mid s_{t,j})P(s' \mid s_{t,j},a)V_t(s') \\
&= V_t(s_{t,j}),
\end{aligned} \tag{43}$$

where the first equality follows from (10), the second equality from the sampling rule in Line 6, the induction hypothesis, and $s_{t,j+1} \sim P(\cdot \mid s_{t,j}, a_{t,j})$, and the last equality from the definition of V_t. For a fixed $s \in \mathcal{S}_h$, the event $\{\mathbb{I}_t(s) = 1\} = \{s_{t,h} = s\}$ belongs to $\mathcal{F}_{t,h}$. Therefore,

$$\begin{aligned}
\mathbb{E}_t\Big[\widehat{L}_{t,h} \;\Big|\; \mathbb{I}_t(s) = 1\Big] &= \mathbb{E}_t\Big[\mathbb{E}\big[\widehat{L}_{t,h} \;\big|\; \mathcal{F}_{t,h}\big] \;\Big|\; \mathbb{I}_t(s) = 1\Big] \\
&= \mathbb{E}_t[V_t(s_{t,h}) \mid \mathbb{I}_t(s) = 1] = V_t(s),
\end{aligned}$$

where the first equality follows from $\{\mathbb{I}_t(s) = 1\} \in \mathcal{F}_{t,h}$, the second equality from (43), and the last equality from $s_{t,h} = s$ on that event and the fact that $V_t(s)$ is determined before episode t. For a fixed pair (s,a), the visit event instead belongs to $\mathcal{F}_{t,h+1}$, and

$$\begin{aligned}
\mathbb{E}_t\Big[\widehat{L}_{t,h+1} \;\Big|\; \mathbb{I}_t(s,a) = 1\Big] &= \mathbb{E}_t\Big[\mathbb{E}\big[\widehat{L}_{t,h+1} \;\big|\; \mathcal{F}_{t,h+1}\big] \;\Big|\; \mathbb{I}_t(s,a) = 1\Big] \\
&= \sum_{s' \in \mathcal{S}_{h+1}} P(s' \mid s,a)V_t(s') = Q_t(s,a) - \ell_t(s,a),
\end{aligned}$$

where the second equality follows from (43) at step $h+1$ and $s_{t,h+1} \sim P(\cdot \mid s,a)$ given $\mathbb{I}_t(s,a) = 1$, and the last equality from the definition of Q_t.

For the second moment, let $X_j = \widehat{g}_t(s_{t,j})$. The calculation of (38) in Lemma 13, conditional on the history $\mathcal{F}_{t,j}$ before sampling at step j, gives

$$0 \le \mathbb{E}[X_j \mid \mathcal{F}_{t,j}] \le 1, \qquad \mathbb{E}\big[X_j^2 \;\big|\; \mathcal{F}_{t,j}\big] \le K/\gamma. \tag{44}$$

Fix the event $E = \{\mathbb{I}_t(s,a) = 1\}$. For $h < k < j < H$, both E and X_k are determined by $\mathcal{F}_{t,j}$. Since $X_k \ge 0$, we have

$$\begin{aligned}
\mathbb{E}_t[X_k X_j \mid E] &= \mathbb{E}_t[X_k \mathbb{E}[X_j \mid \mathcal{F}_{t,j}] \mid E] \le \mathbb{E}_t[X_k \mid E] \\
&= \mathbb{E}_t[\mathbb{E}[X_k \mid \mathcal{F}_{t,k}] \mid E] \le 1, \qquad \mathbb{E}_t\big[X_k^2 \;\big|\; E\big] \le K/\gamma,
\end{aligned} \tag{45}$$

where the equalities follow from $E \in \mathcal{F}_{t,k} \subseteq \mathcal{F}_{t,j}$ and the $\mathcal{F}_{t,j}$-measurability of X_k, and the inequalities from (44) and $X_k \ge 0$. Expanding the square gives

$$\mathbb{E}_t\Big[\widehat{L}^2_{t,h+1} \;\Big|\; E\Big] = \sum_{k=h+1}^{H-1} \mathbb{E}_t\big[X_k^2 \;\big|\; E\big] + 2 \sum_{h<k<j<H} \mathbb{E}_t[X_k X_j \mid E] \le nK/\gamma + n(n-1),$$

where the inequality follows from (45) applied to the n diagonal terms and the $n(n-1)/2$ pairs of distinct steps. For $h = H-1$, $\widehat{L}_{t,H} = 0$ and the bound is zero as well, which completes the proof. □

C.3 Estimator under known transitions

We now take into account the denominator $q_t(s,a)+\delta$ in (10). In (10), the first term in the numerator estimates the immediate loss, and the second term estimates the losses after taking a at s.

Lemma 15. *The estimate in (10) satisfies*

$$\mathbb{E}_t\left[\widehat{Q}_t(s,a)\right] = \frac{q_t(s,a)}{q_t(s,a)+\delta} Q_t(s,a), \tag{46}$$

$$\mathbb{E}_t\left[\widehat{Q}_t(s,a)^2\right] \le \frac{2H(K/\gamma+H)}{q_t(s,a)+\delta}. \tag{47}$$

Proof. If $q_t(s)=0$, then $\widehat{Q}_t(s,a)=0$ almost surely, and the claims follow. Otherwise, $q_t(s,a)>0$ since $\pi_t(a \mid s) \ge \gamma/A$. Taking conditional expectations of the two terms in (10) gives

$$\begin{aligned}
\mathbb{E}_t\left[\widehat{Q}_t(s,a)\right] &= \frac{q_t(s)\pi_t(a \mid s)}{q_t(s,a)+\delta}\mathbb{E}_t\left[\widehat{\ell}_t(s,a) \;\middle|\; \mathbb{I}_t(s)=1\right] \\
&\quad + \frac{q_t(s,a)}{q_t(s,a)+\delta}\mathbb{E}_t\left[\widehat{L}_{t,h(s)+1} \;\middle|\; \mathbb{I}_t(s,a)=1\right] \\
&= \frac{q_t(s,a)}{q_t(s,a)+\delta}\big(\ell_t(s,a)+Q_t(s,a)-\ell_t(s,a)\big) \\
&= \frac{q_t(s,a)}{q_t(s,a)+\delta} Q_t(s,a),
\end{aligned}$$

where the second equality follows from (37) in Lemma 13, (41) in Lemma 14, and $q_t(s)\pi_t(a \mid s) = q_t(s,a)$.

We next prove (47) by bounding the two terms in $\widehat{Q}_t(s,a)$ separately:

$$\mathbb{E}_t\left[\widehat{Q}_t(s,a)^2\right] \le 2\mathbb{E}_t\left[\left(\frac{\mathbb{I}_t(s)\pi_t(a \mid s)\widehat{\ell}_t(s,a)}{q_t(s,a)+\delta}\right)^2\right] + 2\mathbb{E}_t\left[\left(\frac{\mathbb{I}_t(s,a)\widehat{L}_{t,h(s)+1}}{q_t(s,a)+\delta}\right)^2\right], \tag{48}$$

where the inequality follows from $(x+y)^2 \le 2x^2+2y^2$. We bound the first expectation on the right-hand side of (48) as

$$\begin{aligned}
\mathbb{E}_t\left[\left(\frac{\mathbb{I}_t(s)\pi_t(a \mid s)\widehat{\ell}_t(s,a)}{q_t(s,a)+\delta}\right)^2\right] &= \frac{q_t(s)\pi_t(a \mid s)^2}{(q_t(s,a)+\delta)^2}\mathbb{E}_t\left[\widehat{\ell}_t(s,a)^2 \;\middle|\; \mathbb{I}_t(s)=1\right] \\
&\le \frac{K}{\gamma(q_t(s,a)+\delta)},
\end{aligned} \tag{49}$$

where the inequality follows from (37) in Lemma 13 and $q_t(s)\pi_t(a \mid s)^2 \le q_t(s,a) \le q_t(s,a)+\delta$. For the second expectation on the right-hand side of (48), let $n = H-h(s)-1$. Using (42) in Lemma 14, we obtain

$$\begin{aligned}
\mathbb{E}_t\left[\left(\frac{\mathbb{I}_t(s,a)\widehat{L}_{t,h(s)+1}}{q_t(s,a)+\delta}\right)^2\right] &= \frac{q_t(s,a)}{(q_t(s,a)+\delta)^2}\mathbb{E}_t\left[\widehat{L}^2_{t,h(s)+1} \;\middle|\; \mathbb{I}_t(s,a)=1\right] \\
&\le \frac{nK/\gamma+n(n-1)}{q_t(s,a)+\delta}.
\end{aligned} \tag{50}$$

Substituting (49) and (50) into the right-hand side of (48) gives

$$\mathbb{E}_t\left[\widehat{Q}_t(s,a)^2\right] \le \frac{2((n+1)K/\gamma+n(n-1))}{q_t(s,a)+\delta} \le \frac{2H(K/\gamma+H)}{q_t(s,a)+\delta},$$

where the last inequality follows from $n+1 \le H$, which completes the proof. □

C.4 Regret decomposition

We first decompose the expected regret into a sum controlled by FTRL and a sum containing the bonuses. Adding the same uniform component to the comparator and to π_t allows us to bound the remaining terms by the expected bonus at the initial state.

For a best fixed policy π^*, define

$$\pi^*_\gamma(a \mid s) = (1-\gamma)\pi^*(a \mid s) + \gamma/A, \qquad q^*_\gamma = q^{P,\pi^*_\gamma}. \tag{51}$$

These quantities are fixed independently of the interaction under the oblivious-adversary assumption, and Lemma 7 with $\rho = \gamma$ gives

$$V^{\pi^*_\gamma}(s_0;\ell_t) - V^{\pi^*}(s_0;\ell_t) \le \gamma H^2. \tag{52}$$

Subtracting (51) from (7) yields

$$\pi_t(\cdot \mid s) - \pi^*_\gamma(\cdot \mid s) = (1-\gamma)\big(\widetilde{\pi}_t(\cdot \mid s) - \pi^*(\cdot \mid s)\big). \tag{53}$$

The regret is still measured against π^*. The following lemma combines (53) with the bonus argument of Luo et al. (2021). In its proof, the states are weighted by $q^*_\gamma(s)$, while the comparator in the FTRL bound at s is $\pi^*(\cdot \mid s)$.

Lemma 16. *Fix $\eta > 0$ and $\gamma \in (0, 1/2]$. Suppose that the policy update is (11) with $\widetilde{\pi}_1$ uniform, that the executed policy is $\pi_t(a \mid s) = (1-\gamma)\widetilde{\pi}_t(a \mid s) + \gamma/A$, that $\widehat{Q}_t \ge 0$, and that $m_t \ge 0$ and $0 \le M_t \le 9H^2$ are determined before episode t. Let $\tau \in \{1, \dots, T+1\}$ be such that $\{t < \tau\}$ is determined before episode t. On $\{t < \tau\}$, suppose that, for every nonterminal (s,a),*

$$\mathbb{E}_t\Big[\widehat{Q}_t(s,a)\Big] \le Q_t(s,a), \tag{54}$$

$$\sum_a \widetilde{\pi}_t(a \mid s)\Big(Q_t(s,a) - \mathbb{E}_t\big[\widehat{Q}_t(s,a)\big] + \eta\mathbb{E}_t\big[\widehat{Q}_t(s,a)^2\big]\Big) \le m_t(s). \tag{55}$$

Define $\bar{M}_t(s) = \sum_a \pi_t(a \mid s) M_t(s,a)$, with $\bar{M}_t(s_H) = 0$, and suppose also that, on $\{t < \tau\}$,

$$M_t(s,a) \ge m_t(s) + \sum_{s'} P(s' \mid s,a)\bar{M}_t(s'). \tag{56}$$

If $\eta \le 1/(9H^2)$, then

$$\mathbb{E}[\mathsf{Reg}_T] \le \gamma H^2 T + \frac{H \log A}{\eta} + 81\eta H^5 T + \mathbb{E}\left[\sum_{t<\tau} \bar{M}_t(s_0)\right] + HT\mathbb{P}(\tau \le T). \tag{57}$$

Condition (56) holds with equality for the known-transition bonus, and it follows from $P \in \mathcal{C}_{k_t}$ for the unknown-transition bonus. In the applications below, $\tau = T+1$ for known transitions, whereas (80) defines τ as the first confidence-set failure for unknown transitions.

Proof. Splitting the regret at τ and inserting π^*_γ gives

$$\begin{aligned}
\mathbb{E}[\mathsf{Reg}_T] &= \mathbb{E}\left[\sum_{t<\tau}\big(V^{\pi_t}(s_0;\ell_t) - V^{\pi^*}(s_0;\ell_t)\big)\right] + \mathbb{E}\left[\sum_{\tau\le t\le T}\big(V^{\pi_t}(s_0;\ell_t) - V^{\pi^*}(s_0;\ell_t)\big)\right] \\
&\le \gamma H^2 T + HT\mathbb{P}(\tau \le T) + \mathbb{E}\left[\sum_{t<\tau}\big(V^{\pi_t}(s_0;\ell_t) - V^{\pi^*_\gamma}(s_0;\ell_t)\big)\right],
\end{aligned}$$

where the inequality follows from (52) and $0 \le V^\pi(s_0;\ell_t) \le H$. We next decompose the expectation in the last bound:

$$\begin{aligned}
\mathbb{E}[\mathsf{Reg}_T] \le{}& \gamma H^2 T + HT\mathbb{P}(\tau \le T) \\
&+ (1-\gamma)\sum_s q^*_\gamma(s)\mathbb{E}\left[\sum_{t<\tau}\langle \widetilde{\pi}_t(\cdot \mid s) - \pi^*(\cdot \mid s), Q_t(s,\cdot) - M_t(s,\cdot)\rangle\right] \\
&+ \mathbb{E}\left[\sum_{t<\tau}\sum_s q^*_\gamma(s)\langle \pi_t(\cdot \mid s) - \pi^*_\gamma(\cdot \mid s), M_t(s,\cdot)\rangle\right],
\end{aligned} \tag{58}$$

where the inequality follows from the performance difference lemma and (53). We first bound the third term on the right-hand side of (58). For each state s, the FTRL update in (11) has losses $\widehat{Q}_t(s,\cdot) - M_t(s,\cdot)$ satisfying $\eta(\widehat{Q}_t(s,a) - M_t(s,a)) \geq -9\eta H^2 \geq -1$. By the standard regret bound for FTRL with the negative Shannon entropy regularizer (Lattimore & Szepesvári, 2020, Chapter 28), for every $n \in \{0, \dots, T\}$,

$$\sum_{t=1}^{n} \langle \widetilde{\pi}_t(\cdot \mid s) - \pi^*(\cdot \mid s), \widehat{Q}_t(s,\cdot) - M_t(s,\cdot)\rangle \leq \frac{\log A}{\eta} + \eta \sum_{t=1}^{n} \sum_{a} \widetilde{\pi}_t(a \mid s)(\widehat{Q}_t(s,a) - M_t(s,a))^2. \tag{59}$$

Subtracting and adding $\widehat{Q}_t(s,\cdot)$ gives

$$\begin{aligned}
&\langle \widetilde{\pi}_t(\cdot \mid s) - \pi^*(\cdot \mid s), Q_t(s,\cdot) - M_t(s,\cdot)\rangle \\
&= \langle \widetilde{\pi}_t(\cdot \mid s) - \pi^*(\cdot \mid s), \widehat{Q}_t(s,\cdot) - M_t(s,\cdot)\rangle \\
&\quad + \langle \widetilde{\pi}_t(\cdot \mid s), Q_t(s,\cdot) - \widehat{Q}_t(s,\cdot)\rangle + \langle \pi^*(\cdot \mid s), \widehat{Q}_t(s,\cdot) - Q_t(s,\cdot)\rangle.
\end{aligned} \tag{60}$$

For the third term in (60),

$$\begin{aligned}
&\mathbb{E}\left[\mathbb{1}[t < \tau] \sum_a \pi^*(a \mid s)(\widehat{Q}_t(s,a) - Q_t(s,a))\right] \\
&= \mathbb{E}\left[\mathbb{1}[t < \tau] \sum_a \pi^*(a \mid s)(\mathbb{E}_t\left[\widehat{Q}_t(s,a)\right] - Q_t(s,a))\right] \leq 0,
\end{aligned} \tag{61}$$

where the equality follows from the fact that $\mathbb{1}[t < \tau]$ and π^* are determined before episode t, and the inequality from (54) in the statement of this lemma. Taking $n = \tau - 1$ in (59) and substituting (60) and (61) gives

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t<\tau} \langle \widetilde{\pi}_t(\cdot \mid s) - \pi^*(\cdot \mid s), Q_t(s,\cdot) - M_t(s,\cdot)\rangle\right] \\
&\leq \frac{\log A}{\eta} + \mathbb{E}\left[\sum_{t<\tau}\sum_a \widetilde{\pi}_t(a \mid s)\Big(Q_t(s,a) - \mathbb{E}_t\left[\widehat{Q}_t(s,a)\right] + \eta \mathbb{E}_t\left[\widehat{Q}_t(s,a)^2\right]\Big)\right] + 81\eta H^4 T \\
&\leq \frac{\log A}{\eta} + \mathbb{E}\left[\sum_{t<\tau} m_t(s)\right] + 81\eta H^4 T,
\end{aligned}$$

where the first inequality follows from $(\widehat{Q}_t(s,a) - M_t(s,a))^2 \leq \widehat{Q}_t(s,a)^2 + M_t(s,a)^2$ and $M_t(s,a)^2 \leq 81H^4$, and the second inequality from (55) in the statement. Substituting this bound into (58) yields

$$\begin{aligned}
\mathbb{E}[\mathsf{Reg}_T] &\leq \gamma H^2 T + HT\mathbb{P}(\tau \leq T) + \frac{H \log A}{\eta} + 81\eta H^5 T \\
&\quad + \mathbb{E}\left[\sum_{t<\tau}\sum_s q^*_\gamma(s)\left((1-\gamma)m_t(s) + \bar{M}_t(s) - \sum_a \pi^*_\gamma(a \mid s) M_t(s,a)\right)\right],
\end{aligned} \tag{62}$$

where we used $(1-\gamma)\sum_s q^*_\gamma(s) \leq H$ and the definition of $\bar{M}_t$.

It remains to bound the last sum in (62). For every $s' \in \mathcal{S} \setminus \{s_0, s_H\}$, since q^*_γ is the occupancy measure of π^*_γ under P, we have

$$\sum_{s,a} q^*_\gamma(s)\pi^*_\gamma(a \mid s) P(s' \mid s,a) = q^*_\gamma(s'). \tag{63}$$

For every $t < \tau$, it follows that

$$\begin{aligned}
&\sum_s q^*_\gamma(s)\left((1-\gamma)m_t(s) + \bar{M}_t(s) - \sum_a \pi^*_\gamma(a \mid s) M_t(s,a)\right) \\
&\le \sum_s q^*_\gamma(s)\left(\bar{M}_t(s) - \sum_{a,s'} \pi^*_\gamma(a \mid s) P(s' \mid s,a)\bar{M}_t(s') - \gamma m_t(s)\right) \\
&= \bar{M}_t(s_0) - \gamma \sum_s q^*_\gamma(s) m_t(s) \le \bar{M}_t(s_0),
\end{aligned}$$

where the first inequality follows from (56) in the statement, the equality from (63), $q^*_\gamma(s_0) = 1$, and $\bar{M}_t(s_H) = 0$, and the last inequality from $m_t \ge 0$. Bounding the last sum in (62) by $\bar{M}_t(s_0)$ gives (57), which completes the proof. □

C.5 Regret bound under known transitions and parameter choices

We now prove Theorem 2 and specify the parameters used in both transition settings.

Proof of Theorem 2. We first verify the conditions of Lemma 16, with $\tau = T + 1$. We have $\widehat{Q}_t \ge 0$ by the definitions in (9) and (10), and (46) in Lemma 15 gives $\mathbb{E}_t[\widehat{Q}_t(s,a)] \le Q_t(s,a)$, which verifies (54). For (55), we have

$$\begin{aligned}
&\sum_a \widetilde{\pi}_t(a \mid s)\Big(Q_t(s,a) - \mathbb{E}_t\big[\widehat{Q}_t(s,a)\big] + \eta \mathbb{E}_t\big[\widehat{Q}_t(s,a)^2\big]\Big) \\
&\le H\big(\delta + 2\eta(K/\gamma + H)\big)\sum_a \frac{\widetilde{\pi}_t(a \mid s)}{q_t(s,a) + \delta} = m_t(s),
\end{aligned}$$

where the inequality follows from (46) and (47) in Lemma 15 and $Q_t(s,a) \le H$, and the equality from $\delta = 2\eta(K/\gamma + H)$. For the remaining conditions in Lemma 16, (8) gives $0 \le m_t(s) \le 2H$ and $0 \le M_t(s,a) \le 2H^2$, and (56) holds with equality. The lemma therefore applies whenever $\eta \le 1/(9H^2)$. We next bound $\bar{M}_t(s_0) = V^{\pi_t}(s_0; m_t)$ in (57):

$$\begin{aligned}
\bar{M}_t(s_0) &= 2\delta H \sum_{s,a} \frac{q_t(s)\widetilde{\pi}_t(a \mid s)}{q_t(s,a) + \delta} \\
&\le 4\delta H \sum_{s,a} \frac{q_t(s,a)}{q_t(s,a) + \delta} \le 4\delta HSA,
\end{aligned}$$

where the first inequality follows from $\widetilde{\pi}_t(a \mid s) \le \pi_t(a \mid s)/(1-\gamma) \le 2\pi_t(a \mid s)$ and $q_t(s)\pi_t(a \mid s) = q_t(s,a)$, and the second inequality from $q_t(s,a)/(q_t(s,a) + \delta) \le 1$. Substituting this bound into (57) of Lemma 16 gives

$$\begin{aligned}
\mathbb{E}[\mathsf{Reg}_T] &\le \gamma H^2 T + \frac{H \log A}{\eta} + 81\eta H^5 T + 4\delta HSAT \\
&\le \gamma H^2 T + \frac{H \log A}{\eta} + \eta T\left(81H^5 + 24H^2SK^2 + \frac{24HSK^3}{\gamma}\right), \qquad (64)
\end{aligned}$$

where, in the second inequality, we substituted $\delta = 2\eta(K/\gamma + H)$ and $A = K^2$ and replaced the resulting coefficients 8 by 24 so that the same parameter choice can be used in both settings.

For both algorithms, choose

$$\begin{aligned}
\gamma &= \min\left\{\frac{1}{2}, \left(\frac{24SK^3 \log A}{H^2 T}\right)^{1/3}\right\}, \\
\eta &= \min\left\{\frac{1}{9H^2}, \sqrt{\frac{H \log A}{T(81H^5 + 24H^2SK^2 + 24HSK^3/\gamma)}}\right\}, \qquad \delta = 2\eta(K/\gamma + H). \qquad (65)
\end{aligned}$$

To substitute these parameters into (64), abbreviate $D = H \log A$, $B_0 = 81H^5 + 24H^2SK^2$, and $B_1 = 24HSK^3$ in this calculation only. If neither truncation in (65) is active, the right-hand side of (64) is at most

$$\begin{aligned}\gamma H^2 T + 2\sqrt{DT(B_0 + B_1/\gamma)} &\le \gamma H^2 T + 2\sqrt{DB_1 T/\gamma} + 2\sqrt{DB_0 T} \\ &= 3(DB_1 H^2)^{1/3} T^{2/3} + 2\sqrt{DB_0 T}, \end{aligned} \tag{66}$$

where the inequality follows from $\sqrt{x+y} \le \sqrt{x} + \sqrt{y}$ for $x, y \ge 0$, and the equality from substituting the untruncated value of γ in (65). Moreover,

$$DB_1 H^2 = 24H^4 SK^3 \log A, \qquad DB_0 = (81H^6 + 24H^3 SK^2) \log A. \tag{67}$$

If the untruncated value of γ exceeds $1/2$, then $(DB_1H^2)^{1/3}T^{2/3} > H^2T/2 \ge HT/2$. If the untruncated value of η exceeds $1/(9H^2)$, then $D > T(B_0 + B_1/\gamma)/(81H^4) \ge HT$ and hence $\sqrt{DB_0T} > 9H^3T \ge HT$. In either case the trivial bound $\mathsf{Reg}_T \le HT$ is within a universal factor of the claimed minimum. Substituting (67) into (66) and using $\log A = 2 \log K$ gives the bound in Theorem 2, which completes the proof. □

D Policy optimization under unknown transitions

We retain $\widehat{\ell}_t$ in (9), $\widehat{L}_{t,h}$ in (10), and the policy update in (11). The only changes are the visit probabilities used in $\widehat{Q}_t$ and the bonus computed before each episode.

Theorem 17. *Suppose that P is fixed but unknown. With the parameter choices specified in the proof, Algorithm 4 satisfies, for every $T \ge 1$,*

$$\mathbb{E}[\mathsf{Reg}_T] = \widetilde{O}\Big(\min\Big\{HT,\ (H^4SK^3)^{1/3}T^{2/3} + H^2SK\sqrt{T} + H^2S^4K^2\Big\}\Big). \tag{68}$$

The last term depends on T only through the logarithmic factors hidden in $\widetilde{O}$. As in Theorem 8, it comes from the lower-order term of Lemma 10, and we have not attempted to optimize it.

D.1 Algorithm

The confidence sets specify the occupancy bounds used in the estimates and the transition kernels used to compute the bonuses.

Transition confidence sets. We use the counts N, M, N_k, M_k, the epochs, the empirical transition kernels, and the confidence sets $\mathcal{C}_k$ of Appendix B.5, defined in (17)–(19). The counts here are generated by the policy optimization algorithm's own trajectories. In those definitions, replace the confidence parameter denoted by δ with

$$\delta' = \frac{1}{20H^3ST}, \qquad \iota = \log\frac{TSA}{\delta'}. \tag{69}$$

Here δ' controls the probability that a transition confidence set fails to contain P, while $\delta = 2\eta(K/\gamma + H)$ remains the additive term in the denominators of the Q-function estimate. We use $\mathcal{C}_{k_t}$ directly, without the intersections in (20), since the FTRL decision set here is always $\Delta(\mathcal{A})$. These sets are nonempty: for a visited pair (s,a), the empirical distribution $\bar{P}_k(\cdot \mid s,a)$ satisfies the constraints on $P'(\cdot \mid s,a)$, and for an unvisited pair the radius β_k exceeds one, so that every distribution satisfies them. The initialization appears in Line 1 and the count updates in Lines 9–12 of Algorithm 4.

Occupancy bounds and Q-function estimates. Before episode t, the algorithm computes the following occupancy bounds (Line 4), in the notation $q^{P',\pi}$ of Appendix B.5:

$$\bar{q}_t(s) = \max_{P' \in \mathcal{C}_{k_t}} q^{P',\pi_t}(s), \qquad \underline{q}_t(s) = \min_{P' \in \mathcal{C}_{k_t}} q^{P',\pi_t}(s), \qquad \begin{cases} \bar{q}_t(s,a) = \pi_t(a \mid s)\bar{q}_t(s), \\ \underline{q}_t(s,a) = \pi_t(a \mid s)\underline{q}_t(s). \end{cases} \tag{70}$$

The maximization is the same as in (22), with the policy fixed and the kernel ranging over $\mathcal{C}_{k_t}$, and the minimization is analogous. Both can be computed by the dynamic program described in

Algorithm 4: Policy optimization under unknown transitions

Input: $\eta > 0$, $\gamma \in (0, 1/2]$, $\delta = 2\eta(K/\gamma + H)$, and δ' in (69)

1 Set $\widetilde{\pi}_1(a \mid s) = 1/A$, $k = 1$, all counts N, M, N_1, M_1 to zero, and $\mathcal{C}_1$ equal to all layered kernels;

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

3 Set $k_t = k$ and $\pi_t(a \mid s) = (1-\gamma)\widetilde{\pi}_t(a \mid s) + \gamma/A$;

4 Compute $\bar{q}_t, \underline{q}_t$ by (70);

5 Compute m_t, M_t by (72) and (73);

6 Execute π_t as in Lines 5–7 of Algorithm 2;

7 Compute $\widehat{\ell}_t$ and $\widehat{L}_{t,h}$ as in Lines 8–9 of Algorithm 2, and compute $\widehat{Q}_t$ by (71);

8 Compute $\widetilde{\pi}_{t+1}$ by (11);

9 Increment $N(s,a)$ and $M(s,a,s')$ for each observed transition (s,a,s');

10 **if** $t < T$ *and* $N(s,a) \geq \max\{1, 2N_k(s,a)\}$ *for a visited* (s,a) **then**

11 Set $k \leftarrow k+1$, $N_k \leftarrow N$, and $M_k \leftarrow M$;

12 Form $\mathcal{C}_k$ by (17)–(19), with confidence parameter δ';

Appendix B.5. The relation $q^{P',\pi_t}(s,a) = \pi_t(a \mid s) q^{P',\pi_t}(s)$ shows that a kernel attaining the maximum (or the minimum) at a state s attains it at every pair (s,a) as well.

After observing the episode, the algorithm computes $\widehat{\ell}_t$ and $\widehat{L}_{t,h}$ from (9) and (10) and sets

$$\widehat{Q}_t(s,a) = \frac{\mathbb{I}_t(s)\pi_t(a \mid s)\widehat{\ell}_t(s,a) + \mathbb{I}_t(s,a)\widehat{L}_{t,h(s)+1}}{\bar{q}_t(s,a) + \delta}. \tag{71}$$

As in (10), the first term uses visits to s and the second uses visits to (s,a). The only difference from (10) is that the unknown $q_t(s,a)$ in the denominator is replaced by its upper bound $\bar{q}_t(s,a)$. When $P \in \mathcal{C}_{k_t}$, (75) in Lemma 18 gives $\mathbb{E}_t[\widehat{Q}_t(s,a)] \leq Q_t(s,a)$ and bounds the difference $Q_t(s,a) - \mathbb{E}_t[\widehat{Q}_t(s,a)]$ in terms of δ and $\bar{q}_t(s,a) - \underline{q}_t(s,a)$.

Bonus term. Define $M_t(s_H, \cdot) = 0$ and compute the following quantities backward through the layers:

$$m_t(s) = H \sum_a \widetilde{\pi}_t(a \mid s) \frac{2\delta + \bar{q}_t(s,a) - \underline{q}_t(s,a)}{\bar{q}_t(s,a) + \delta}, \tag{72}$$

$$M_t(s,a) = m_t(s) + \left(1 + \frac{1}{H}\right) \max_{P' \in \mathcal{C}_{k_t}} \sum_{s'} P'(s' \mid s,a) \sum_{a'} \pi_t(a' \mid s') M_t(s',a'). \tag{73}$$

The difference $\bar{q}_t(s,a) - \underline{q}_t(s,a)$ in (72) accounts for the additional bias from estimating P. The definition in (73) follows Luo et al. (2021), and the factor $1 + 1/H$ increases the weight of m_t at each later step. Whenever $P \in \mathcal{C}_{k_t}$, the maximization includes P, and hence (56) in Lemma 16 holds.

The policy, occupancy bounds, and bonuses are computed using only observations from preceding episodes (Lines 3–5). The algorithm then collects feedback and updates the policy (Lines 6–8). The count updates in Lines 9–12 are the same as those in Algorithm 3, except that no intersection of confidence sets is taken. The policy update is not restarted when a new epoch starts.

D.2 Estimation and transition uncertainty

The moments of $\widehat{L}_{t,h}$ in Lemma 14 do not require knowledge of P. We now account for the upper bound $\bar{q}_t(s,a)$ in the denominator of (71).

Lemma 18. *Suppose that $P \in \mathcal{C}_{k_t}$. Then the estimate in (71) satisfies*

$$\mathbb{E}_t\Big[\widehat{Q}_t(s,a)\Big] = \frac{q_t(s,a)}{\bar{q}_t(s,a)+\delta} Q_t(s,a), \tag{74}$$

$$0 \le Q_t(s,a) - \mathbb{E}_t\Big[\widehat{Q}_t(s,a)\Big] \le \frac{H(\delta + \bar{q}_t(s,a) - \underline{q}_t(s,a))}{\bar{q}_t(s,a)+\delta}, \tag{75}$$

$$\mathbb{E}_t\Big[\widehat{Q}_t(s,a)^2\Big] \le \frac{2H(K/\gamma + H)}{\bar{q}_t(s,a)+\delta}. \tag{76}$$

Proof. If $q_t(s) = 0$, then $\widehat{Q}_t(s,a) = 0$ almost surely, and (74) holds since $q_t(s,a) = 0$. Otherwise, $q_t(s,a) > 0$ since $\pi_t(a \mid s) \ge \gamma/A$, and taking conditional expectations of the two terms in (71) gives

$$\begin{aligned}\mathbb{E}_t\Big[\widehat{Q}_t(s,a)\Big] &= \frac{q_t(s)\pi_t(a \mid s)}{\bar{q}_t(s,a)+\delta}\mathbb{E}_t\Big[\widehat{\ell}_t(s,a) \;\Big|\; \mathbb{I}_t(s) = 1\Big] + \frac{q_t(s,a)}{\bar{q}_t(s,a)+\delta}\mathbb{E}_t\Big[\widehat{L}_{t,h(s)+1} \;\Big|\; \mathbb{I}_t(s,a) = 1\Big] \\ &= \frac{q_t(s,a)}{\bar{q}_t(s,a)+\delta}\big(\ell_t(s,a) + Q_t(s,a) - \ell_t(s,a)\big) = \frac{q_t(s,a)}{\bar{q}_t(s,a)+\delta} Q_t(s,a),\end{aligned}$$

where the second equality follows from (37) in Lemma 13, (41) in Lemma 14, and $q_t(s)\pi_t(a \mid s) = q_t(s,a)$. This proves (74).

Under the assumption $P \in \mathcal{C}_{k_t}$, we have $\underline{q}_t(s,a) \le q_t(s,a) \le \bar{q}_t(s,a)$, and hence the coefficient in (74) lies in $[0,1]$. Since $Q_t(s,a) \in [0,H]$, (74) gives

$$0 \le Q_t(s,a) - \mathbb{E}_t\Big[\widehat{Q}_t(s,a)\Big] = \frac{\bar{q}_t(s,a) + \delta - q_t(s,a)}{\bar{q}_t(s,a)+\delta} Q_t(s,a) \le \frac{H(\delta + \bar{q}_t(s,a) - \underline{q}_t(s,a))}{\bar{q}_t(s,a)+\delta},$$

where the last inequality follows from $q_t(s,a) \ge \underline{q}_t(s,a)$. This proves (75), including states with $q_t(s) = 0$.

We next prove (76). It suffices to consider $q_t(s) > 0$. Let $n = H - h(s) - 1$. Bounding the squares of the two terms in (71) separately gives

$$\begin{aligned}\mathbb{E}_t\Big[\widehat{Q}_t(s,a)^2\Big] &\le \frac{2q_t(s)\pi_t(a \mid s)^2}{(\bar{q}_t(s,a)+\delta)^2}\mathbb{E}_t\Big[\widehat{\ell}_t(s,a)^2 \;\Big|\; \mathbb{I}_t(s) = 1\Big] \\ &\quad + \frac{2q_t(s,a)}{(\bar{q}_t(s,a)+\delta)^2}\mathbb{E}_t\Big[\widehat{L}^2_{t,h(s)+1} \;\Big|\; \mathbb{I}_t(s,a) = 1\Big] \\ &\le \frac{2K/\gamma + 2(nK/\gamma + n(n-1))}{\bar{q}_t(s,a)+\delta} \le \frac{2H(K/\gamma + H)}{\bar{q}_t(s,a)+\delta},\end{aligned}$$

where the first inequality follows from $(x+y)^2 \le 2x^2 + 2y^2$, the second inequality from (37) in Lemma 13, (42) in Lemma 14, and $q_t(s)\pi_t(a \mid s)^2 \le q_t(s,a) \le \bar{q}_t(s,a) + \delta$, and the last inequality from $n + 1 \le H$. At a state with $q_t(s) = 0$, $\widehat{Q}_t(s,a) = 0$ almost surely and the bound holds without conditioning, which completes the proof. □

D.3 Occupancy bounds from transition confidence sets

We use Lemma 10 with the executed policies of Algorithm 4 and confidence parameter δ' to bound the expected sum of the differences between the maximum and minimum in (70).

Lemma 19. *For the confidence sets above, there is a universal constant C such that*

$$\mathbb{P}\{P \notin \mathcal{C}_{k_t} \text{ for some } t \le T\} \le CH\delta', \tag{77}$$

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{s,a}(\bar{q}_t(s,a) - \underline{q}_t(s,a))\right] \le C\Big(HS\sqrt{AT}\,\iota + HS^4A\iota^2 + HST\delta'\Big). \tag{78}$$

Proof. By (26) with confidence parameter δ', the probability in (77) is at most $4\delta' \le CH\delta'$. If $T < A$, then $\sum_{t,s,a}(\bar{q}_t(s,a) - \underline{q}_t(s,a)) \le ST \le HS\sqrt{AT}$. It therefore remains to prove (78) for

$T \geq A$. Let $\mathcal{E}$ be the event on which $P \in \mathcal{C}_k$ for every epoch k and (27) holds for every choice of the kernels P_t^s, with confidence parameter δ'. By Lemma 10, $\mathbb{P}(\mathcal{E}^c) \leq C_0 H \delta'$.

Choose kernels $P_t^{s,+}, P_t^{s,-} \in \mathcal{C}_{k_t}$ such that

$$q^{P_t^{s,+},\pi_t}(s) = \bar{q}_t(s), \qquad q^{P_t^{s,-},\pi_t}(s) = \underline{q}_t(s). \tag{79}$$

Such kernels exist since $\mathcal{C}_{k_t}$ is compact. For each episode,

$$\begin{aligned}
\sum_{s,a}(\bar{q}_t(s,a) - \underline{q}_t(s,a)) &= \sum_s \sum_a \pi_t(a \mid s)(\bar{q}_t(s) - \underline{q}_t(s)) \\
&= \sum_s (\bar{q}_t(s) - \underline{q}_t(s)) \\
&\leq \sum_s \left| q^{P_t^{s,+},\pi_t}(s) - q_t(s) \right| + \sum_s \left| q^{P_t^{s,-},\pi_t}(s) - q_t(s) \right|,
\end{aligned}$$

where the first equality follows from (70), the second equality from $\sum_a \pi_t(a \mid s) = 1$, and the inequality from (79) and the triangle inequality after adding and subtracting $q_t(s)$. On $\mathcal{E}$, summing this inequality over t gives

$$\sum_{t=1}^{T} \sum_{s,a} (\bar{q}_t(s,a) - \underline{q}_t(s,a)) \leq 2C\Big(HS\sqrt{AT}\,\iota + HS^4 A\iota^2\Big),$$

where we used (27) in Lemma 10 once for each of the two choices $P_t^{s,+}$ and $P_t^{s,-}$. Regardless of $\mathcal{E}$, both $\bar{q}_t(s)$ and $\underline{q}_t(s)$ lie in $[0,1]$, and hence the left-hand side is at most ST. Consequently,

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{T} \sum_{s,a} (\bar{q}_t(s,a) - \underline{q}_t(s,a))\right] \\
&\leq 2C\Big(HS\sqrt{AT}\,\iota + HS^4 A\iota^2\Big) + ST\mathbb{P}(\mathcal{E}^c) \\
&\leq 2C\Big(HS\sqrt{AT}\,\iota + HS^4 A\iota^2\Big) + CHST\delta',
\end{aligned}$$

where the last inequality follows from $\mathbb{P}(\mathcal{E}^c) \leq CH\delta'$ by Lemma 10. Increasing C gives (78), which completes the proof. □

D.4 Regret analysis (Proof of Theorem 17)

The proof combines Lemmas 16 and 19.

Proof of Theorem 17. We first verify the conditions of Lemma 16 before the first confidence-set failure. Define

$$\tau = \min\big(\{t \in [T] : P \notin \mathcal{C}_{k_t}\} \cup \{T+1\}\big). \tag{80}$$

The time τ is used only in the proof, and the learner does not need to detect a confidence-set failure. For the fixed true kernel P, $\{t < \tau\} = \bigcap_{u=1}^{t} \{P \in \mathcal{C}_{k_u}\}$ is determined before episode t. The moment bounds in Lemmas 13 and 14 hold whether or not P belongs to the confidence set.

For $t < \tau$, we have $\widehat{Q}_t \geq 0$ by the definitions in (9) and (71), and (75) in Lemma 18 gives $\mathbb{E}_t[\widehat{Q}_t(s,a)] \leq Q_t(s,a)$, as required by (54). For (55), we have

$$\begin{aligned}
&\sum_a \widetilde{\pi}_t(a \mid s)\Big(Q_t(s,a) - \mathbb{E}_t\big[\widehat{Q}_t(s,a)\big] + \eta \mathbb{E}_t\big[\widehat{Q}_t(s,a)^2\big]\Big) \\
&\leq H \sum_a \widetilde{\pi}_t(a \mid s) \frac{\delta + \bar{q}_t(s,a) - \underline{q}_t(s,a) + 2\eta(K/\gamma + H)}{\bar{q}_t(s,a) + \delta} = m_t(s),
\end{aligned}$$

where the inequality follows from (75) and (76) in Lemma 18, and the equality from $\delta = 2\eta(K/\gamma + H)$ and (72). We next verify the remaining conditions on m_t and M_t in Lemma 16. Since $0 \leq$

$2\delta + \bar{q}_t(s,a) - \underline{q}_t(s,a) \le 3(\bar{q}_t(s,a) + \delta)$ and $\sum_a \widetilde{\pi}_t(a \mid s) = 1$, we have $0 \le m_t(s) \le 3H$ in (72), and unrolling (73) gives

$$0 \le M_t(s,a) \le 3H \sum_{j=0}^{H-h(s)-1} (1+1/H)^j \le 3\mathrm{e}H^2 \le 9H^2,$$

where we used $(1+1/H)^H \le \mathrm{e} < 3$. For $t < \tau$, we have $P \in \mathcal{C}_{k_t}$, and hence

$$M_t(s,a) \ge m_t(s) + (1+1/H)\sum_{s'} P(s' \mid s,a)\bar{M}_t(s') \ge m_t(s) + \sum_{s'} P(s' \mid s,a)\bar{M}_t(s'),$$

where the first inequality follows from $P \in \mathcal{C}_{k_t}$ in (73), and the second inequality from $\bar{M}_t \ge 0$. We have verified the conditions of Lemma 16, whose bound (57) gives

$$\mathbb{E}[\mathsf{Reg}_T] \le \gamma H^2 T + \frac{H \log A}{\eta} + 81\eta H^5 T + \mathbb{E}\left[\sum_{t<\tau} \bar{M}_t(s_0)\right] + HT\mathbb{P}(\tau \le T). \tag{81}$$

To bound the fourth term on the right-hand side of (81), choose $\widetilde{P}_t \in \mathcal{C}_{k_t}$ satisfying

$$\sum_{s'} \widetilde{P}_t(s' \mid s,a)\bar{M}_t(s') = \max_{P' \in \mathcal{C}_{k_t}} \sum_{s'} P'(s' \mid s,a)\bar{M}_t(s') \quad \text{for every } (s,a),$$

where such a choice exists since the constraints in (19) are imposed separately on each $P'(\cdot \mid s,a)$. Expanding (73) with $\widetilde{P}_t$ gives

$$\begin{aligned}\bar{M}_t(s_0) &= \sum_{h=0}^{H-1}(1+1/H)^h \sum_{s \in \mathcal{S}_h} q^{\widetilde{P}_t,\pi_t}(s) m_t(s) \\ &\le \mathrm{e}\, V^{\widetilde{P}_t,\pi_t}(s_0; m_t) \le 3\sum_s q^{\widetilde{P}_t,\pi_t}(s) m_t(s),\end{aligned} \tag{82}$$

where the first inequality follows from $(1+1/H)^h \le \mathrm{e}$, and the second inequality from $\mathrm{e} < 3$ and the definition of the value function. Substituting m_t from (72) yields

$$\begin{aligned}\sum_s q^{\widetilde{P}_t,\pi_t}(s) m_t(s) &= H \sum_{s,a} \frac{q^{\widetilde{P}_t,\pi_t}(s)\widetilde{\pi}_t(a \mid s)(2\delta + \bar{q}_t(s,a) - \underline{q}_t(s,a))}{\bar{q}_t(s,a) + \delta} \\ &\le 2H \sum_{s,a} \frac{q^{\widetilde{P}_t,\pi_t}(s,a)(2\delta + \bar{q}_t(s,a) - \underline{q}_t(s,a))}{\bar{q}_t(s,a) + \delta} \\ &\le 4\delta HSA + 2H \sum_{s,a} (\bar{q}_t(s,a) - \underline{q}_t(s,a)),\end{aligned} \tag{83}$$

where the first inequality follows from $\widetilde{\pi}_t(a \mid s) \le 2\pi_t(a \mid s)$ and $q^{\widetilde{P}_t,\pi_t}(s)\pi_t(a \mid s) = q^{\widetilde{P}_t,\pi_t}(s,a)$, and the second inequality from $q^{\widetilde{P}_t,\pi_t}(s,a)/(\bar{q}_t(s,a) + \delta) \le 1$, which holds by (70).

Substituting (82) and (83) into (81) and extending the sums over $t < \tau$ to $t \le T$, which is possible since these terms are nonnegative, gives

$$\begin{aligned}\mathbb{E}[\mathsf{Reg}_T] \le{}& \gamma H^2 T + \frac{H \log A}{\eta} + 81\eta H^5 T + 12\delta HSAT \\ &+ 6H\mathbb{E}\left[\sum_{t=1}^{T}\sum_{s,a}(\bar{q}_t(s,a) - \underline{q}_t(s,a))\right] + HT\mathbb{P}(\tau \le T).\end{aligned} \tag{84}$$

The first four terms in (84) equal the right-hand side of (64) after substituting $\delta = 2\eta(K/\gamma + H)$ and $A = K^2$. For the remaining two terms, (77) and (78) in Lemma 19 give

$$\begin{aligned}&6H\mathbb{E}\left[\sum_{t=1}^{T}\sum_{s,a}(\bar{q}_t(s,a) - \underline{q}_t(s,a))\right] + HT\mathbb{P}(\tau \le T) \\ &\le C\left(H^2 S\sqrt{AT}\,\iota + H^2 S^4 A\iota^2 + H^2 ST\delta'\right) \\ &\le C\left(H^2 S\sqrt{AT}\,\iota + H^2 S^4 A\iota^2 + 1\right),\end{aligned} \tag{85}$$

where the last inequality follows from $\delta' = 1/(20H^3ST)$ in (69). When neither truncation in (65) is active, substituting (85) into (84) and using (66) and (67) in the proof of Theorem 2 yields

$$\mathbb{E}[\mathsf{Reg}_T] \le O\Big((H^4SK^3\log K)^{1/3}T^{2/3} + (H^3 + H^{3/2}K\sqrt{S})\sqrt{T\log K} + H^2SK\sqrt{T}\,\iota + H^2S^4K^2\iota^2 + 1\Big).$$

Since $S \ge H+1$ and $K \ge 2$, we have $H^3 + H^{3/2}K\sqrt{S} \le 2H^2SK$, and the square-root terms are thus absorbed into $H^2SK\sqrt{T}\,\iota$. If either truncation is active, the same comparison with HT as in the proof of Theorem 2 applies. Combining these bounds with $\mathsf{Reg}_T \le HT$ gives (68), which completes the proof. □

The second term inside the parentheses in (85) comes from the same bound on the occupancy measures as for global optimization (Lemma 10), and the additional factor H comes from the coefficient $6H$ in (84). This term is dominated by $H^2S\sqrt{AT}\,\iota$ when $T \ge S^6A\iota^2$. We keep it explicit in (68) and do not claim that its dependence on S is sharp.

E Proof of the regret lower bound (Theorem 3)

This section provides the proof of Theorem 3. Throughout, $\mathrm{D}(\mathbb{P}, \mathbb{P}')$ denotes the Kullback–Leibler (KL) divergence and $\mathrm{kl}(p, q)$ the KL divergence between Bernoulli distributions with means p and q.

Proof of Theorem 3. Fix a learning algorithm. We construct a layered MDP in which comparisons that provide information about a Borda winner force the learner into states with larger losses for the rest of the episode. The transition kernel is known, and each state's preference matrix is fixed over all episodes.

The instance is shown in Figure 1, and its construction is given in the first three steps below. We then lower bound the regret averaged over a random choice of the preference matrices in two ways, through the perturbation of the matrices, which the learner cannot detect without many comparisons that reveal the Borda winner, and through the losses at the penalty states, which every such comparison incurs for the rest of the episode. Adding the two bounds and choosing the perturbation size gives the theorem.

The instance has M chains, and the $H-1$ intermediate layers of each chain are split into a first part of L layers and a final part of R layers (Figure 1), where

$$M = \left\lfloor \frac{S-2}{2(H-1)} \right\rfloor, \qquad L = \left\lfloor \frac{H-1}{2} \right\rfloor, \qquad R = H - 1 - L\,.$$

The assumption $S \ge 2H$ ensures that $M \ge 1$. The preference matrices in the first part are unknown to the learner, and the final part penalizes the comparisons that reveal them.

States and transitions. For every $c \in [M]$ and $h \in [H-1]$, introduce two states $x_{c,h}$ and $z_{c,h}$ in layer h (Figure 1(b)). We call the states $z_{c,h}$ penalty states. Together with s_0 and s_H, this uses $2 + 2M(H-1) \le S$ states. If the inequality is strict, add unreachable states to the intermediate layers, with arbitrary transitions to the next layer. Let

$$\mathcal{X} = \{x_{c,h} : c \in [M],\ h \in [L]\}$$

be the set of the states $x_{c,h}$ in the first L intermediate layers.

Recall that $\mathcal{A} = [K] \times [K]$. Let $\widetilde{K} = \lfloor K/2 \rfloor$, $\mathcal{K}_g = [\widetilde{K}]$, and $\mathcal{K}_b = [K] \setminus \mathcal{K}_g$. Define

$$\mathcal{A}_b = \mathcal{A} \setminus (\mathcal{K}_g \times \mathcal{K}_g)$$

so that $a \in \mathcal{A}_b$ exactly when at least one coordinate of a belongs to $\mathcal{K}_b$.

For every action a, the transition from s_0 is uniform over $x_{1,1}, \dots, x_{M,1}$ (Figure 1(a)). For $c \in [M]$, $a \in \mathcal{A}$, and $h \in [H-2]$, the remaining nonzero entries of the transition kernel P are

$$\begin{aligned} P(x_{c,h+1} \mid x_{c,h}, a) &= 1 \quad \text{if } h \le L \text{ and } a \notin \mathcal{A}_b\,, \\ P(z_{c,h+1} \mid x_{c,h}, a) &= 1 \quad \text{if } h \le L \text{ and } a \in \mathcal{A}_b\,, \\ P(x_{c,h+1} \mid x_{c,h}, a) &= 1 \quad \text{if } h > L\,, \\ P(z_{c,h+1} \mid z_{c,h}, a) &= 1 \quad \text{for every } h \in [H-2]\,. \end{aligned}$$

All states in layer $H-1$ lead to s_H. Hence a pair in $\mathcal{A}_b$ selected at a state in $\mathcal{X}$ moves the trajectory to the penalty states, where it stays for the rest of the episode, and in particular for the final R layers. The transition kernel is the same for every choice of the preference matrices below, so that knowing it gives the learner no information about the matrices.

Preference matrices. We now define the matrices placed at the states in $\mathcal{X}$. Let $K_{\mathrm{b}} = K - \widetilde{K}$. For $\varepsilon \in (0, 1/40]$, define

$$\begin{aligned} \mathsf{P} &= \left(\begin{array}{c|c} 0.5\,\mathbf{1}_{\widetilde{K}\times\widetilde{K}} & 0.9\,\mathbf{1}_{\widetilde{K}\times K_{\mathrm{b}}} \\ \hline 0.1\,\mathbf{1}_{K_{\mathrm{b}}\times\widetilde{K}} & 0.5\,\mathbf{1}_{K_{\mathrm{b}}\times K_{\mathrm{b}}} \end{array}\right), \\ \mathsf{P}^{(m)} &= \left(\begin{array}{c|c} 0.5\,\mathbf{1}_{\widetilde{K}\times\widetilde{K}} & 0.9\,\mathbf{1}_{\widetilde{K}\times K_{\mathrm{b}}} + 2\varepsilon e_m \mathbf{1}_{K_{\mathrm{b}}}^{\top} \\ \hline 0.1\,\mathbf{1}_{K_{\mathrm{b}}\times\widetilde{K}} - 2\varepsilon \mathbf{1}_{K_{\mathrm{b}}} e_m^{\top} & 0.5\,\mathbf{1}_{K_{\mathrm{b}}\times K_{\mathrm{b}}} \end{array}\right), \qquad m \in \mathcal{K}_g\,. \end{aligned} \tag{86}$$

Here e_m is the m-th standard basis vector in $\mathbb{R}^{\widetilde{K}}$. We also define the matrix at the penalty states,

$$\mathsf{U} = 0.5\,\mathbf{1}_{K\times K}\,. \tag{87}$$

For a preference matrix Q, we write the shifted Borda score and the corresponding loss as

$$b_{\mathsf{Q}}(i) = \frac{1}{K}\sum_{r=1}^{K}(\mathsf{Q}[i,r] - 1), \qquad \ell_{\mathsf{Q}}(i,k) = -\frac{1}{2}\big(b_{\mathsf{Q}}(i) + b_{\mathsf{Q}}(k)\big). \tag{88}$$

Under P, all arms in $\mathcal{K}_g$ have the same Borda score β_g and all arms in $\mathcal{K}_b$ have the same Borda score β_b, where

$$\beta_g = \frac{0.5\widetilde{K} + 0.9K_{\mathrm{b}}}{K} - 1\,, \qquad \beta_b = \frac{0.1\widetilde{K} + 0.5K_{\mathrm{b}}}{K} - 1 = \beta_g - \frac{2}{5}\,.$$

Hence

$$b_{\mathsf{P}}(i) - b_{\mathsf{P}}(k) = \frac{2}{5} \qquad \text{for every } i \in \mathcal{K}_g \text{ and } k \in \mathcal{K}_b\,. \tag{89}$$

Under $\mathsf{P}^{(m)}$, the Borda scores are

$$b_{\mathsf{P}^{(m)}}(i) = \begin{cases} \beta_g + \Delta & \text{if } i = m\,, \\ \beta_g & \text{if } i \in \mathcal{K}_g \setminus \{m\}\,, \\ \beta_b - 2\varepsilon/K & \text{if } i \in \mathcal{K}_b\,, \end{cases} \qquad \Delta = \frac{2\varepsilon K_{\mathrm{b}}}{K} \ge \varepsilon\,.$$

Consequently, (m,m) is the unique optimal pair under $\mathsf{P}^{(m)}$, and for every other pair $a = (i,k)$,

$$\ell_{\mathsf{P}^{(m)}}(a) - \ell_{\mathsf{P}^{(m)}}(m,m) = \frac{1}{2}(b_{\mathsf{P}^{(m)}}(m) - b_{\mathsf{P}^{(m)}}(i) + b_{\mathsf{P}^{(m)}}(m) - b_{\mathsf{P}^{(m)}}(k)) \ge \frac{\varepsilon}{2}\,. \tag{90}$$

At the penalty states, the matrix is U, so that $b_{\mathsf{U}}(i) = -1/2$ and $\ell_{\mathsf{U}}(a) = 1/2$ for every i and a. Since

$$\frac{1}{2} - \ell_{\mathsf{P}}(1,1) = \frac{1}{2} + \beta_g = \frac{2K_{\mathrm{b}}}{5K} \ge \frac{1}{5} \tag{91}$$

and $\ell_{\mathsf{P}^{(m)}}(m,m) \le \ell_{\mathsf{P}}(1,1)$, the loss $1/2$ at a penalty state exceeds the minimal loss under P and under $\mathsf{P}^{(m)}$ by at least $1/5$.

The preference matrices are those of Saha et al. (2021, Lemma 14), with perturbation size 2ε and group sizes $\widetilde{K}$ and K_{b}. In our construction, every pair to which P and $\mathsf{P}^{(m)}$ assign different probabilities belongs to $\mathcal{A}_b$ and therefore changes the subsequent state trajectory.

Random choice of the preference matrices. We choose the matrices at random before the interaction begins. Set $\mathsf{P}^{(0)} = \mathsf{P}$, let $T' = T/M$ be the expected number of episodes in which a given chain is visited, and let

$$\varepsilon = c_0 \min\left\{1, \left(\frac{KR}{T'}\right)^{1/3}\right\}. \tag{92}$$

The numerical constant $c_0 \in (0, 1/40]$ will be fixed at the end of the proof.

For every $x \in \mathcal{X}$, let $J_x \in \{0\} \cup \mathcal{K}_g$ specify the matrix at x. Draw the indices independently, once before the first episode, according to

$$\Pr[J_x = 0] = \frac{1}{2}, \qquad \Pr[J_x = m] = \frac{1}{2\widetilde{K}} \quad \text{for every } m \in \mathcal{K}_g. \tag{93}$$

Let $J = (J_x)_{x\in\mathcal{X}}$ and let $\mathcal{D}_{\mathrm{pref}}$ denote its product distribution. The indices are not revealed to the learner.

For a realization $j = (j_x)_{x\in\mathcal{X}}$, define, for every episode t, the preference function

$$\mathsf{P}_t^{(j)}(s,i,k) = \begin{cases} \mathsf{P}^{(j_x)}[i,k] & \text{if } s = x \in \mathcal{X}, \\ \mathsf{P}[i,k] & \text{if } s = s_0 \text{ or } s = x_{c,h} \text{ with } h > L, \\ \mathsf{U}[i,k] & \text{otherwise}. \end{cases}$$

Thus the matrices at all penalty states and additional unreachable states are U as shown in Figure 1. Let $\mathcal{M}_{\mathsf{pref}}(j)$ be the resulting instance and $\ell^{(j)}$ its Borda loss function. For fixed j, the preference matrices and $\ell^{(j)}$ are fixed over all episodes. The regret is measured with $\ell^{(j)}$, while the learner observes only the Bernoulli outcomes of its comparisons. We write $\mathbb{E}^{(j)}$ for expectation over the learner's randomization, transitions, and preference outcomes under this fixed instance.

Optimal policy. The comparator in the regret is the best fixed policy for the instance $\mathcal{M}_{\mathsf{pref}}(j)$, which we now identify. Define $a_0^\star = (1,1)$ and $a_m^\star = (m,m)$ for $m \in \mathcal{K}_g$, and let $\pi_j^\star$ be the policy that selects $a_{j_x}^\star$ at each $x \in \mathcal{X}$, $(1,1)$ at s_0 and at the states $x_{c,h}$ with $h > L$, and an arbitrary pair elsewhere.

At $x \in \mathcal{X}$, the pair $a_{j_x}^\star$ minimizes the current loss and leads to $x_{c,h+1}$. An action in $\mathcal{A}_b$ has no smaller current loss and leads instead to penalty states, whose losses exceed the corresponding losses under $\pi_j^\star$. At the remaining states $x_{c,h}$, the transition does not depend on the action, and minimizing the current loss is therefore sufficient. Backward induction gives $\pi_j^\star \in \arg\min_{\pi\in\Pi} V^\pi(s_0; \ell^{(j)})$. Since $\ell^{(j)}$ is fixed over episodes, this policy attains the minimum in the regret definition for each fixed j, before averaging over J.

Loss difference at one state. We now turn to the analysis of the regret on this instance. We first express the loss difference at a layer using the learner's visit probabilities. Let $q_t = q^{\pi_t}$ be the occupancy measure of the policy selected in episode t under the common transition kernel, and define

$$N_T(y,a) = \sum_{t=1}^{T} q_t(y,a) \qquad \text{for } y \in \mathcal{S} \setminus \{s_H\} \text{ and } a \in \mathcal{A}. \tag{94}$$

Since the policies depend on past observations, $N_T(y,a)$ may be random. Fix $x = x_{c,h} \in \mathcal{X}$ and write $z = z_{c,h}$. Every policy enters chain c with probability $1/M$ and then reaches exactly one of x and z at layer h as shown in Figure 1, and thus $q_t(x) + q_t(z) = 1/M$ and

$$\sum_{a\in\mathcal{A}} (N_T(x,a) + N_T(z,a)) = \frac{T}{M} = T', \qquad 0 \le N_T(x,a) \le T'. \tag{95}$$

Recall that $\mathbb{I}_t(y,a)$ indicates whether (y,a) is visited in episode t. Given the observations before episode t, which determine π_t, the pair (y,a) is visited with probability $q_t(y,a)$, that is,

$$\mathbb{E}^{(j)}[\mathbb{I}_t(y,a) \mid \pi_t] = q_t(y,a).$$

Consequently,

$$\mathbb{E}^{(j)}\left[\sum_{t=1}^{T}\mathbb{I}_t(y,a)\right]=\mathbb{E}^{(j)}\left[\sum_{t=1}^{T}q_t(y,a)\right]=\mathbb{E}^{(j)}[N_T(y,a)]\,. \tag{96}$$

Fix $j_{-x}=(j_{x'})_{x'\in\mathcal{X}\setminus\{x\}}$. For $u\in\{0\}\cup\mathcal{K}_g$, let $\mathcal{M}_{\mathsf{pref}}(u,j_{-x})$ be the instance in which the matrix at x is $\mathsf{P}^{(u)}$ and the other indices are j_{-x}. Thus the instances considered below differ only in the matrix at x. Let $\mathbb{P}_{u,j_{-x}}$ be the joint distribution of the learner's randomization, the selected policies, the state-action trajectories, and the observed preference outcomes, and let $\mathbb{E}^{(u,j_{-x})}$ be the corresponding expectation. Since x and j_{-x} remain fixed in this step and the next, we drop them from the notation and write

$$\mathbb{P}_0=\mathbb{P}_{0,j_{-x}}\,,\quad \mathbb{P}_m=\mathbb{P}_{m,j_{-x}}\,,\quad \mathbb{E}_0=\mathbb{E}^{(0,j_{-x})}\,,\quad \mathbb{E}_m=\mathbb{E}^{(m,j_{-x})}\,. \tag{97}$$

Both lower bounds below are expressed through the expected number of comparisons in $\mathcal{A}_b$ at x under the baseline matrix P,

$$B_x(j_{-x})=\sum_{a\in\mathcal{A}_b}\mathbb{E}_0[N_T(x,a)]\,. \tag{98}$$

Let $R_{x,u}(j_{-x})$ be the expected loss difference between the learner and the comparator $\pi_j^\star$ with $j_x=u$ at layer h of chain c, that is, at the states x and z, under the instance $\mathcal{M}_{\mathsf{pref}}(u,j_{-x})$. The comparator visits x with probability $1/M$ and never visits z, so that its expected loss at this layer of the chain is $T'\ell_{\mathsf{P}^{(u)}}(a_u^\star)$, and hence

$$\begin{aligned}
R_{x,u}(j_{-x})&=\mathbb{E}^{(u,j_{-x})}\left[\sum_{a\in\mathcal{A}}N_T(x,a)\ell_{\mathsf{P}^{(u)}}(a)+\frac{1}{2}\sum_{a\in\mathcal{A}}N_T(z,a)\right]-T'\ell_{\mathsf{P}^{(u)}}(a_u^\star)\\
&=\mathbb{E}^{(u,j_{-x})}\left[\sum_{a\in\mathcal{A}}N_T(x,a)(\ell_{\mathsf{P}^{(u)}}(a)-\ell_{\mathsf{P}^{(u)}}(a_u^\star))\right]\\
&\quad+\mathbb{E}^{(u,j_{-x})}\left[\sum_{a\in\mathcal{A}}N_T(z,a)\left(\frac{1}{2}-\ell_{\mathsf{P}^{(u)}}(a_u^\star)\right)\right],
\end{aligned}$$

where the second equality follows from (95). Both gaps in the last expression are nonnegative. When the matrix at x is $\mathsf{P}^{(m)}$, we have

$$\begin{aligned}
R_{x,m}(j_{-x})&\ge\frac{\varepsilon}{2}\mathbb{E}_m\left[\sum_{a\neq(m,m)}N_T(x,a)+\sum_{a\in\mathcal{A}}N_T(z,a)\right]\\
&=\frac{\varepsilon}{2}(T'-\mathbb{E}_m[N_T(x,(m,m))])\,,
\end{aligned} \tag{99}$$

where the inequality follows from the gap of at least $\varepsilon/2$ for every pair $a\neq(m,m)$ in (90) and from the gap of at least $1/5\ge\varepsilon/2$ at the penalty states shown after (91), and the equality from (95).

Change of measure at one state. This step shows that, on average over m, the learner cannot select (m,m) much more often under $\mathsf{P}^{(m)}$ than under P unless $B_x(j_{-x})$ is large. We compare the two expectations of $N_T(x,(m,m))$. Since this variable takes values in $[0,T']$, the variational characterization of total variation distance and Pinsker's inequality give

$$\begin{aligned}
\frac{1}{T'}\mathbb{E}_m[N_T(x,(m,m))]&\le\frac{1}{T'}\mathbb{E}_0[N_T(x,(m,m))]+\mathrm{TV}(\mathbb{P}_0,\mathbb{P}_m)\\
&\le\frac{1}{T'}\mathbb{E}_0[N_T(x,(m,m))]+\sqrt{\frac{1}{2}\mathrm{D}(\mathbb{P}_0,\mathbb{P}_m)}\,,
\end{aligned} \tag{100}$$

where $\mathrm{TV}(\cdot,\cdot)$ denotes the total variation distance.

Under $\mathbb{P}_0$ and $\mathbb{P}_m$, the learner's decision rules and the transition kernel are the same, and the conditional distribution of a preference outcome differs only at x for the pairs (m,i) and (i,m) with

$i \in \mathcal{K}_b$. By the chain rule for the KL divergence and (96),

$$\begin{aligned}\mathrm{D}(\mathbb{P}_0, \mathbb{P}_m) &= \sum_{i \in \mathcal{K}_b} \mathbb{E}_0[N_T(x,(m,i))]\mathrm{kl}(0.9, 0.9+2\varepsilon) \\ &\quad + \sum_{i \in \mathcal{K}_b} \mathbb{E}_0[N_T(x,(i,m))]\mathrm{kl}(0.1, 0.1-2\varepsilon) \\ &\le C_{\mathrm{kl}}\varepsilon^2 \sum_{i \in \mathcal{K}_b} \mathbb{E}_0[N_T(x,(m,i)) + N_T(x,(i,m))]\,, \end{aligned} \tag{101}$$

where $C_{\mathrm{kl}} > 0$ is a universal constant, and the last inequality follows from $\mathrm{kl}(p,q) \le (p-q)^2/(q(1-q))$ and $\varepsilon \le 1/40$, which give $q \in [0.05, 0.95]$ and hence $q(1-q) \ge 0.05 \cdot 0.95$ for both divergences, so that $C_{\mathrm{kl}} = 85$ suffices.

Since $\widetilde{K} = \lfloor K/2 \rfloor \ge K/3$ for $K \ge 5$, (95) implies

$$\frac{1}{\widetilde{K}} \sum_{m \in \mathcal{K}_g} \mathbb{E}_0[N_T(x,(m,m))] \le \frac{T'}{\widetilde{K}} \le \frac{3T'}{K}\,. \tag{102}$$

Moreover, summing (101) over m and using (98) yields

$$\sum_{m \in \mathcal{K}_g} \mathrm{D}(\mathbb{P}_0, \mathbb{P}_m) \le C_{\mathrm{kl}}\varepsilon^2 B_x(j_{-x})\,.$$

Jensen's inequality then gives

$$\begin{aligned}\frac{1}{\widetilde{K}} \sum_{m \in \mathcal{K}_g} \sqrt{\frac{1}{2}\mathrm{D}(\mathbb{P}_0, \mathbb{P}_m)} &\le \sqrt{\frac{1}{2\widetilde{K}} \sum_{m \in \mathcal{K}_g} \mathrm{D}(\mathbb{P}_0, \mathbb{P}_m)} \\ &\le \varepsilon \sqrt{\frac{3C_{\mathrm{kl}} B_x(j_{-x})}{2K}}\,. \end{aligned} \tag{103}$$

Averaging (100) over m and using (102) and (103), we obtain

$$\frac{1}{\widetilde{K}} \sum_{m \in \mathcal{K}_g} \mathbb{E}_m[N_T(x,(m,m))] \le \frac{3T'}{K} + T'\varepsilon \sqrt{\frac{3C_{\mathrm{kl}} B_x(j_{-x})}{2K}}\,. \tag{104}$$

Substituting this inequality into (99), and using $1 - 3/K \ge 2/5$ for $K \ge 5$, gives

$$\begin{aligned}\frac{1}{2\widetilde{K}} \sum_{m \in \mathcal{K}_g} R_{x,m}(j_{-x}) &\ge \frac{\varepsilon}{4}\left(T' - \frac{1}{\widetilde{K}} \sum_{m \in \mathcal{K}_g} \mathbb{E}_m[N_T(x,(m,m))]\right) \\ &\ge \frac{\varepsilon T'}{10} - C_1 \varepsilon^2 T' \sqrt{\frac{B_x(j_{-x})}{K}}\,, \end{aligned} \tag{105}$$

where $C_1 = \sqrt{3C_{\mathrm{kl}}/2}/4$. Thus the average regret at x under the perturbed matrices is of order $\varepsilon T'$ unless $B_x(j_{-x})$ is of order K/ε^2.

Two lower bounds on the average regret. We bound the regret averaged over the independent choices of the preference matrices. Let

$$\mathfrak{R}_T = \mathbb{E}_{J \sim \mathcal{D}_{\mathrm{pref}}}\left[\mathbb{E}^{(J)}[\mathsf{Reg}_T]\right].$$

We first use the losses incurred in the first L intermediate layers. For a fixed realization j, the expected regret equals the total expected loss difference from $\pi_j^\star$. The contribution at layer h in chain c is $R_{x_{c,h}, j_{x_{c,h}}}(j_{-x_{c,h}})$, while the contributions at s_0 and in the last R intermediate layers are nonnegative, so that

$$\mathbb{E}^{(j)}[\mathsf{Reg}_T] \ge \sum_{x \in \mathcal{X}} R_{x, j_x}(j_{-x})\,. \tag{106}$$

For each fixed j_{-x}, independence of the indices gives

$$\mathbb{E}_J[R_{x,J_x}(J_{-x}) \mid J_{-x} = j_{-x}] = \frac{1}{2}R_{x,0}(j_{-x}) + \frac{1}{2\widetilde{K}}\sum_{m\in\mathcal{K}_g} R_{x,m}(j_{-x}).$$

Taking expectations in (106), dropping the nonnegative terms with $J_x = 0$, and using (105) yields

$$\begin{aligned}\mathfrak{R}_T &\geq \sum_{x\in\mathcal{X}} \mathbb{E}_{J_{-x}}\left[\frac{1}{2\widetilde{K}}\sum_{m\in\mathcal{K}_g} R_{x,m}(J_{-x})\right] \\ &\geq \frac{\varepsilon T' LM}{10} - C_1\varepsilon^2 T' \sum_{x\in\mathcal{X}} \mathbb{E}_{J_{-x}}\left[\sqrt{\frac{B_x(J_{-x})}{K}}\right].\end{aligned}$$

We next use the losses in the last R intermediate layers. Fix j. If the learner selects a pair in $\mathcal{A}_b$ at some $x \in \mathcal{X}$ in an episode, which happens at most once per episode since the trajectory then stays at the penalty states, the learner incurs the loss $1/2$ at each of the last R layers, whereas the comparator incurs $\ell_{\mathsf{P}}(1,1)$, and the difference is at least $R/5$ by (91). All other contributions to $V^{\pi_t}(s_0;\ell^{(j)}) - V^{\pi_j^\star}(s_0;\ell^{(j)})$ are nonnegative, since the comparator minimizes the loss at every state $x_{c,h}$ and the losses at the penalty states are larger. Hence, counting only the states with $j_x = 0$ so that the counts match $B_x(j_{-x})$ after averaging over J,

$$V^{\pi_t}(s_0;\ell^{(j)}) - V^{\pi_j^\star}(s_0;\ell^{(j)}) \geq \frac{R}{5}\sum_{x\in\mathcal{X}} \mathbb{1}[j_x = 0] \sum_{a\in\mathcal{A}_b} q_t(x,a)\,.$$

Summing over t and taking expectations gives

$$\mathbb{E}^{(j)}[\mathsf{Reg}_T] \geq \frac{R}{5}\sum_{x\in\mathcal{X}} \mathbb{1}[j_x = 0] \sum_{a\in\mathcal{A}_b} \mathbb{E}^{(j)}[N_T(x,a)]\,.$$

Averaging this inequality over J and using (93) and (98) gives

$$\begin{aligned}\mathfrak{R}_T &\geq \frac{R}{5}\sum_{x\in\mathcal{X}} \mathbb{E}_J\left[\mathbb{1}[J_x = 0] \sum_{a\in\mathcal{A}_b} \mathbb{E}^{(J)}[N_T(x,a)]\right] \\ &= \frac{R}{10}\sum_{x\in\mathcal{X}} \mathbb{E}_{J_{-x}}[B_x(J_{-x})]\,.\end{aligned}$$

Combining the two bounds. The first bound is useful when the counts B_x are small, and the second when they are large. Adding the two inequalities gives

$$2\mathfrak{R}_T \geq \frac{\varepsilon T' LM}{10} + \sum_{x\in\mathcal{X}} \mathbb{E}_{J_{-x}}\left[\frac{R}{10}B_x(J_{-x}) - C_1\varepsilon^2 T'\sqrt{\frac{B_x(J_{-x})}{K}}\right]. \tag{107}$$

For every $B \geq 0$,

$$\begin{aligned}\frac{R}{10}B - C_1\varepsilon^2 T'\sqrt{\frac{B}{K}} &= \frac{R}{10}\left(\sqrt{B} - \frac{5C_1\varepsilon^2 T'}{R\sqrt{K}}\right)^2 - \frac{5C_1^2\varepsilon^4 (T')^2}{2KR} \\ &\geq -\frac{5C_1^2\varepsilon^4 (T')^2}{2KR}\,.\end{aligned} \tag{108}$$

Since $|\mathcal{X}| = LM$, (107) and (108) imply

$$2\mathfrak{R}_T \geq \varepsilon T' LM\left(\frac{1}{10} - \frac{5C_1^2}{2}\frac{\varepsilon^3 T'}{KR}\right).$$

The choice in (92) gives $\varepsilon^3 T'/(KR) \leq c_0^3$. Choose $c_0 \leq 1/40$ sufficiently small that $(5C_1^2/2)c_0^3 \leq 1/20$. Then

$$\mathfrak{R}_T \geq \frac{\varepsilon T' LM}{40}\,. \tag{109}$$

By (92) and $MT' = T$,

$$\varepsilon T' LM = c_0 LM \min\Big\{T', (KR)^{1/3}(T')^{2/3}\Big\}$$
$$= c_0 \min\Big\{LT, L(KRM)^{1/3}T^{2/3}\Big\}. \tag{110}$$

For $H \geq 3$, $L \geq (H-1)/3 \geq 2H/9$ and $R \geq (H-1)/2 \geq H/3$. Also, since $S \geq 2H$ and $\lfloor v \rfloor \geq v/2$ for $v \geq 1$, we have

$$M \geq \frac{S-2}{4(H-1)} \geq \frac{S}{8H}.$$

It follows that $LT \geq (2/9)HT$ and

$$L(KRM)^{1/3} \geq \frac{2H}{9}\left(K \cdot \frac{H}{3} \cdot \frac{S}{8H}\right)^{1/3} \geq c_2 H(SK)^{1/3}, \tag{111}$$

for a universal constant $c_2 > 0$. Combining (109)–(111) gives

$$\mathfrak{R}_T \geq c_3 \min\Big\{HT, (H^3SK)^{1/3}T^{2/3}\Big\},$$

where $c_3 > 0$ is a universal constant. Finally, $\mathfrak{R}_T$ is the average of $\mathbb{E}^{(j)}[\mathsf{Reg}_T]$ over fixed realizations j. Hence some realization $j^\star$ satisfies the same bound, and $\mathcal{M}_{\mathsf{pref}}(j^\star)$ is the required state-wise PbMDP, which completes the proof. □

F Pair-independent transitions

To isolate the effect of transitions that depend on the selected pair, we consider here state-wise PbMDPs whose transition kernel does not depend on it, that is, $P(s' \mid s, a) = P(s' \mid s)$ for all (s, a, s'). The state-visit probabilities are then the same for every policy, and the learner faces a separate dueling bandit problem at each state. We show that the minimax regret under this assumption is $\widetilde{\Theta}(\min\{HT, (H^2SK)^{1/3}T^{2/3}\})$ for $H \geq 2$ and $K \geq 5$ (Proposition 20 and Theorem 21). Under the assumptions of Theorem 3, the $T^{2/3}$ term is smaller by a factor of $H^{1/3}$ than that of the minimax regret $\widetilde{\Theta}(\min\{HT, (H^3SK)^{1/3}T^{2/3}\})$ of general state-wise PbMDPs (Theorems 1 and 3), and this factor is thus the cost of transitions that depend on the selected pair.

F.1 Upper bound

The algorithm runs a separate copy of Dueling-EXP3 (Saha et al., 2021, Theorem 2) at each non-terminal state. The copy at s uses only the comparisons observed at s and is updated only when s is visited. Since a state is visited at most once per episode, the distributions over pairs at all states are determined before each episode, and the algorithm thus selects a policy $\pi_t \in \Pi$ at the start of episode t. The algorithm does not use the transition kernel. Since the parameters of Dueling-EXP3 depend on the number of rounds, each copy applies the doubling trick to its own number of visits. In a phase shorter than $2K \log K$, the copy makes uniform comparisons, and in a longer phase it runs Dueling-EXP3 with the parameters for the phase length. After n visits, the expected regret of the copy against any fixed arm is $O((K \log K)^{1/3} n^{2/3})$.

Proposition 20. *Suppose that $P(s' \mid s, a) = P(s' \mid s)$ for all (s, a, s'). The algorithm above satisfies, for every $T \geq 1$,*

$$\mathbb{E}[\mathsf{Reg}_T] = O\Big(\min\Big\{HT, (H^2SK \log K)^{1/3}T^{2/3}\Big\}\Big).$$

Proof. Let $n_s = \sum_{t=1}^{T} \mathbb{I}_t(s)$ be the number of visits to s, and let $q(s) = q^\pi(s)$ be the visit probability of s, which does not depend on π. Since the visit probabilities do not depend on the policy and the loss of a pair is the average of the losses of its two arms, the comparator can be taken to play (i_s^*, i_s^*) at each state s, where $i_s^* \in \arg\min_{i \in [K]} \sum_{t=1}^{T} \ell_t(s, (i, i))$. Under the oblivious-adversary assumption, the episodes in which s is visited do not depend on the learner, and conditionally on these episodes

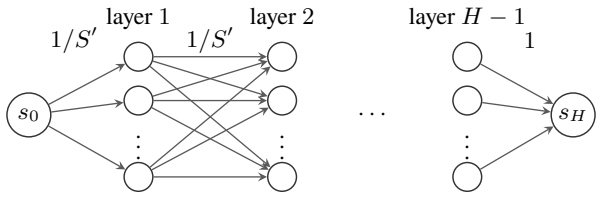


Figure 2: Pair-independent instance with $S' = \lfloor (S-2)/(H-1) \rfloor$ reachable states per intermediate layer. Transitions to the next intermediate layer are uniform for every pair, and all states in layer $H-1$ lead to s_H.

the copy at s faces a fixed sequence of preference matrices. Since the shift in (1) cancels in loss differences, the regret bound of the copy applies to the losses ℓ_t. Therefore,

$$\begin{aligned}
\mathbb{E}[\mathsf{Reg}_T] &= \sum_s \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{I}_t(s)\big(\ell_t(s, a_{t,h(s)}) - \ell_t(s, (i_s^*, i_s^*))\big)\right] \\
&\le C(K \log K)^{1/3} \sum_s \mathbb{E}\left[n_s^{2/3}\right] \\
&\le C(K \log K)^{1/3} T^{2/3} \sum_s q(s)^{2/3} \\
&\le C(H^2 SK \log K)^{1/3} T^{2/3},
\end{aligned}$$

where C is a universal constant, the equality follows from $\mathbb{E}_t[\mathbb{I}_t(s)\ell_t(s, a_{t,h(s)})] = q(s) \sum_a \pi_t(a \mid s)\ell_t(s,a)$, the first inequality from the regret bound of the copy at s conditionally on its visits, the second inequality from Jensen's inequality and $\mathbb{E}[n_s] = Tq(s)$, and the third inequality from $\sum_s q(s)^{2/3} \le S^{1/3}(\sum_s q(s))^{2/3} = (SH^2)^{1/3}$. Together with $\mathsf{Reg}_T \le HT$, this completes the proof. □

Note that under known transitions, Algorithm 1 with adjusted parameters attains a bound of the same order. Since the visit probabilities do not depend on the policy, the policy π_ρ of Lemma 7 satisfies

$$V^{\pi_\rho}(s_0; \ell) - V^{\pi}(s_0; \ell) \le \rho H.$$

Using this bound in place of Lemma 7 in the proof of Theorem 1 gives

$$\mathbb{E}[\mathsf{Reg}_T] \le H \log(SA)/\eta + \eta SKT/(2\rho) + \rho HT.$$

Choosing $\eta = (H \log(SA)/(T\sqrt{HSK}))^{2/3}$ and $\rho = \min\{1/2, \sqrt{\eta SK/H}\}$ gives $\mathbb{E}[\mathsf{Reg}_T] = O(\min\{HT, (H^2 SK \log(SK))^{1/3} T^{2/3}\})$ for every $T \ge 1$.

F.2 Lower bound

Theorem 21. *Suppose that $H \ge 2$ and $K \ge 5$. For any learning algorithm and any $T \ge 1$, there exists an episodic state-wise PbMDP with Borda scores and a known transition kernel with $P(s' \mid s, a) = P(s' \mid s)$ for all (s, a, s'), such that*

$$\mathbb{E}[\mathsf{Reg}_T] = \Omega\Big(\min\Big\{HT, (H^2 SK)^{1/3} T^{2/3}\Big\}\Big).$$

The instance is the layered MDP in Figure 2, in which each intermediate layer has $S' = \lfloor (S-2)/(H-1) \rfloor$ reachable states and every policy visits each of them with probability $1/S'$ in every episode. At each of these states, the preference matrix is P or $\mathsf{P}^{(m)}$ of (86), drawn at random independently across states. The proof is a simpler version of the proof of Theorem 3. Since the learner's choices at a state may depend on the observations at the other states, the change of measure is again applied one state at a time, with the matrices at the other states fixed. The difference is that a pair in $\mathcal{A}_b$ now incurs a loss gap of at least $1/5$ at the state itself under P, instead of $R/5$ at the penalty states, so that the two bounds obtained under P and $\mathsf{P}^{(m)}$ concern the regret at the same state and are combined state by state.

Proof. Fix a learning algorithm. We use the notation of the proof of Theorem 3, namely $\widetilde{K} = \lfloor K/2 \rfloor$, $\mathcal{K}_g = [\widetilde{K}]$, $\mathcal{K}_b = [K] \setminus \mathcal{K}_g$, $\mathcal{A}_b = \mathcal{A} \setminus (\mathcal{K}_g \times \mathcal{K}_g)$, the matrices P, $\mathsf{P}^{(m)}$, and U of (86) and (87), and the Borda scores and losses b_{Q} and ℓ_{Q} of (88).

Instance. Let $S' = \lfloor (S-2)/(H-1) \rfloor$. Since every intermediate layer is nonempty, $S \geq H+1$ and hence $S' \geq 1$. Each intermediate layer $h \in [H-1]$ contains the states $x_{h,1}, \ldots, x_{h,S'}$ (Figure 2), and if $2 + (H-1)S' < S$, we add unreachable states to the intermediate layers. For every pair a, the transition from s_0 is uniform over $x_{1,1}, \ldots, x_{1,S'}$, the transition from $x_{h,i}$ with $h \leq H-2$ is uniform over $x_{h+1,1}, \ldots, x_{h+1,S'}$, the transition from an unreachable state is uniform over the next layer, and every state in layer $H-1$ leads to s_H. Thus the transition kernel does not depend on the selected pair, and every policy visits each of the states in

$$\mathcal{X} = \{x_{h,i} : h \in [H-1],\ i \in [S']\}$$

with probability $1/S'$ in every episode.

Let $T' = T/S'$ be the expected number of episodes in which a given state in $\mathcal{X}$ is visited, and set

$$\varepsilon = \frac{1}{40} \min\left\{1, \left(\frac{K}{T'}\right)^{1/3}\right\}. \tag{112}$$

For every $x \in \mathcal{X}$, draw an index $J_x \in \{0\} \cup \mathcal{K}_g$ according to (93), independently across states and before the first episode, and let $J = (J_x)_{x \in \mathcal{X}}$ and $\mathcal{D}_{\mathrm{pref}}$ be its product distribution. For a realization $j = (j_x)_{x \in \mathcal{X}}$, the preference function is fixed over episodes and given by $\mathsf{P}^{(j_x)}$ at $x \in \mathcal{X}$, by P at s_0, and by U at the unreachable states, where $\mathsf{P}^{(0)} = \mathsf{P}$. Let $\mathcal{M}_{\mathsf{pref}}(j)$ be the resulting instance, $\ell^{(j)}$ its Borda loss function, and $\mathbb{E}^{(j)}$ the expectation over the learner's randomization, transitions, and preference outcomes under this instance.

The comparator is the following policy. Define $a_0^\star = (1,1)$ and $a_m^\star = (m,m)$ for $m \in \mathcal{K}_g$, and let $\pi_j^\star$ be the policy that selects $a_{j_x}^\star$ at each $x \in \mathcal{X}$, $(1,1)$ at s_0, and an arbitrary pair elsewhere.

Since the state-visit probabilities are the same for every policy, $V^\pi(s_0; \ell^{(j)}) = \sum_s q^\pi(s) \sum_a \pi(a \mid s) \ell^{(j)}(s,a)$ is minimized by minimizing the loss at every state separately. Under P, every pair in $\mathcal{K}_g \times \mathcal{K}_g$ minimizes the loss by (89), and under $\mathsf{P}^{(m)}$, the pair (m,m) is the unique minimizer by (90). Hence $\pi_j^\star \in \arg\min_{\pi \in \Pi} V^\pi(s_0; \ell^{(j)})$, and since $\ell^{(j)}$ is fixed over episodes, $\pi_j^\star$ attains the minimum in the regret definition for each fixed j.

Regret at one state. Let $q_t = q^{\pi_t}$ be the occupancy measure of the policy selected in episode t, so that $q_t(x,a) = \pi_t(a \mid x)/S'$ for $x \in \mathcal{X}$, and define $N_T(x,a) = \sum_{t=1}^T q_t(x,a)$ as in (94). Then

$$\sum_{a \in \mathcal{A}} N_T(x,a) = T', \qquad 0 \leq N_T(x,a) \leq T', \tag{113}$$

and (96) holds for the same reason. Since π_t and $\pi_j^\star$ visit every state with the same probability, the loss difference in episode t is $V^{\pi_t}(s_0; \ell^{(j)}) - V^{\pi_j^\star}(s_0; \ell^{(j)}) = \sum_s q_t(s) \sum_a (\pi_t(a \mid s) - \pi_j^\star(a \mid s)) \ell^{(j)}(s,a)$. The term for s_0 is nonnegative, because $(1,1)$ minimizes ℓ_{P}, and the terms for the unreachable states vanish. Summing over episodes and taking expectations gives

$$\mathbb{E}^{(j)}[\mathsf{Reg}_T] \geq \sum_{x \in \mathcal{X}} R_x(j), \qquad R_x(j) = \mathbb{E}^{(j)}\left[\sum_{a \in \mathcal{A}} N_T(x,a)\big(\ell_{\mathsf{P}^{(j_x)}}(a) - \ell_{\mathsf{P}^{(j_x)}}(a_{j_x}^\star)\big)\right]. \tag{114}$$

Fix $x \in \mathcal{X}$ and a realization $j_{-x} = (j_{x'})_{x' \in \mathcal{X} \setminus \{x\}}$ of the indices at the other states. For $u \in \{0\} \cup \mathcal{K}_g$, let $\mathcal{M}_{\mathsf{pref}}(u, j_{-x})$ be the instance in which the matrix at x is $\mathsf{P}^{(u)}$ and the other indices are j_{-x}, and let $\mathbb{P}_{u,j_{-x}}$ and $\mathbb{E}^{(u,j_{-x})}$ be the distribution of the interaction and the expectation under this instance, defined as in the proof of Theorem 3. As in (97), we drop x and j_{-x} from the notation and write $\mathbb{P}_0$, $\mathbb{P}_m$, $\mathbb{E}_0$, and $\mathbb{E}_m$. The instances compared below differ only in the matrix at x.

Since the learner's choices at x may depend on the observations made at the other states, the comparison concerns the distribution of the entire interaction. Define $B_x(j_{-x}) = \sum_{a \in \mathcal{A}_b} \mathbb{E}_0[N_T(x,a)]$

as in (98) and

$$R_{x,u}(j_{-x}) = \mathbb{E}^{(u,j_{-x})}\left[\sum_{a\in\mathcal{A}} N_T(x,a)(\ell_{\mathsf{P}^{(u)}}(a) - \ell_{\mathsf{P}^{(u)}}(a_u^\star))\right],$$

so that $R_x(j) = R_{x,j_x}(j_{-x})$, and $R_{x,u}(j_{-x}) \geq 0$ for every u.

Under P, (89) gives $\ell_{\mathsf{P}}(a) - \ell_{\mathsf{P}}(1,1) = \frac{1}{5}(\mathbb{1}[a^{\mathrm{L}} \in \mathcal{K}_b] + \mathbb{1}[a^{\mathrm{R}} \in \mathcal{K}_b])$, which is at least $1/5$ for $a \in \mathcal{A}_b$. Hence

$$R_{x,0}(j_{-x}) \geq \frac{1}{5} B_x(j_{-x}). \tag{115}$$

Under $\mathsf{P}^{(m)}$, we have

$$R_{x,m}(j_{-x}) \geq \frac{\varepsilon}{2}\mathbb{E}_m\left[\sum_{a\neq(m,m)} N_T(x,a)\right] = \frac{\varepsilon}{2}(T' - \mathbb{E}_m[N_T(x,(m,m))]), \tag{116}$$

where the inequality follows from the gap of at least $\varepsilon/2$ for every pair $a \neq (m,m)$ in (90), and the equality from (113). The bound on $\frac{1}{\widetilde{K}}\sum_{m\in\mathcal{K}_g}\mathbb{E}_m[N_T(x,(m,m))]$ derived in the proof of Theorem 3, (100)–(104), uses only the following facts, with $T' = T/S'$ in place of T/M: (i) $\varepsilon \leq 1/40$; (ii) $N_T(x,(m,m)) \in [0,T']$; (iii) the two instances share the transition kernel and differ only in the outcome distributions at x for the pairs (m,i) and (i,m) with $i \in \mathcal{K}_b$; (iv) the identities (96) and (113) and the definition of $B_x(j_{-x})$; (v) $\widetilde{K} \geq K/3$.

This bound therefore applies to the present instance. Substituting it into (116) and using $1 - 3/K \geq 2/5$ for $K \geq 5$, as in (105), gives

$$\frac{1}{2\widetilde{K}}\sum_{m\in\mathcal{K}_g} R_{x,m}(j_{-x}) \geq \frac{\varepsilon T'}{10} - C_1\varepsilon^2 T'\sqrt{\frac{B_x(j_{-x})}{K}}, \tag{117}$$

where $C_1 = \sqrt{3C_{\mathrm{kl}}/2}/4$.

Lower bound at one state. Let

$$\overline{R}_x(j_{-x}) = \frac{1}{2}R_{x,0}(j_{-x}) + \frac{1}{2\widetilde{K}}\sum_{m\in\mathcal{K}_g} R_{x,m}(j_{-x})$$

be the regret at x averaged over J_x according to (93). If $B_x(j_{-x}) \geq K/(40C_{\mathrm{kl}}\varepsilon^2)$, then (115) gives

$$\overline{R}_x(j_{-x}) \geq \frac{1}{2}R_{x,0}(j_{-x}) \geq \frac{1}{10}B_x(j_{-x}) \geq \frac{K}{400C_{\mathrm{kl}}\varepsilon^2}.$$

Otherwise, $C_1\varepsilon^2 T'\sqrt{B_x(j_{-x})/K} < C_1\varepsilon T'/\sqrt{40C_{\mathrm{kl}}} \leq \varepsilon T'/20$, since $C_1/\sqrt{40C_{\mathrm{kl}}} = \sqrt{3/80}/4 < 1/20$, and (117) gives

$$\overline{R}_x(j_{-x}) \geq \frac{1}{2\widetilde{K}}\sum_{m\in\mathcal{K}_g} R_{x,m}(j_{-x}) \geq \frac{\varepsilon T'}{20}.$$

In both cases, $\overline{R}_x(j_{-x}) \geq c\min\{K/\varepsilon^2, \varepsilon T'\}$ with $c = \min\{1/(400C_{\mathrm{kl}}), 1/20\}$. By (112), if $T' \leq K$, then $\varepsilon = 1/40$ and the minimum equals $T'/40$. If $T' > K$, then $\varepsilon = (K/T')^{1/3}/40$ and the minimum equals $K^{1/3}(T')^{2/3}/40$. Hence, for every $x \in \mathcal{X}$ and every j_{-x},

$$\overline{R}_x(j_{-x}) \geq \frac{c}{40}\min\left\{T', K^{1/3}(T')^{2/3}\right\}. \tag{118}$$

Summing over the states. Since the indices are independent, conditioning on $J_{-x} = j_{-x}$ leaves J_x distributed as in (93), and hence $\mathbb{E}_J[R_x(J) \mid J_{-x} = j_{-x}] = \overline{R}_x(j_{-x})$. Taking expectations in (114) and using (118), $|\mathcal{X}| = (H-1)S'$, and $S'T' = T$, we obtain

$$\begin{aligned}
\mathbb{E}_{J\sim\mathcal{D}_{\mathrm{pref}}}\left[\mathbb{E}^{(J)}[\mathsf{Reg}_T]\right] &\geq \sum_{x\in\mathcal{X}} \mathbb{E}_{J_{-x}}\left[\overline{R}_x(J_{-x})\right] \\
&\geq \frac{c}{40}(H-1)S'\min\left\{T', K^{1/3}(T')^{2/3}\right\} \\
&= \frac{c}{40}(H-1)\min\left\{T, (S'K)^{1/3}T^{2/3}\right\}.
\end{aligned} \tag{119}$$

Finally, $H-1 \geq H/2$ for $H \geq 2$, and since $\lfloor v \rfloor \geq v/2$ for $v \geq 1$, we have

$$S' \geq (S-2)/(2(H-1)) \geq S/(4H),$$

where the last inequality is equivalent to $S(H+1) \geq 4H$ and follows from $S \geq H+1$. It follows that

$$(H-1)(S'K)^{1/3} \geq \frac{H}{2}(SK/(4H))^{1/3} \geq \frac{1}{4}(H^2SK)^{1/3},$$

and the right-hand side of (119) is at least $c' \min\{HT, (H^2SK)^{1/3}T^{2/3}\}$ for a universal constant $c' > 0$.

The left-hand side of (119) is an average of $\mathbb{E}^{(j)}[\mathsf{Reg}_T]$ over fixed realizations j. Hence some realization $j^\star$ satisfies the same bound, and $\mathcal{M}_{\mathsf{pref}}(j^\star)$ is the required state-wise PbMDP, which completes the proof. □

G Experiments

We evaluate global optimization and policy optimization, Algorithms 1 and 2, with known transitions on small instances of the lower-bound construction in Appendix E. The experiments address three questions: (i) whether the regret grows with T at the rate $T^{2/3}$ of Theorems 1 and 2, (ii) how the regret at a fixed T depends on S, H, and K, and (iii) how the running times of the two algorithms compare. The code, all settings including the random seeds, and the results of all runs are included in the supplementary material.

G.1 Setup

Instances. An instance is specified by H, M, K, and T, and has $S = 2 + 2M(H-1)$ states, including the M unreachable states $z_{c,1}$, and $A = K^2$ pairs. As in Appendix E, let $L = \lfloor (H-1)/2 \rfloor$ and $R = H - 1 - L$. The preference matrices are drawn from the distribution in (93) before each run and fixed during the run. For comparison with Theorem 3, we set ε in (86) to

$$\varepsilon_T = \min\left\{\frac{1}{40}, \left(\frac{8KRM}{6375T}\right)^{1/3}\right\},$$

which maximizes the explicit lower bound

$$B_\varepsilon(T) = \frac{\varepsilon LT}{20} - \frac{1275}{128}\frac{\varepsilon^4 LT^2}{KRM}$$

obtained from (107) and (108).

Evaluation. For each run, we compute the regret $\mathsf{Reg}_T = \sum_{t=1}^{T}(V^{\pi_t}(s_0;\ell) - V^{\pi^*}(s_0;\ell))$ of the executed policies exactly from the true losses, which are used only for this evaluation. We report the cumulative regret, the regret per episode Reg_T/T, or the regret per step $\mathsf{Reg}_T/(HT)$.

Parameters. We use the parameters of Theorem 1 and (65) in the proof of Theorem 2 with the exploration rate and the learning rate multiplied by constants c_e and c_η, respectively. For global optimization,

$$\eta_0 = \left(\frac{\log(SA)}{T\sqrt{SK}}\right)^{2/3}, \qquad \eta = c_\eta \eta_0, \qquad \rho = \min\left\{\frac{1}{2}, c_e\frac{\sqrt{\eta_0 SK}}{H}\right\},$$

and for policy optimization,

$$\gamma = \min\left\{\frac{1}{2}, c_e\left(\frac{24SK^3\log A}{H^2T}\right)^{1/3}\right\},$$

$$\eta = \min\left\{\frac{1}{9H^2}, c_\eta \min\left\{\frac{1}{9H^2}, \sqrt{\frac{H\log A}{T(81H^5 + 24H^2SK^2 + 24HSK^3/\gamma)}}\right\}\right\},$$

$$\delta = 2\eta(K/\gamma + H).$$

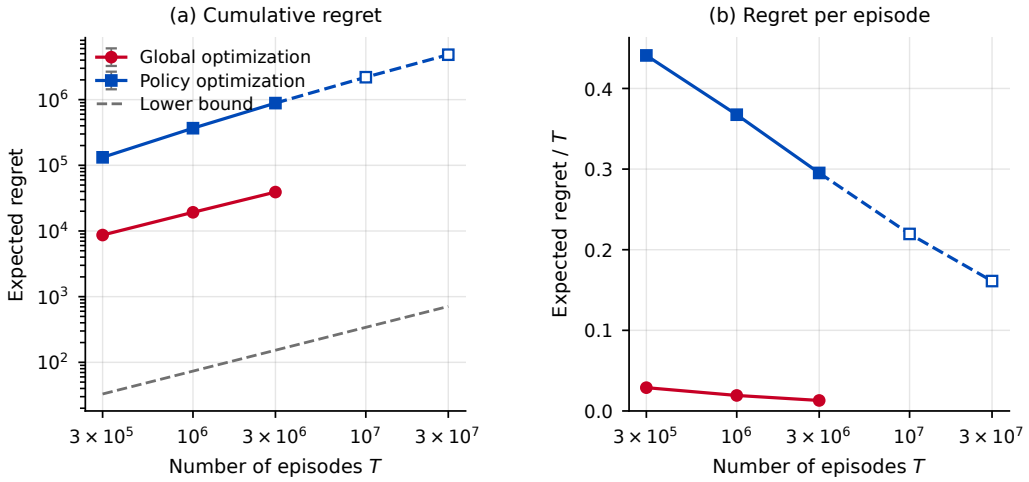


Figure 3: Cumulative regret and regret per episode versus the number of episodes T ($H = 3$, $S = 6$, and $K = 6$). Error bars show one standard error of the weighted mean over ten runs and are smaller than the markers. Open squares and dashed segments show the additional runs of policy optimization with $T \in \{10^7, 3 \times 10^7\}$. The gray curve is the lower bound $B_{\varepsilon_T}(T)$.

The estimators and the bonus are as in Algorithms 1 and 2. The multipliers were chosen in preliminary runs with other seeds so as to minimize the regret per step, and then fixed: $(c_e, c_\eta) = (1, 64)$ for global optimization, $(1, 4)$ for policy optimization in the experiment varying T, and $(0.14, 4)$ for policy optimization in the other experiments.

The exploration rates ρ and γ decrease with T but grow with S and K, and the truncation at $1/2$ is active unless T is large relative to the instance. In the truncated regime, the algorithms select a uniformly random pair with probability $1/2$ at every visited state, so that the regret is dominated by these comparisons and the bounds are vacuous. The experiment varying T therefore uses a small instance and large values of T, and in the other experiments the multiplier $c_e = 0.14$ keeps γ below $1/2$. The truncation is inactive in all reported runs.

G.2 Dependence on the number of episodes

We first examine whether the regret grows with T at the rate $T^{2/3}$ of Theorems 1 and 2. We set $H = 3$, $M = 1$, and $K = 6$, so that $S = 6$, and run both algorithms with $T \in \{3 \times 10^5, 10^6, 3 \times 10^6\}$ and policy optimization also with $T \in \{10^7, 3 \times 10^7\}$. The runs for different values of T are separate, each with its own ε_T and parameters, so that Figure 3 compares a family of instances indexed by T. These instances have the smallest H and S allowed in Theorem 3, and only the preference matrix at $x_{1,1}$ is random, with index $J \in \{0, 1, 2, 3\}$. Instead of drawing J, for each algorithm and T we use four learner seeds for $J = 0$ and two for each other value of J, and average the mean regrets for the four values of J with their probabilities $1/2, 1/6, 1/6, 1/6$ in (93). The gray curve in Figure 3 is the lower bound $B_{\varepsilon_T}(T)$, which applies to the regret averaged over J and the learner's randomness.

Figure 3(b) shows that the regret per episode decreases with T for both algorithms. In Figure 3(a), a least-squares fit of the logarithm of the cumulative regret against $\log T$ has slope 0.65 for global optimization over its three values of T, consistent with the exponent $2/3$ in Theorem 1, and slope 0.78 for policy optimization over its five values. For policy optimization, the slope between consecutive values of T decreases from 0.85 to 0.72 as T increases, approaching the exponent $2/3$ in Theorem 2.

G.3 Dependence on the problem size

We next examine how the regret at a fixed T depends on S, H, and K. We fix $T = 20{,}000$ and vary one of S, H, and K at a time.

- For S, we set $H = 5$ and $K = 6$ and vary $M \in \{1, 2, 4, 8\}$, so that $S \in \{10, 18, 34, 66\}$.

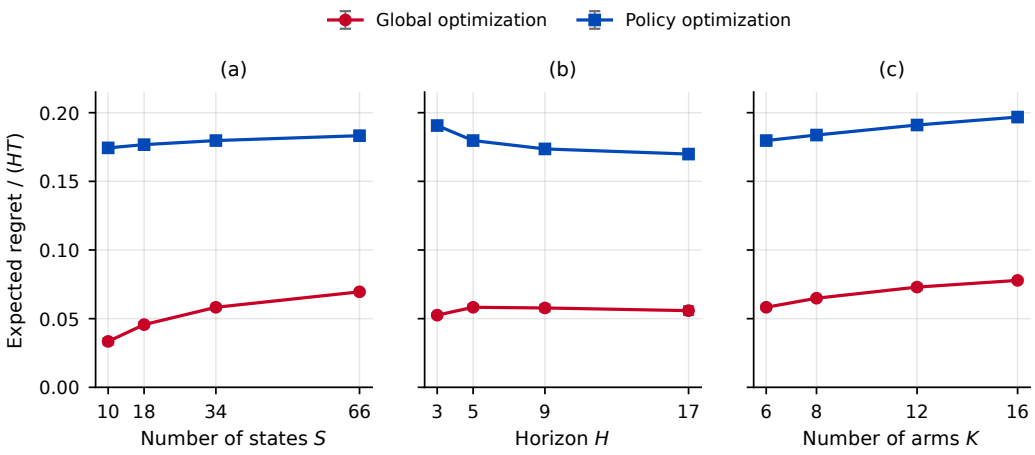


Figure 4: Regret per step for $T = 20{,}000$ as S, H, and K vary. Error bars show one standard error over ten runs and are smaller than the markers.

- For H, we set $S = 34$ and $K = 6$ and vary $(H, M) \in \{(3, 8), (5, 4), (9, 2), (17, 1)\}$, where M changes with H so that S is fixed, and hence the number $S - M$ of reachable states changes as well.
- For K, we set $H = 5$ and $M = 4$ and vary $K \in \{6, 8, 12, 16\}$.

Each configuration uses ten independent draws of the instance and of the learner's random numbers.

Figure 4 shows that the regret per step increases with S and with K for both algorithms, whereas its dependence on H is weak, with a slight decrease for policy optimization. These trends are observed at a fixed T and do not determine the dependence of the worst-case regret on S, H, and K.

G.4 Running time

We finally compare the running times of the two algorithms, since avoiding the convex program of global optimization is a motivation for policy optimization. We measure the time for $20{,}000$ episodes in five configurations (H, S, K), with five repetitions each. In each repetition, both algorithms use the same instance and the same random numbers, and the order in which they run alternates between repetitions. The measurements were taken on a single thread of an ARM64 macOS machine with Python 3.8.12, NumPy 1.24.4, SciPy 1.10.1, and Numba 0.58.1. They include initialization, policy updates, and trajectory simulation, but not the evaluation of the regret. Our implementation of global optimization eliminates the flow constraints of Ω_ρ using the chain structure of the instances and solves the resulting convex program with SLSQP, and our implementation of policy optimization performs the occupancy and Bellman computations with Numba. The experiment varying T uses faster implementations specialized to $H = 3$ and $M = 1$, which compute the same updates and are not timed.

Figure 5 shows that policy optimization is faster in every configuration for our implementations, and the median ratio of the running time of global optimization to that of policy optimization over paired repetitions ranges from 4.8 to 17.1.

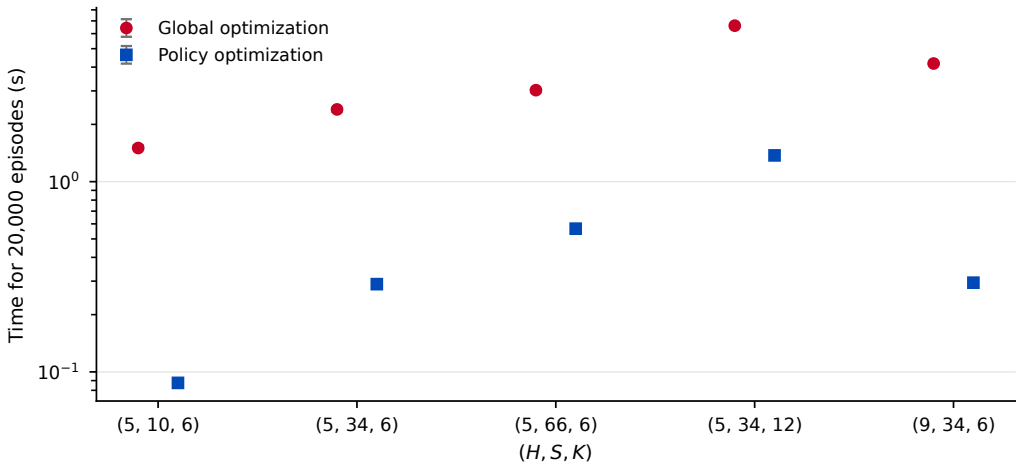


Figure 5: Running time for 20,000 episodes, excluding the evaluation of the regret, for five configurations (H, S, K). Points show medians. Error bars show interquartile ranges over five paired repetitions and are smaller than the markers.