
Follow-the-Perturbed-Leader for Adversarial Markov Decision Processes with Bandit Feedback

Yan Dai

Tsinghua University
yan-dai20@mails.tsinghua.edu.cn

Haipeng Luo

University of Southern California
haipengl@usc.edu

Liyu Chen

University of Southern California
liyuc@usc.edu

Abstract

We consider regret minimization for Adversarial Markov Decision Processes (AMDPs), where the loss functions are changing over time and adversarially chosen, and the learner only observes the losses for the visited state-action pairs (i.e., bandit feedback). While there has been a surge of studies on this problem using Online-Mirror-Descent (OMD) methods, very little is known about the Follow-the-Perturbed-Leader (FTPL) methods, which are usually computationally more efficient and also easier to implement since it only requires solving an offline planning problem. Motivated by this, we take a closer look at FTPL for learning AMDPs, starting from the standard episodic finite-horizon setting. We find some unique and intriguing difficulties in the analysis and propose a workaround to eventually show that FTPL is also able to achieve near-optimal regret bounds in this case. More importantly, we then find two significant applications: First, the analysis of FTPL turns out to be readily generalizable to delayed bandit feedback with order-optimal regret, while OMD methods exhibit extra difficulties (Jin et al., 2022). Second, using FTPL, we also develop the first no-regret algorithm for learning communicating AMDPs in the infinite-horizon setting with bandit feedback and stochastic transitions. Our algorithm is efficient assuming access to an offline planning oracle, while even for the easier full-information setting, the only existing algorithm (Chandrasekaran and Tewari, 2021) is computationally inefficient.

1 Introduction

Markov Decision Processes (MDPs) have long been used to model problems in reinforcement learning, where the agent takes sequential actions in an environment, leading to transitions among different states and observations on loss (or reward equivalently) signals. While the classical MDP model assumes a fixed loss function, there has been increasing interest in studying regret minimization under non-stationary or even adversarial loss functions via the Adversarial MDP (AMDP) model, starting from the work of Even-Dar et al. (2009).

Similar to other regret minimization problems, there are typically two categories of algorithms for AMDPs: those based on the Follow-the-Perturbed-Leader (FTPL) framework (Even-Dar et al., 2009; Neu et al., 2010, 2012; Chandrasekaran and Tewari, 2021) and those based on the Online-Mirror-Descent (OMD) or the closely related Follow-the-Regularized-Leader (FTRL) framework (Zimin and Neu, 2013; Rosenberg and Mansour, 2019a,b; Jin et al., 2020, 2021, 2022). FTPL methods are usually computationally more efficient and easier to implement as it only requires solving an offline

Table 1: An overview of the proposed algorithms/results and comparisons with related works.

Setting	Transition	Feedback	Algorithm	Regret ^a	Method	Effi. ^b
Episodic H -horizon AMDPs	Known	Bandit	Zimin and Neu (2013)	$\tilde{O}(H\sqrt{SAK})$	OMD	✓
			This work (Theorem 4)	$\tilde{O}(H^{3/2}\sqrt{SAK})$	FTPL	✓
			Jin et al. (2020)	$\tilde{O}(H^2S\sqrt{AK})$	OMD	✓
			This work (Theorem 5)	$\tilde{O}(H^2S\sqrt{AK})$	FTPL	✓
	Unknown	Bandit & Delayed	Delayed HEDGE (Jin et al., 2022)	$\tilde{O}(H^2S\sqrt{AK} + H^{3/2}\sqrt{S\mathfrak{D}})$	OMD	✗
			Delayed UOB-FTRL (Jin et al., 2022)	$\tilde{O}(H^2S\sqrt{AK} + H^{3/2}SA\sqrt{\mathfrak{D}})$	OMD	✓
			Delayed UOB-REPS (Jin et al., 2022)	$\tilde{O}(H^2S\sqrt{AK} + H^{3/4}(SA)^{1/4}\sqrt{\mathfrak{D}})$	OMD	✓
		This work (Theorem 6)	$\tilde{O}(H^2S\sqrt{AK} + H^{3/2}\sqrt{SA\mathfrak{D}})$	FTPL	✓	
Infinite-horizon AMDPs	Known	Full-info	Even-Dar et al. (2009)	$\tilde{O}(\tau^2\sqrt{T})$ (Ergodic)	OMD	✓
			Chandrasekaran and Tewari (2021)	$\tilde{O}(S^4\sqrt{T})$ (Deterministic)	FTPL	✓
			Chandrasekaran and Tewari (2021)	$\tilde{O}(D^2\sqrt{ST})$ (Commu)	FTPL	✗
		Bandit	Neu et al. (2014)	$\tilde{O}(\sqrt{\tau^3AT})$ (Ergodic)	OMD	✓
			Dekel and Hazan (2013)	$\tilde{O}(S^3AT^{2/3})$ (Deterministic)	OMD	✓
			This work (Theorem 7)	$\tilde{O}(A^{1/2}(SD)^{2/3}T^{5/6})$ (Commu)	FTPL	✓!
			This work (Theorem 8)	$\tilde{O}(A^{1/3}(SDT)^{2/3})$ (Commu)	OMD	✗
			Dekel et al. (2014)	$\Omega(S^{1/3}T^{2/3})$ (if only Commu)	—	—

^aHere, S and A are the number of states and actions respectively, K is the number of episodes, T is the total number of steps, \mathfrak{D} is the total amount of delay, τ is the mixing time of an ergodic MDP, and D is the diameter of a communicating MDP. Several related works use different notations from ours, and their regret bounds have been converted based on our notations. For infinite-horizon AMDPs, the extra assumptions are listed after the regret bounds, with ‘‘Ergodic’’ standing for ergodic MDPs, ‘‘Deterministic’’ standing for MDPs with deterministic transitions, and ‘‘Commu’’ standing for communicating MDPs (the weakest assumption).

^bThis column indicates the algorithm’s efficiency: ✓ means polynomial (in all parameters) time complexity, ✗ means $\Omega(A^S)$ time complexity, and ✓! means efficient assuming access to a planning oracle (that returns the best policy given all the MDP’s parameters). Note that FTPL-based algorithms are usually easier to implement compared to OMD/FTPL-based ones (both treated as OMD-based in this table as they are quite similar).

optimization problem (a.k.a. a *planning* problem in the MDP literature). In contrast, OMD/FTPL methods require solving convex optimization problems over a complicated occupancy measure space.

Despite its computational advantages and ease in implementation, FTPL methods are much less studied (especially for learning AMDPs) since they are harder to analyze, less versatile, and are believed to suffer worse regret compared to OMD/FTPL methods. A recent work by Wang and Dong (2020) disputes the last common belief and shows that, for episodic AMDPs with full-information feedback, FTPL also enjoys near-optimal regret, similarly to OMD/FTPL. Nevertheless, little is known about FTPL for learning AMDPs with the more challenging bandit feedback — to our knowledge, the only FTPL algorithm for this case is by Neu et al. (2010). However, that algorithm is analyzed under a strong assumption that every state is reachable by any policy with at least a constant probability $\alpha > 0$. Such an exploratory assumption is too strong to be used in realistic applications.

Motivated by this fact, we take a closer look at FTPL for learning AMDPs under bandit feedback, aiming at showing strong regret guarantees while enjoying its computational advantages. We start with the standard episodic finite-horizon setting and indeed find some intriguing difficulties compared to OMD/FTPL. After addressing these difficulties, we then show critical applications of FTPL methods to two more challenging setups: episodic AMDPs with *delayed* bandit feedback and *infinite-horizon* AMDPs with only communicating assumptions, with the latter result advancing the state-of-the-art. More specifically, our contributions are (see also Table 1 for a summary):

1. We start with the heavily studied episodic setting with K episodes, H steps in each episode, S states, and A actions. Our first intriguing observation is that: since the loss of each policy is linear in a non-binary vector (i.e., the occupancy measure), existing analysis for the stability term of FTPL fails, even though it works for the binary case (e.g., Neu and Bartók (2016)). Our next important observation is that there exists a simple fix to this issue that only leads to an extra H factor. This eventually leads to $\tilde{O}(H^{3/2}\sqrt{SAK})$ regret when the transition is known (Algorithm 1, Theorem 4), which is only \sqrt{H} factor larger than the near-optimal regret achieved

by OMD (Zimin and Neu, 2013), and $\tilde{O}(H^2 S \sqrt{AK})$ regret when the transition is unknown (Algorithm 3, Theorem 5), matching the state-of-the-art again achieved by OMD (Jin et al., 2020). See Section 3 for details.

2. We next find that compared to OMD, the analysis of FTPL is much easier to be generalized to the delayed feedback setting where losses for episode k are observed only at the end of episode $k + d_k$ for some $d_k \geq 0$ (Lancewicki et al., 2022; Jin et al., 2022). Indeed, these two prior works demonstrate the difficulty of analyzing OMD with delay feedback, with Lancewicki et al. (2022) only achieving $\tilde{O}((K + \mathfrak{D})^{2/3})$ regret (where $\mathfrak{D} = \sum_k d_k$ is the total amount of delay; dependence on other parameters is omitted) and Jin et al. (2022) improving it to $\tilde{O}(\sqrt{K + \mathfrak{D}})$ via either an inefficient algorithm or an efficient OMD-based algorithm with more involved analysis and/or new delayed-adapted loss estimators. FTPL, on the other hand, achieves $\tilde{O}(\sqrt{K + \mathfrak{D}})$ regret by a simple extension of the analysis (Theorem 6). The dependence on S and A is also better than the OMD method of (Jin et al., 2022) with the same kind of standard loss estimators (though worse than their best result with the delayed-adapted estimators; see Table 1 and Section 4 for details).
3. While our results above do not improve the best existing ones, our final application of FTPL provides the *first* result for learning infinite-horizon communicating AMDPs with bandit feedback and known stochastic transitions. Specifically, our algorithm achieves $\tilde{O}(A^{1/2}(SD)^{2/3}T^{5/6})$ regret (Algorithm 6, Theorem 7), where D is the diameter of the MDP and T is the total number of steps. It is efficient assuming access to an offline planning oracle (that returns the best stationary policy given a fixed transition function and a sequence of loss functions for each step). Previous results either only handle deterministic transitions (Dekel and Hazan, 2013) or full-information loss feedback (Chandrasekaran and Tewari, 2021). Moreover, the FTPL algorithm of Chandrasekaran and Tewari (2021) for stochastic transitions is inefficient even given the same planning oracle (since it explicitly adds independent noise to every policy). For completeness, we also provide an inefficient algorithm (Algorithm 7) that achieves $\tilde{O}(A^{1/2}(SDT)^{2/3})$ regret in our bandit setting, matching the $\Omega(T^{2/3})$ lower bound of Dekel et al. (2014) in terms of T . See Section 5 for details.

1.1 Related Work

Follow-the-Perturbed-Leader: FTPL is first proposed by Hannan (1957) and later popularized by Kalai and Vempala (2005). It has proven to be extremely powerful for structured online learning problems (such as online shortest path) since its implementation is as easy as solving the corresponding offline optimization problem (such as finding the shortest path of a given graph). Over the years, FTPL has been extended to problems with semi-bandit feedback (Neu, 2015; Neu and Bartók, 2016), contextual information (Syrkanis et al., 2016), non-linear losses (Dudík et al., 2020), smoothed adversaries (Block et al., 2022; Haghtalab et al., 2022), and others. However, FTPL for learning AMDPs under bandit feedback is poorly understood, which motivates this work. As we successfully show, improving our understanding of FTPL is indeed beneficial since it at least leads to new results for the infinite-horizon setting (in addition to its computational advantages for other settings). Below, we briefly review the literature of AMDPs for the three settings we consider.

Episodic Finite-Horizon AMDPs: Earlier works on this topic focus on the easier known transition case. In particular, the OMD-based O-REPS algorithm by Zimin and Neu (2013) achieves $\tilde{O}(H\sqrt{K})$ regret with full-information feedback and $\tilde{O}(H\sqrt{SAK})$ regret with bandit feedback, both optimal up to logarithmic factors. On the other hand, FTPL is recently shown to achieve $\tilde{O}(H^2\sqrt{K})$ regret with full-information feedback (Wang and Dong, 2020). As mentioned, the only FTPL algorithm for bandit feedback is by Neu et al. (2010), which guarantees $\tilde{O}(H^2\sqrt{AK}/\alpha)$ regret assuming that all states are reachable by any policy with a probability of at least α . In contrast, our FTPL algorithm removes this requirement and achieves $\tilde{O}(H^{3/2}\sqrt{SAK})$ regret, which is only \sqrt{H} away from optimal.

When the transition is unknown, with full-information feedback, the OMD-based algorithm UC-O-REPS (Rosenberg and Mansour, 2019a) achieves $\tilde{O}(H^2 S \sqrt{AK})$ regret, while the FTPL-based FPOP (Neu et al., 2012) is shown to achieve $\tilde{O}(H^2 S \sqrt{AK})$ regret as well (Wang and Dong, 2020). With bandit feedback, the OMD-based algorithm UOB-REPS (Jin et al., 2020) also achieves the same $\tilde{O}(H^2 S \sqrt{AK})$ regret. At the same time, our algorithm enjoys the same guarantee and is the first FTPL algorithm for bandit feedback and unknown transition. However, the current best lower bound for this problem is $\Omega(H^{3/2}\sqrt{SAK})$ (Jin et al., 2018), so there is still an $\mathcal{O}(\sqrt{HS})$ gap.

Besides OMD and FTPL, there is, in fact, another category of algorithms for learning AMDPs: policy optimization (Shani et al., 2020; Luo et al., 2021), which performs OMD in each *state* and is also efficient. However, the regret bounds are worse by at least an H factor (Luo et al., 2021).

Delayed Feedback: The most related works are Lancewicki et al. (2022) and Jin et al. (2022), and we refer the reader to the references therein for the literature on delayed feedback for different problems. Importantly, Jin et al. (2022) point out the unique difficulty when analyzing OMD/FTRL for AMDPs with delayed feedback. Circumventing this difficulty one way or another, they develop three algorithms: the first one, Delayed HEDGE, is inefficient; the second one, Delayed UOB-FTRL, achieves worse regret (\sqrt{SA} larger for the delay-related term) compared to ours; and the third one makes use of a delay-adapted estimator and achieves the best bound (see Table 1). We emphasize again that our FTPL analysis is much simpler and a direct extension of the non-delayed case. The current best lower bound for this problem is $\Omega(H^{3/2}\sqrt{SAK} + H\sqrt{\mathcal{D}})$ (Lancewicki et al., 2022).

Infinite-Horizon AMDPs: Learning AMDPs becomes significantly more difficult in the infinite horizon setting. As far as we know, all works in this line (including ours) assume a known transition function. Earlier works focus on the simpler case with a strong *ergodic* assumption (Even-Dar et al., 2009; Neu et al., 2014). For the more general *communicating* assumptions, a recent work (Chandrasekaran and Tewari, 2021) considers full-information feedback and develops an efficient FTPL algorithm for deterministic transitions with $\tilde{\mathcal{O}}(S^4\sqrt{T})$ regret and another inefficient FTPL algorithm for stochastic transitions with $\tilde{\mathcal{O}}(D^2\sqrt{ST})$ regret. Under bandit feedback, prior works only study deterministic transitions (Arora et al., 2012; Dekel and Hazan, 2013), with Dekel and Hazan (2013) achieving $\tilde{\mathcal{O}}(S^3AT^{2/3})$ regret, matching the lower bound (Dekel et al., 2014) for the T -dependency. Our results are the first for bandit feedback and stochastic transitions. Note that since bandit feedback is only more general, our oracle-efficient algorithm can also be applied to the full-information setting, while the only existing algorithm (Chandrasekaran and Tewari, 2021) is computationally inefficient.

2 Preliminaries

General Notations: We use $[N]$ to denote the set $\{1, 2, \dots, N\}$. For a (finite) set X , we use $\Delta(X) \triangleq \{x \in \mathbb{R}_{\geq 0}^{|X|} \mid \sum_{i=1}^{|X|} x_i = 1\}$ to denote the probability simplex over the set X . We use $\tilde{\mathcal{O}}(\cdot)$ to hide all terms logarithmic in H, S, A, K and T . Laplace(η) denotes the Laplace (also known as double-exponential) distribution with center 0 and parameter η , whose probability density is $f(x) = \frac{\eta}{2} \exp(-\eta|x|)$, $\forall x \in \mathbb{R}$. For an event \mathcal{E} , let $\mathbb{1}[\mathcal{E}]$ be its indicator. In episodic settings, let $\{\mathcal{F}_k\}_{k=0}^K$ be the natural filtration such that \mathcal{F}_k contains the history of episodes $1, \dots, k$. With a slight abuse of notation, in the infinite-horizon setting, we also use $\{\mathcal{F}_t\}_{t=0}^T$ to denote the natural filtration.

Episodic Adversarial Markov Decision Process: An episodic Adversarial Markov Decision Process (AMDP) is defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbb{P}, \ell, K, H, s^1)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathbb{P}: [H] \times \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition function, $\ell: [K] \times [H] \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the loss function unknown to the agent but fixed before the game (i.e., we are assuming an *oblivious* adversary),¹ K is the number of episodes, H is the horizon length, and $s^1 \in \mathcal{S}$ is the initial state. Denote by $S = |\mathcal{S}| < \infty$ and $A = |\mathcal{A}| < \infty$, the number of states and actions, respectively.

The agent interacts with the environment for K episodes. For the k -th one ($k \leq K$), she starts from the initial state s^1 and sequentially interacts with the environment for H steps. At the h -th step (where $h \in [H]$), the agent observes state $s_k^h \in \mathcal{S}$, chooses an action $a_k^h \in \mathcal{A}$, observes and suffers the loss $\ell_k^h(s_k^h, a_k^h)$ (bandit feedback),² and then transits to state s_k^{h+1} according to the probability distribution $\mathbb{P}^h(\cdot \mid s_k^h, a_k^h)$. After H steps, the episode ends and the agent proceeds to episode $k + 1$.

A (deterministic) *policy* of the agent is defined by $\pi = \{\pi^h: \mathcal{S} \rightarrow \mathcal{A}\}_{h \in [H]}$. Denote the set of all deterministic policies by Π . The expected loss incurred by policy $\pi \in \Pi$ for an episode with loss function $\hat{\ell}$ is denoted by $V(\pi; \hat{\ell}) \triangleq \mathbb{E} \left[\sum_{h=1}^H \hat{\ell}^h(s^h, \pi^h(s^h)) \mid s^{h+1} \sim \mathbb{P}^h(\cdot \mid s^h, \pi^h(s^h)), \forall h < H \right]$. Suppose the agent uses policies $\pi_1, \pi_2, \dots, \pi_K$ for episodes $1, 2, \dots, K$, respectively. The total expected loss of the agent is then $\mathbb{E} \left[\sum_{k=1}^K V(\pi_k; \ell_k) \right]$, where the expectation is taken with respect to

¹Note that the loss function can vary arbitrarily for different (k, h) -pairs, instead of being stochastic.

²On the other hand, in the easier full-information setting, the entire ℓ_k^h is revealed.

the agent’s private randomness. The baseline is the best deterministic policy in hindsight, defined by $\pi^* \in \operatorname{argmin}_{\pi \in \Pi} \sum_{k=1}^K V(\pi; \ell_k)$. The goal of the agent is to minimize her *regret* over K episodes, which is the difference between her total loss and that of π^* , formally defined as

$$\mathcal{R}_K \triangleq \mathbb{E} \left[\sum_{k=1}^K V(\pi_k; \ell_k) \right] - \sum_{k=1}^K V(\pi^*; \ell_k).$$

Episodic AMDPs with Delayed Feedback: This setup is exactly the same as the episodic AMDPs, except that the feedback $\{\ell_k^h(s_k^h, a_k^h)\}_{h=1}^H$ for episode k is only available after d_k episodes, i.e., at the end of the $(k + d_k)$ -th episode. Define $\mathcal{D} = \sum_{k=1}^K d_k$ to be the total feedback delay, assumed to be known to the agent as this assumption can be easily relaxed via a doubling trick (Thune et al., 2019).³

Infinite-Horizon AMDPs: Similar to episodic AMDPs, infinite-horizon AMDPs is defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbb{P}, \ell, T, s^1)$. Here, starting from the initial state $s^1 \in \mathcal{S}$, the agent interacts with the environment for T total steps without any reset, under the transition model $\mathbb{P}: \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ (which does not vary over time) and loss functions $\ell: [T] \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$. More specifically, at time $t \in [T]$, the agent observes state $s^t \in \mathcal{S}$, chooses an action $a^t \in \mathcal{A}$, observes and suffers loss $\ell^t(s^t, a^t)$, and then transits to $s^{t+1} \sim \mathbb{P}(\cdot | s^t, a^t)$. Her goal is also to minimize the regret, defined as

$$\mathcal{R}_T \triangleq \mathbb{E} \left[\sum_{t=1}^T \ell^t(s^t, a^t) \middle| s^{t+1} \sim \mathbb{P}(\cdot | s^t, a^t) \right] - \min_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^T \ell^t(s^t, \pi(s^t)) \middle| s^{t+1} \sim \mathbb{P}(\cdot | s^t, \pi(s^t)) \right], \quad (1)$$

where Π is now the set of all deterministic policies mapping from \mathcal{S} to \mathcal{A} . As pointed out by Bartlett and Tewari (2009), without any extra assumptions, sublinear regret is impossible for this problem due to the lack of resets. Earlier works make a strong *ergodic* assumption such that, intuitively, any mistake will be forgiven after logarithmic steps (Even-Dar et al., 2009). Here, we instead focus on the much weaker *communicating* assumption as in Chandrasekaran and Tewari (2021):

Definition 1 (Communicating MDP). *We call an MDP \mathcal{M} communicating if it has a finite diameter $D \triangleq \max_{s \neq s'} \min_{\pi \in \Pi} \mathbb{E}[T(s' | \mathcal{M}, \pi, s)]$ where $T(s' | \mathcal{M}, \pi, s)$ is the (random) time step when state s' is first reached by policy π starting from state s .*

Just like Chandrasekaran and Tewari (2021), for technical reasons, we also need the following mild assumption saying that there exists a special state for the agent to “park” there without moving.

Assumption 2. *There exist state $s^* \in \mathcal{S}$ and action $a^* \in \mathcal{A}$ such that $\mathbb{P}(s^* | s^*, a^*) = 1$.*

3 FTPL for Episodic AMDPs

In this section, we consider the basic (non-delayed) episodic setting. To best illustrate the unique difficulty we meet when analyzing FTPL and the way we address it, we first discuss the known-transition case (i.e., $\{\mathbb{P}^h\}_{h=1}^H$ is known to the agent), and then move on to unknown transitions.

3.1 Known Transition

Our algorithm follows the standard FTPL framework (see Algorithm 1). Ahead of time (as the adversary is oblivious), we sample a perturbation vector $z: [H] \times \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ so that $z^h(s, a)$ is an independent sample from Laplace(η) for some parameter η . At the beginning of episode k , given the loss estimators $\widehat{\ell}_1, \dots, \widehat{\ell}_{k-1}$ from previous episodes (whose construction will be specified later), we simply play the policy that minimizes the cumulative perturbed estimated loss (break tie arbitrarily):

$$\pi_k = \operatorname{argmin}_{\pi \in \Pi} \left(V(\pi; z) + \sum_{k'=1}^{k-1} V(\pi; \widehat{\ell}_{k'}) \right) = \operatorname{argmin}_{\pi \in \Pi} V(\pi; \widehat{\ell}_{0:k-1}),$$

where we use $\widehat{\ell}_{l:r}$ (where $0 \leq l \leq r \leq K$) as a shorthand notation for $\sum_{k'=l}^r \widehat{\ell}_{k'}$ and $\widehat{\ell}_0$ as an alias for z for notational convenience. This optimization over $\pi \in \Pi$ is a simple planning problem and can be solved by dynamic programming efficiently.

³As in Jin et al. (2022), we only consider delayed loss feedback, but not delayed trajectory feedback, since the latter only affects the transition estimation and can be handled similarly to Lancewicki et al. (2022).

Algorithm 1 FTPL for Episodic AMDPs with Bandit Feedback and Known Transition

Require: Laplace distribution parameter η . Geometric Re-sampling parameter L .

- 1: Sample perturbation $\widehat{\ell}_0 = z$ such that $z^h(s, a)$ is an independent sample of $\text{Laplace}(\eta)$.
 - 2: **for** $k = 1, 2, \dots, K$ **do**
 - 3: Calculate $\pi_k = \operatorname{argmin}_{\pi \in \Pi} V(\pi; \widehat{\ell}_{0:k-1})$ (via dynamic programming).
 - 4: **for** $h = 1, 2, \dots, H$ **do**
 - 5: Observe s_k^h , play $a_k^h = \pi_k(s_k^h)$, suffer and observe loss $\ell_k^h(s_k^h, a_k^h)$.
 - 6: Calculate loss estimator $\widehat{\ell}_k^h$ via Geometric Re-sampling (Neu and Bartók, 2016):
 - 7: **for** $M_k^h = 1, 2, \dots, L$ **do**
 - 8: Sample a fresh perturbation \tilde{z} in the same way as z .
 - 9: Calculate $\pi'_k = \operatorname{argmin}_{\pi \in \Pi} V(\pi; \widehat{\ell}_{1:k-1} + \tilde{z})$.
 - 10: Simulate π'_k for h steps starting from s^1 and following transitions $\mathbb{P}^1, \dots, \mathbb{P}^h$.
 - 11: **if** (s_k^h, a_k^h) is visited at step h or $M_k^h = L$ **then**
 - 12: Set $\widehat{\ell}_k^h(s, a) = M_k^h \cdot \ell_k^h(s_k^h, a_k^h) \cdot \mathbb{1}[(s_k^h, a_k^h) = (s, a)]$ and break.
-

Upon seeing s_k^h , a_k^h , and $\ell_k^h(s_k^h, a_k^h)$, we construct the loss estimator $\widehat{\ell}_k^h$ using the Geometric Re-sampling technique (Neu and Bartók, 2016). The idea is to repeat the sampling procedure (Line 8 to 10) until the same pair (s_k^h, a_k^h) is visited again at step h or this has been repeated L times for some parameter L . Let the total number of trials be M_k^h , then the estimator is defined as $\widehat{\ell}_k^h(s, a) = M_k^h \cdot \ell_k^h(s_k^h, a_k^h) \cdot \mathbb{1}[(s_k^h, a_k^h) = (s, a)]$ (Line 12). Note that the sampling procedure can be done freely without interacting with the environment as the transition is known. The rationale behind this estimator is that as long as L is reasonably large, M_k^h is a good approximation of the inverse probability of visiting (s_k^h, a_k^h) (which is hard to calculate directly for FTPL), making $\widehat{\ell}_k^h$ a good (and efficient) approximation of the standard importance weighted estimator (Zimin and Neu, 2013).

Analysis Sketch: While our algorithm follows the standard FTPL framework, we find some intriguing difficulty in the analysis that is unique to MDPs and undiscovered before. To illustrate this difficulty, let us first describe an overview of the analysis. First, since the loss estimators are almost unbiased (as shown by Neu and Bartók (2016)), we only need to focus on the regret with respect to the estimated losses, that is, $\mathbb{E} \left[\sum_{k=1}^K V(\pi_k; \widehat{\ell}_k) - \sum_{k=1}^K V(\pi^*; \widehat{\ell}_k) \right]$. Adding and subtracting $\mathbb{E} \left[\sum_{k=1}^K V(\pi_{k+1}; \widehat{\ell}_k) \right]$ (the loss of an imaginary “leader” that looks one episode ahead), our next goal is to bound the so-called *stability term* $\mathbb{E} \left[\sum_{k=1}^K V(\pi_k; \widehat{\ell}_k) - \sum_{k=1}^K V(\pi_{k+1}; \widehat{\ell}_k) \right]$ (the rest, usually referred as the *error term*, can be bounded by the standard “be-the-leader” lemma).

For the stability term, fix an episode k and define $p_k(\pi)$ as the probability of selecting π as π_k w.r.t. the randomness of the perturbation z . Further introduce the notion of *occupancy measures* (Altman, 1999; Neu et al., 2012): each policy $\pi \in \Pi$ induces H occupancy measures $\mu_\pi^h \in \Delta(\mathcal{S} \times \mathcal{A})$, $\forall h \in [H]$, where $\mu_\pi^h(s, a)$ denotes the probability of visiting (s, a) at step h if one executes policy π starting from the initial state s^1 . With these notations, each summand for the stability term becomes:

$$\mathbb{E} \left[V(\pi_k; \widehat{\ell}_k) - V(\pi_{k+1}; \widehat{\ell}_k) \right] = \mathbb{E} \left[\sum_{\pi \in \Pi} (p_k(\pi) - p_{k+1}(\pi)) \langle \mu_\pi, \widehat{\ell}_k \rangle \right],$$

where $\langle \mu_\pi, \widehat{\ell}_k \rangle \triangleq \sum_{h=1}^H \langle \mu_\pi^h, \widehat{\ell}_k^h \rangle$. This stability term is exactly in the same form as that in Lemma 8 of Neu and Bartók (2016) or Lemma 10 of Syrgkanis et al. (2016) for (contextual) semi-bandit problems, except that in their contexts, μ_π is a *binary* vector. This seemingly slight difference turns out to be important! Specifically, in these two prior works, they both show (using our notations):

$$p_{k+1}(\pi) \geq p_k(\pi) \exp \left(-\eta \langle \mu_\pi, \widehat{\ell}_k \rangle \right), \quad (2)$$

which, together with the fact $\exp(-x) \geq 1 - x$, implies

$$\mathbb{E} \left[V(\pi_k; \widehat{\ell}_k) - V(\pi_{k+1}; \widehat{\ell}_k) \right] \leq \eta \mathbb{E} \left[\sum_{\pi \in \Pi} p_k(\pi) \langle \mu_\pi, \widehat{\ell}_k \rangle^2 \right]. \quad (3)$$

Readers familiar with the online learning literature would have recognized the last expression, since it is also the standard stability term achieved by (inefficiently) running the classical HEDGE

algorithm (Freund and Schapire, 1997) over all policies (see e.g. Theorem 7.3 of Bubeck (2011)). Indeed, this term is small enough and can be shown to be of order $\mathcal{O}(\eta HSA)$ in our context after plugging in the definition of the loss estimators, which would then basically complete the proof.

However, not only do we realize that the proof of Eq. (2) heavily rely on the binary nature of μ_π , we in fact also find a counterexample where Eq. (3) is simply *incorrect* when μ_π is non-binary (see Appendix B.1.5 for the counterexample). We find this fact intriguing, because Eq. (3) holds for the aforementioned inefficient HEDGE algorithm regardless whether μ_π is binary or not.

Further examining the proof of Neu and Bartók (2016) and Syrgkanis et al. (2016), however, one can prove the following weaker version of Eq. (2) and Eq. (3) (namely Eq. (4) and Eq. (5) respectively).

Lemma 3 (Single-Step Stability). *For all $k \in [K]$ and $\pi \in \Pi$, we have*

$$p_{k+1}(\pi) \geq p_k(\pi) \exp\left(-\eta \sum_{h=1}^H \|\widehat{\ell}_k^h\|_1\right), \quad (4)$$

and thus

$$\mathbb{E}\left[V(\pi_k; \widehat{\ell}_k) - V(\pi_{k+1}; \widehat{\ell}_k)\right] \leq \eta \mathbb{E}\left[\left(\sum_{h=1}^H \|\widehat{\ell}_k^h\|_1\right) \sum_{\pi \in \Pi} p_k(\pi) \langle \mu_\pi, \widehat{\ell}_k \rangle\right]. \quad (5)$$

Fortunately, while Eq. (5) looks seemingly much larger than the classic bound Eq. (3), it is in fact at most larger by an H factor, that is, the right-hand side of Eq. (5) can be shown to be of order $\mathcal{O}(\eta H^2 SA)$ (see Lemma 12 in the appendix). Putting everything together, this allows us to prove the following regret guarantee for Algorithm 1, which is \sqrt{H} larger than the optimal bound (Zimin and Neu, 2013) due to the weakened stability bound. One may refer to Appendix B.1 for the formal proof.

Theorem 4. *For episodic AMDPs with bandit feedback and known transitions, Algorithm 1 with $\eta = 1/\sqrt{HSAK}$ and $L = \sqrt{SAK}/H$ ensures $\mathcal{R}_T = \widetilde{\mathcal{O}}(H^{3/2}\sqrt{SAK})$.*

3.2 Unknown Transition

To handle unknown transitions, we mostly follow existing ideas. First, for each episode k we maintain a confidence set \mathcal{P}_k of the transition function as Jin et al. (2022), whose construction is given in Appendix B.2.1. These confidence sets ensure that i) $\mathbb{P} \in \mathcal{P}_k$ with high probability and ii) $\mathcal{P}_{k+1} \subseteq \mathcal{P}_k$. Generalizing the notation $V(\pi; \widehat{\ell})$, we use $V(\pi; \widehat{\ell}, P)$ to denote the expected loss of policy π for an episode with loss function $\widehat{\ell}$ and transition P (so $V(\pi; \widehat{\ell}) = V(\pi; \widehat{\ell}, \mathbb{P})$). Then deploying the idea of optimism, we replace Line 3 of Algorithm 1 with $\pi_k = \operatorname{argmin}_{\pi \in \Pi} \min_{P \in \mathcal{P}_k} V(\pi; \widehat{\ell}_{0:k-1}, P)$, which can be efficiently found using Extended Value Iteration (Jaksch et al., 2010). As Wang and Dong (2020) argues, this is far more efficient than performing OMD over occupancy measure spaces.

We also need to modify the Geometric Re-sampling procedure accordingly since Line 10 requires using the true transition. To do so, we combine the procedure with the idea of *upper occupancy measures* from Jin et al. (2020). Specifically, in each trial we sample π'_k in the same way as π_k but with a fresh perturbation, then find the optimistic transition within \mathcal{P}_k that maximizes the probability of π'_k visiting (s_k^h, a_k^h) (which can be done efficiently using dynamic programming as shown by Jin et al. (2020)), and finally simulate π'_k for h steps following this optimistic transition.

Due to space limit, the full algorithm, Algorithm 3, is deferred to Appendix B.2. The analysis of the extra regret caused by the transition estimation error can be handled similarly to Jin et al. (2022) (more specifically, their Delayed HEDGE algorithm). As in previous works, this happens to be of order $\widetilde{\mathcal{O}}(H^2 S \sqrt{AK})$ and becomes the dominating term of the regret. This makes our final regret the same as the state-of-the-art (Jin et al., 2020), despite the weaker single-step stability lemma discussed in Section 3.1 (since this part is dominated now). Formally, we have the following regret guarantee.

Theorem 5. *For episodic AMDPs with bandit feedback and unknown transitions, Algorithm 3 with $\eta = 1/\sqrt{HSAK}$ and $L = \sqrt{SAK}/H$ ensures $\mathcal{R}_T = \widetilde{\mathcal{O}}(H^2 S \sqrt{AK})$.*

4 FTPL for Episodic AMDPs with Delayed Feedback

In this section, we show how our FTPL algorithm and analysis can be easily extended to the delayed feedback setting where the losses for episode k are only observed at the end of episode $k + d_k$.

The only change to the algorithm is to naturally delay the loss estimator construction until the loss feedback is received, and at each episode k only use the estimators constructed so far, i.e., $\Omega_k \triangleq \{k' \mid k' + d_{k'} < k\}$, to compute the current policy π_k . See Algorithm 4 in Appendix C.

To show how the analysis works, we focus on the known transition case at this moment for simplicity. Similar to the non-delayed case, the key is to bound the stability term, which was $\mathbb{E} \left[\sum_{k=1}^K V(\pi_k; \widehat{\ell}_k) - \sum_{k=1}^K V(\pi_{k+1}; \widehat{\ell}_k) \right]$ in Section 3.1, but now becomes $\mathbb{E} \left[\sum_{k=1}^K V(\pi_k; \widehat{\ell}_k) - \sum_{k=1}^K V(\widetilde{\pi}_{k+1}; \widehat{\ell}_k) \right]$ where $\widetilde{\pi}_{k+1} = \operatorname{argmin}_{\pi \in \Pi} V(\pi; \widehat{\ell}_{0:k})$ is a ‘‘cheating policy’’ (Gyorgy and Joulani, 2021; Jin et al., 2022) that uses all loss estimators from the first k episodes (which matches π_{k+1} for the non-delayed case). By the exact same analysis as Eq. (4) and Eq. (5), one can show

$$\mathbb{E} \left[V(\pi_k; \widehat{\ell}_k) - V(\widetilde{\pi}_{k+1}; \widehat{\ell}_k) \right] \leq \eta \mathbb{E} \left[\underbrace{\left(\sum_{k' \in [k] \setminus \Omega_k} \sum_{h=1}^H \|\widehat{\ell}_{k'}^h\|_1 \right)}_{\text{DIFF}} \sum_{\pi \in \Pi} p_k(\pi) \langle \mu_\pi, \widehat{\ell}_k \rangle \right],$$

where the DIFF term is the cumulative ℓ_1 norms of all the estimators used in computing $\widetilde{\pi}_{k+1}$ but not π_k (again, a direct generalization of Eq. (5) where only k satisfies such conditions for k'). It is then not hard to imagine that when summed over k , the DIFF term is eventually related to the total amount of delay $\mathfrak{D} = \sum_k d_k$. Indeed, the sum of all stability terms over K episodes can be shown to be of order $\mathcal{O}(\eta H^2 S A (K + \mathfrak{D}))$. This is basically all the extra elements we need in the proof. More generally for unknown transitions, we prove the following guarantee (see Appendix C for the proof).

Theorem 6. *For episodic AMDPs with delayed bandit feedback and unknown transitions, Algorithm 4 with $\eta = 1/\sqrt{HSA(K+\mathfrak{D})}$ and $L = \sqrt{HSA/H}$ ensures $\mathcal{R}_T = \widetilde{\mathcal{O}}(H^2 S \sqrt{AK} + H^{3/2} \sqrt{SA\mathfrak{D}})$.*

The simplicity of our analysis is similar to the Delayed HEDGE algorithm (Jin et al., 2022), but the latter is inefficient with time complexity $\Omega(A^S)$. The efficient Delayed UOB-FTRL algorithm (Jin et al., 2022) requires a more complicated analysis and only achieves $\widetilde{\mathcal{O}}(H^2 S \sqrt{AK} + H^{3/2} S A \sqrt{\mathfrak{D}})$ regret (which is worse than ours), while its improved variant Delayed UOB-REPS with a new *delay-adapted* estimator achieves the current best bound $\widetilde{\mathcal{O}}(H^2 S \sqrt{AK} + H^{5/4} (SA)^{1/4} \sqrt{\mathfrak{D}})$. However, it is unclear to us whether such delay-adapted estimators can help improve FTPL. Finally, we again remark that the current best lower bound is $\Omega(H^{3/2} \sqrt{SAK} + H \sqrt{\mathfrak{D}})$ (Lancewicki et al., 2022).

5 FTPL for Infinite-Horizon AMDPs

At last, we discuss how FTPL can be used to derive the first no-regret algorithm for infinite-horizon communicating AMDPs with bandit feedback and (known) stochastic transition. Note that learning infinite-horizon AMDPs is much more difficult due to the lack of resets (in a sense, this is like a finite-horizon problem but with only one long episode with T steps). Another way to see the difficulty is that the benchmark in the regret definition Eq. (1) is evaluated on states generated by following π^* repeatedly for T rounds, without any resets. From a technical viewpoint, this requires the algorithm to also make sure that, when following a policy π , its suffered loss is indeed close to the total loss if π has been followed since the very beginning, which is unnatural without ergodic assumptions.

Chandrasekaran and Tewari (2021) resolve this issue by the combination of two ideas. First, under the mild Assumption 2, they show that whenever the agent wants to switch the current policy to another policy π , there exists a procedure to make sure that after $\mathcal{O}(D^2)$ steps of a transition phase, the agent’s state distribution is exactly the same as that induced by following π from the very beginning. That is, after this switching procedure, the agent can ‘‘pretend’’ that she has followed π all the time. Second, since this procedure requires a cost of $\mathcal{O}(D^2)$ steps (where the loss of the agent can be arbitrarily bad and only trivially bounded by $\mathcal{O}(D^2)$), the algorithm needs to switch its policy infrequently.

Our algorithm follows the same ideas. However, while low-switching is relatively easy to ensure in the full-information case without paying extra regret, it is known that with bandit feedback there is an unavoidable trade-off between the number of switches and the regret, which can be optimally balanced via a simple epoching scheme (Dekel et al., 2014). To this end, we divide the total T steps into $J = o(T)$ epochs, each with length $H = T/J = \omega(D^2)$. At the beginning of the j -th epoch, we compute a new policy π_j , apply the switching procedure of Chandrasekaran and Tewari (2021) to adjust the state distribution (see Algorithm 5), and finally follow the same policy π_j for the rest of the epoch. This clearly only introduces J switches, which contributes to at most $\mathcal{O}(JD^2)$ extra regret.

It remains to specify how to find π_j in epoch j using FTPL. The key difference compared to the episodic case is that, due to the lack of resets, we need to add perturbation to *every* time step instead of just to each of the H steps of an episode. We then still play the policy that minimizes the cumulative estimated losses plus all the perturbed losses. Formally, π_j is defined as:

$$\pi_j = \operatorname{argmin}_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^{(j-1)T/J} \widehat{\ell}^t(s^t, \pi(s^t)) + \sum_{t=1}^T z^t(s^t, \pi(s^t)) \middle| s^{t+1} \sim \mathbb{P}(\cdot \mid s^t, \pi(s^t)), \forall t \right], \quad (6)$$

where $\{z^t : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}\}_{t \in [T]}$ is such that each $z^t(s, a)$ is an independent sample of Laplace(η), and each $\widehat{\ell}^t$ is the estimator of ℓ^t constructed from the Geometric Re-sampling procedure.

Unfortunately, as far as we know, there is in fact no existing polynomial time algorithm for solving Eq. (6) (the difficulty comes from the restriction on *stationary* policies whose behavior does not vary over time). Even if the losses are stochastic, the problem is only known to be P-hard (Papadimitriou and Tsitsiklis, 1987; Mundhenk et al., 2000) and no polynomial algorithm has been developed.

However, note that this optimization is exactly in the same form as the benchmark in the regret definition Eq. (1). Following many prior works such as Dudík et al. (2020); Block et al. (2022); Haghtalab et al. (2022), we thus assume access to a planning oracle that solves this offline problem, making our algorithm only oracle-efficient instead of truly polynomial-time-efficient. Note that even given this oracle, the algorithm of Chandrasekaran and Tewari (2021) is inefficient since it creates independent perturbation for *each* of the A^S policies, while our perturbation is much more compact.

In terms of the analysis, the key extra challenge is caused by having T perturbed losses. Indeed, the same analysis from the episodic case (Lemma 44) would lead to a term of order $\widetilde{O}(T/\eta)$, which is prohibitively large. Instead, inspired by Syrgkanis et al. (2016), we provide a different analysis showing that this can be improved to $\widetilde{O}(S\sqrt{AT}/\eta)$, which has worse dependencies on S and A but better dependency on T , the key to ensure sub-linear regret eventually. To conclude, our FTPL algorithm achieves the following guarantee (see Appendix D.1 for the full algorithm and analysis).

Theorem 7. *For infinite-horizon AMDPs with bandit feedback and known transitions, Algorithm 6 with $\eta = \frac{S^{1/3}}{D^{2/3}T^{1/3}}$, $J = \frac{S^{2/3}A^{1/2}T^{5/6}}{D^{4/3}}$ and $L = \frac{S^{1/3}A^{1/2}T^{1/6}}{D^{2/3}}$ ensures $\mathcal{R}_T = \widetilde{O}(A^{1/2}(SD)^{2/3}T^{5/6})$.*

We emphasize again that this is the first (oracle-efficient) algorithm for this setting. Even in the easier full-information setting (where ℓ^t is fully revealed at the end of time t), our algorithm also has its computational advantages compared to that of Chandrasekaran and Tewari (2021), since, as mentioned, their algorithm requires $\Omega(A^S)$ complexity (albeit with a better regret bound $\widetilde{O}(D^2\sqrt{ST})$).

The best lower bound for this setting is $\Omega(S^{1/3}T^{2/3})$ (Dekel et al., 2014). Dekel and Hazan (2013) achieve $\widetilde{O}(S^3AT^{2/3})$ but only when the transition is deterministic. For completeness, we provide a HEDGE-based *inefficient* algorithm (Appendix D.2) for general stochastic transitions, which achieves the optimal regret in terms of the dependence on T , improving our oracle-efficient FTPL algorithm.

Theorem 8. *For infinite-horizon AMDPs with bandit feedback and known transitions, Algorithm 7 with $\eta = \frac{S^{1/3}}{A^{1/3}(DT)^{2/3}}$ and $J = \frac{(ST)^{2/3}A^{1/3}}{D^{4/3}}$ ensures $\mathcal{R}_T = \widetilde{O}(A^{1/3}(SDT)^{2/3})$.*

6 Conclusion

In this paper, we designed FTPL-based algorithms for adversarial MDPs with bandit feedback in various settings, including episodic settings, delayed feedback settings and infinite-horizon settings. Our algorithms are easy to implement as they only require solving the offline planning problem, and in some cases they match the state-of-the-art performance or are even the first ever no-regret algorithms.

One interesting open question is whether, despite our counterexample, Eq. (3) can still hold with a larger constant for the right-hand side, either with our current algorithm or via some modified versions (for example with a different kind of perturbation). Achieving this would lead to an improved version of Lemma 3 and thus give the near-optimal delay-related regret term $\widetilde{O}(H^{3/2}\sqrt{\mathfrak{D}})$ for the delayed feedback setting, which is not currently achieved by any existing algorithms.

An alternative direction is to try to equip our Algorithm 4 (for episodic AMDPs with delayed feedback) with the “delay-adapted” loss estimators proposed by Jin et al. (2022). As their analysis

heavily relies on the exponential weight scheme (see their Lemma D.7, which bounds KL divergences between consecutive policies), it is unclear to us whether FTPL enjoys a similar property.

Another important future direction is to improve our results in the infinite-horizon setting, such as improving the $\tilde{O}(T^{5/6})$ oracle-efficient regret upper bound, removing the usage of oracles, or dealing with the unknown transition case (which has not yet been studied at all).

There are also several possible generalizations of our setting. For example, we only assume the losses to be adversarial. Further incorporating evolving transition is an important next step. There is already an FTPL-based algorithm (Yu and Mannor, 2009) for evolving dynamics (though they are assuming ergodic infinite-horizon MDPs), which builds upon the FTPL analysis by Even-Dar et al. (2009) (see their Lemma III.3). Although our work directly improves the performance guarantee of Even-Dar et al. (2009), it is highly unclear whether we can adopt the algorithm of Yu and Mannor (2009) for unknown-transition episodic MDPs (they assumed the transitions to be revealed after each episode) or infinite-horizon weakly communicating MDPs. Solving either case will be interesting. Moreover, considering dynamic regret instead of static regret can also be challenging.

Acknowledgments and Disclosure of Funding

We greatly acknowledge Vasilis Syrgkanis for the helpful discussion about whether their single-step stability lemma (Syrgkanis et al., 2016, Lemma 10) holds for non-binary action spaces. We also thank the anonymous reviewers for their insightful comments, which we greatly benefit from. HL is supported by NSF Award IIS-1943607 and a Google Faculty Research Award.

References

- Eitan Altman. *Constrained Markov decision processes: stochastic modeling*. Routledge, 1999.
- Raman Arora, Ofer Dekel, and Ambuj Tewari. Deterministic mdps with adversarial rewards and bandit feedback. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 93–101, 2012.
- Peter Bartlett and Ambuj Tewari. Regal: a regularization based algorithm for reinforcement learning in weakly communicating mdps. In *Uncertainty in Artificial Intelligence: Proceedings of the 25th Conference*, pages 35–42. AUAI Press, 2009.
- Adam Block, Yuval Dagan, Noah Golowich, and Alexander Rakhlin. Smoothed online learning is as easy as statistical learning. In *Conference on Learning Theory*, pages 1716–1786. PMLR, 2022.
- Sébastien Bubeck. Introduction to online optimization. *Lecture Notes*, 2, 2011.
- Gautam Chandrasekaran and Ambuj Tewari. Online learning in adversarial mdps: Is the communicating case harder than ergodic? *arXiv preprint arXiv:2111.02024*, 2021.
- Ofer Dekel and Elad Hazan. Better rates for any adversarial deterministic mdp. In *International Conference on Machine Learning*, pages 675–683. PMLR, 2013.
- Ofer Dekel, Jian Ding, Tomer Koren, and Yuval Peres. Bandits with switching costs: $T^{2/3}$ regret. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 459–467, 2014.
- Miroslav Dudík, Nika Haghtalab, Haipeng Luo, Robert E Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. Oracle-efficient online learning and auction design. *Journal of the ACM (JACM)*, 67(5):1–57, 2020.
- Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Online markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Andras Gyorgy and Pooria Joulani. Adapting to delays and data in adversarial multi-armed bandits. In *International Conference on Machine Learning*, pages 3988–3997. PMLR, 2021.

- Nika Haghtalab, Yanjun Han, Abhishek Shetty, and Kunhe Yang. Oracle-efficient online learning for beyond worst-case adversaries. *arXiv preprint arXiv:2202.08549*, 2022.
- James Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3(2):97–139, 1957.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11:1563–1600, 2010.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? *Advances in neural information processing systems*, 31, 2018.
- Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu. Learning adversarial markov decision processes with bandit feedback and unknown transition. In *International Conference on Machine Learning*, pages 4860–4869. PMLR, 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 34*, pages 20491–20502, 2021.
- Tiancheng Jin, Tal Lancelwicki, Haipeng Luo, Yishay Mansour, and Aviv Rosenberg. Near-optimal regret for adversarial mdp with delayed bandit feedback. *arXiv preprint arXiv:2201.13172*, 2022.
- Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- Tal Lancelwicki, Aviv Rosenberg, and Yishay Mansour. Learning adversarial markov decision processes with delayed feedback. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022*, pages 7281–7289. AAAI Press, 2022.
- Haipeng Luo, Chen-Yu Wei, and Chung-Wei Lee. Policy optimization in adversarial mdps: Improved exploration via dilated bonuses. *Advances in Neural Information Processing Systems*, 34, 2021.
- Martin Mundhenk, Judy Goldsmith, Christopher Lusena, and Eric Allender. Complexity of finite-horizon markov decision process problems. *Journal of the ACM (JACM)*, 47(4):681–720, 2000.
- Gergely Neu. First-order regret bounds for combinatorial semi-bandits. In *Conference on Learning Theory*, pages 1360–1375. PMLR, 2015.
- Gergely Neu and Gábor Bartók. An efficient algorithm for learning with semi-bandit feedback. In *International Conference on Algorithmic Learning Theory*, pages 234–248. Springer, 2013.
- Gergely Neu and Gábor Bartók. Importance weighting without importance weights: An efficient algorithm for combinatorial semi-bandits. *Journal of Machine Learning Research*, 17(154):1–21, 2016.
- Gergely Neu, András György, and Csaba Szepesvári. The online loop-free stochastic shortest-path problem. In *COLT 2010 - The 23rd Conference on Learning Theory*, pages 231–243. Omnipress, 2010.
- Gergely Neu, Andras Gyorgy, and Csaba Szepesvári. The adversarial stochastic shortest path problem with unknown transition probabilities. In *Artificial Intelligence and Statistics*, pages 805–813. PMLR, 2012.
- Gergely Neu, András György, Csaba Szepesvári, and András Antos. Online markov decision processes under bandit feedback. *IEEE Transactions on Automatic Control*, 59:676–691, 2014.
- Christos H Papadimitriou and John N Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.
- Aviv Rosenberg and Yishay Mansour. Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*, pages 5478–5486. PMLR, 2019a.
- Aviv Rosenberg and Yishay Mansour. Online stochastic shortest path with bandit feedback and unknown transition function. *Advances in Neural Information Processing Systems*, 32, 2019b.

- Lior Shani, Yonathan Efroni, Aviv Rosenberg, and Shie Mannor. Optimistic policy optimization with bandit feedback. In *International Conference on Machine Learning*, pages 8604–8613. PMLR, 2020.
- Vasilis Syrgkanis, Akshay Krishnamurthy, and Robert Schapire. Efficient algorithms for adversarial contextual learning. In *International Conference on Machine Learning*, pages 2159–2168. PMLR, 2016.
- Tobias Sommer Thune, Nicolò Cesa-Bianchi, and Yevgeny Seldin. Nonstochastic multiarmed bandits with unrestricted delays. *Advances in Neural Information Processing Systems*, 32, 2019.
- Yuanhao Wang and Kefan Dong. Refined analysis of fpl for adversarial markov decision processes. *arXiv preprint arXiv:2008.09251*, 2020.
- Jia Yuan Yu and Shie Mannor. Arbitrarily modulated markov decision processes. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 2946–2953. IEEE, 2009.
- Alexander Zimin and Gergely Neu. Online learning in episodic markovian decision processes by relative entropy policy search. *Advances in neural information processing systems*, 26, 2013.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) Check the Related Work section (especially Table 1) and the Conclusion section.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[No\]](#) Pure theoretical paper.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#) Check the Preliminaries section.
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#) A sketched analysis of Theorem 4 is presented in Section 3.1 while all detailed proofs are contained in the appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[N/A\]](#)
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[N/A\]](#)
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[N/A\]](#)
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[N/A\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[N/A\]](#)
 - (b) Did you mention the license of the assets? [\[N/A\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)

5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Supplementary Materials

A	Notations	14
B	Analysis of Episodic AMDP Algorithms	15
B.1	Known Transition Case (Theorem 4)	15
B.2	Unknown Transition Case (Theorem 5)	19
C	Analysis of Episodic AMDP Algorithms with Delayed Feedback (Theorem 6)	28
C.1	Regret Decomposition	28
D	Analysis of Infinite-horizon AMDP Algorithms	31
D.1	FTPL-Based Efficient Algorithm (Theorem 7)	31
D.2	Hedge-Based Inefficient Algorithm (Theorem 8)	37
E	Auxiliary Lemmas	39
E.1	Geometric Re-sampling Properties	39
E.2	Importance Weighting Properties	41
E.3	Auxiliary Lemmas for Error Terms	41

A Notations

We summarize our notations used in the appendix below:

- For a policy $\pi \in \Pi$ and a transition $P: [H] \times \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$, the occupancy measure of π at the h -th step ($h \in [H]$) is defined as

$$\mu_{\pi}^h(s, a; P) = \Pr\{(s^h, a^h) = (s, a) \mid a^h = \pi^h(s^h), s^{h+1} \sim P^h(\cdot \mid s^h, a^h), s^1\}.$$

We will use $\mu_{\pi}^h(P)$ to denote the vector $\{\mu_{\pi}^h(s, a; P)\}_{(s,a) \in \mathcal{S} \times \mathcal{A}}$. Specifically, if P is the true transition \mathbb{P} , we will abbreviate it as μ_{π}^h for simplicity.

- With a slight abuse of notation, for infinite-horizon AMDPs, we will also use the same notation $\mu_{\pi}^t \in \Delta(\mathcal{S} \times \mathcal{A})$ ($t \in [T]$) to refer to the occupancy measure of π at time slot t , starting from the first state s^1 and following the transition \mathbb{P} (as we do not consider unknown transition cases for infinite-horizon AMDPs, we will always abbreviate the transitions).
- For a policy $\pi \in \Pi$, a transition $P: [H] \times \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ and a loss function $\widehat{\ell}: [H] \times \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$, the value function is defined as

$$V(\pi; \widehat{\ell}, P) = \mathbb{E} \left[\widehat{\ell}^h(s^h, a^h) \mid a^h = \pi(s^h), s^{h+1} \sim P^h(s^h, a^h), s^1 \right] = \sum_{h=1}^H \langle \mu_{\pi}^h(P), \widehat{\ell}^h \rangle.$$

- A perturbation $z: \mathbb{R}^{[H] \times \mathcal{S} \times \mathcal{A}}$ is a fresh sample such that

$$z^h(s, a) \sim \text{Laplace}(\eta) \text{ and each entry is independently sampled.}$$

For simplicity in notations, we use $\widehat{\ell}_0$ as an alias of z .

- For a sequence of loss functions $\widehat{\ell}_1, \widehat{\ell}_2, \dots, \widehat{\ell}_k$, we use $\widehat{\ell}_{1:k}$ to denote $\sum_{k'=1}^k \widehat{\ell}_{k'}$.

Algorithm 2 FTPL for Episodic AMDPs with Bandit Feedback and Known Transition

Require: Laplace distribution parameter η . Geometric Re-sampling parameter L .

- 1: Sample perturbation $\widehat{\ell}_0 = z$ such that $z^h(s, a)$ is an independent sample of $\text{Laplace}(\eta)$.
 - 2: **for** $k = 1, 2, \dots, K$ **do**
 - 3: Calculate $\pi_k = \operatorname{argmin}_{\pi \in \Pi} V(\pi; \widehat{\ell}_{0:k-1})$ (via dynamic programming).
 - 4: **for** $h = 1, 2, \dots, H$ **do**
 - 5: Observe s_k^h , play $a_k^h = \pi_k(s_k^h)$, suffer and observe loss $\ell_k^h(s_k^h, a_k^h)$.
 - 6: Calculate loss estimator $\widehat{\ell}_k^h$ via Geometric Re-sampling (Neu and Bartók, 2016):
 - 7: **for** $M_k^h = 1, 2, \dots, L$ **do**
 - 8: Sample a fresh perturbation \tilde{z} in the same way as z .
 - 9: Calculate $\pi'_k = \operatorname{argmin}_{\pi \in \Pi} V(\pi; \widehat{\ell}_{1:k-1} + \tilde{z})$.
 - 10: Simulate π'_k for h steps starting from s^1 and following transitions $\mathbb{P}^1, \dots, \mathbb{P}^h$.
 - 11: **if** (s_k^h, a_k^h) is visited at step h or $M_k^h = L$ **then**
 - 12: Set $\widehat{\ell}_k^h(s, a) = M_k^h \cdot \ell_k^h(s_k^h, a_k^h) \cdot \mathbb{1}[(s_k^h, a_k^h) = (s, a)]$ and break.
-

B Analysis of Episodic AMDP Algorithms

B.1 Known Transition Case (Theorem 4)

For convenience, we restate the algorithm for episodic AMDPs with bandit feedback and known transitions in Algorithm 2. As shown by Syrgkanis et al. (2016, Appendix A.2), for an oblivious adversary (which is our case), it suffices to draw the perturbations once at the beginning of the interaction (i.e., the perturbation z is fixed throughout the game).

Then we give the proof of Theorem 4. As sketched in the main text, we define the following probability, as-if we are resampling a perturbation z for each round:

$$p_k(\pi) = \Pr_z\{\pi_k = \pi \mid \widehat{\ell}_1, \widehat{\ell}_2, \dots, \widehat{\ell}_{k-1}\}.$$

Note that, as mentioned in (Syrgkanis et al., 2016, Appendix A.2), p_k is just the probability of picking π at episode k given all history from episodes $1, 2, \dots, k-1$. Now, we decompose our regret \mathcal{R}_K into the following three terms:

$$\begin{aligned} \mathcal{R}_K &= \mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \sum_{\pi \in \Pi} p_k(\pi) \langle \mu_\pi^h, \ell_k^h \rangle - \sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi^*}^h, \ell_k^h \rangle \right] \\ &= \underbrace{\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi_k}^h, \ell_k^h - \widehat{\ell}_k^h \rangle + \sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi^*}^h, \widehat{\ell}_k^h - \ell_k^h \rangle \right]}_{\text{GR error term}} \\ &\quad + \underbrace{\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \sum_{\pi \in \Pi} p_{k+1}(\pi) \langle \mu_\pi^h, \widehat{\ell}_k^h \rangle - \sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi^*}^h, \widehat{\ell}_k^h \rangle \right]}_{\text{Error term}} \\ &\quad + \underbrace{\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \sum_{\pi \in \Pi} (p_k(\pi) - p_{k+1}(\pi)) \langle \mu_\pi^h, \widehat{\ell}_k^h \rangle \right]}_{\text{Stability term}}. \end{aligned}$$

B.1.1 Bounding the GR Error Term

Lemma 9 (Bounding GR Error Term). *The GR error term is bounded by*

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi_k}^h, \ell_k^h - \widehat{\ell}_k^h \rangle + \sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi^*}^h, \widehat{\ell}_k^h - \ell_k^h \rangle \right] \leq \frac{SAHK}{eL}.$$

Proof. First notice that, from Lemma 38, $\mathbb{E}[\widehat{\ell}_k^h(s, a) \mid \mathcal{F}_{k-1}] \leq \ell_k^h(s, a)$. Moreover, as π^* is deterministic (i.e., it does not depend on the randomness from the algorithm), the second term

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi^*}^h, \widehat{\ell}_k^h - \ell_k^h \rangle \right] = \mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi^*}^h, \mathbb{E}[\widehat{\ell}_k^h \mid \mathcal{F}_{k-1}] - \ell_k^h \rangle \right] \leq 0.$$

For the first term, again by Lemma 38, we have

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi_k}^h, \ell_k^h - \widehat{\ell}_k^h \rangle \right] = \sum_{k=1}^K \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbb{E} [\mu_{\pi_k}^h(s, a) \cdot (1 - q_k^h(s, a))^L \ell_k^h(s, a)],$$

where $q_k^h(s, a)$ is the probability of visiting (s, a) in a single trial of the Geometric Re-sampling process, which is just (note that q_k^h itself is also a random variable as p_k is non-deterministic)

$$q_k^h(s, a) = \mathbb{E}[\mu_{\pi_k}^h(s, a)] = \sum_{\pi \in \Pi} p_k(\pi) \mu_{\pi}^h(s, a)$$

in our case. By noticing that $q(1 - q)^L \leq qe^{-Lq} \leq \frac{1}{eL}$ for all $q \geq 0$ (Neu and Bartók, 2016), we have

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi_k}^h, \ell_k^h - \widehat{\ell}_k^h \rangle \right] \leq HKSA \frac{1}{eL} = \frac{SAHK}{eL},$$

as claimed. \square

B.1.2 Bounding the Error Term

Lemma 10 (Bounding Error Term). *The error term is bounded by*

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \sum_{\pi \in \Pi} p_{k+1}(\pi) \langle \mu_{\pi}^h, \widehat{\ell}_k^h \rangle \right] - \sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi^*}^h, \widehat{\ell}_k^h \rangle \leq \frac{2H}{\eta} (1 + \ln(SA)).$$

Proof. The proof uses the standard “be-the-leader” technique. For simplicity, we rewrite the error term as

$$\mathbb{E} \left[\sum_{k=1}^K V(\pi_{k+1}; \widehat{\ell}_k, \mathbb{P}) - \sum_{k=1}^K V(\pi^*; \widehat{\ell}_k, \mathbb{P}) \right].$$

Now consider the summation inside the expectation. If we add an extra term $V(\pi_1; \widehat{\ell}_0, \mathbb{P}) - V(\pi^*; \widehat{\ell}_0, \mathbb{P})$ where $\widehat{\ell}_0 = z$ is the perturbation, we will have

$$\begin{aligned} & \sum_{k=0}^K V(\pi_{k+1}; \widehat{\ell}_k, \mathbb{P}) - V(\pi^*; \widehat{\ell}_{0:K}, \mathbb{P}) \\ & \stackrel{(a)}{\leq} \sum_{k=0}^K V(\pi_{k+1}; \widehat{\ell}_k, \mathbb{P}) - V(\pi_{K+1}; \widehat{\ell}_{0:K}, \mathbb{P}) = \sum_{k=0}^{K-1} V(\pi_{k+1}; \widehat{\ell}_k, \mathbb{P}) - V(\pi_{K+1}; \widehat{\ell}_{0:K-1}, \mathbb{P}) \\ & \stackrel{(b)}{\leq} \sum_{k=0}^{K-1} V(\pi_{k+1}; \widehat{\ell}_k, \mathbb{P}) - V(\pi_K; \widehat{\ell}_{0:K-1}, \mathbb{P}) = \sum_{k=0}^{K-2} V(\pi_{k+1}; \widehat{\ell}_k, \mathbb{P}) - V(\pi_K; \widehat{\ell}_{0:K-2}, \mathbb{P}) \\ & \leq \dots \leq V(\pi_1; \widehat{\ell}_0, \mathbb{P}) - V(\pi_2; \widehat{\ell}_0, \mathbb{P}) \stackrel{(c)}{\leq} 0, \end{aligned}$$

where (a) used the optimality of π_{K+1} w.r.t. $\widehat{\ell}_{0:K}$, (b) used the optimality of π_K w.r.t. $\widehat{\ell}_{0:K-1}$ and so on, until the last step (c) where the optimality of π_1 w.r.t. $\widehat{\ell}_0$ is used. So we have

$$\mathbb{E} \left[\sum_{k=1}^K V(\pi_{k+1}; \widehat{\ell}_k, \mathbb{P}) - \sum_{k=1}^K V(\pi^*; \widehat{\ell}_k, \mathbb{P}) \right] \leq \mathbb{E}[V(\pi^*; \widehat{\ell}_0, \mathbb{P}) - V(\pi_1; \widehat{\ell}_0, \mathbb{P})].$$

By the notation of occupancy measures, we can rewrite it as

$$\mathbb{E} \left[\sum_{h=1}^H \langle \mu_{\pi^*}^h, \widehat{\ell}_0^h \rangle - \sum_{h=1}^H \langle \mu_{\pi_1}^h, \widehat{\ell}_0^h \rangle \right] \leq 2 \sum_{h=1}^H \mathbb{E}[\|\widehat{\ell}_0^h\|_\infty].$$

Recall that $\ell_0^h(s, a) \sim \text{Laplace}(\eta)$, so we have

$$\mathbb{E}[\|\widehat{\ell}_0^h\|_\infty] = \mathbb{E} \left[\max_{s,a} |\ell_0^h(s, a)| \right] \leq \frac{1 + \ln(SA)}{\eta},$$

where the last step is due to the fact that $|\ell_0^h(s, a)|$ is an exponential distribution and Lemma 44. \square

B.1.3 Bounding the Stability Term

For the stability term, we first prove the following ‘‘single-step stability’’ lemma that we stated without proof in the main body.

Lemma 11 (Single-Step Stability). *For all $k \in [K]$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$,*

$$p_{k+1}(\pi) \geq p_k(\pi) \exp \left(-\eta \sum_{h=1}^H \|\widehat{\ell}_k^h\|_1 \right), \quad \forall \pi \in \Pi.$$

Proof. For simplicity, we use $\pi = \text{best}(\ell)$ to denote $\pi = \text{argmin}_{\pi \in \Pi} V(\pi; \ell, \mathbb{P})$. Then we have

$$\begin{aligned} p_k(\pi) &= \int_z \mathbb{1} \left[\pi = \text{best} \left(\widehat{\ell}_{1:k-1} + z \right) \right] f(z) \, dz \\ &= \int_z \mathbb{1} \left[\pi = \text{best} \left(\widehat{\ell}_{1:k-1} + (z + \widehat{\ell}_k) \right) \right] f(z + \widehat{\ell}_k) \, dz \\ &= \int_z \mathbb{1} \left[\pi = \text{best} \left(\widehat{\ell}_{1:k} + z \right) \right] f(z + \widehat{\ell}_k) \, dz, \end{aligned}$$

where $f(z)$ is the probability density function of z and the second step made use of the fact that $z + \widehat{\ell}_k$ is still linear in z . Moreover,

$$p_{k+1}(\pi) = \int_z \mathbb{1} \left[\pi = \text{best} \left(\widehat{\ell}_{1:k} + z \right) \right] f(z) \, dz.$$

Recall that the definition of $f(z)$ is just $f(z) = \prod_{h=1}^H \sum_{s,a} \exp(-\eta|z^h(s, a)|) = \prod_{h=1}^H \exp(-\eta\|z^h\|_1)$ as each entry of z is i.i.d. We thus have

$$f(z + \widehat{\ell}_k) = \prod_{h=1}^H \exp \left(-\eta \left(\|z^h + \widehat{\ell}_k\|_1 - \|z^h\|_1 \right) \right) f(z),$$

which gives

$$\frac{f(z + \widehat{\ell}_k)}{f(z)} \in \left[\exp \left(-\eta \sum_{h=1}^H \|\widehat{\ell}_k^h\|_1 \right), \exp \left(\eta \sum_{h=1}^H \|\widehat{\ell}_k^h\|_1 \right) \right]$$

by triangle inequality. Therefore, $p_{k+1}(\pi)/p_k(\pi)$ lies in this interval as well, which is just our claim. \square

Lemma 12 (Bounding Stability Term). *The stability term is bounded by*

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \sum_{\pi \in \Pi} (p_k(\pi) - p_{k+1}(\pi)) \langle \mu_{\pi}^h, \widehat{\ell}_k^h \rangle \right] \leq 3\eta H^2 SAK.$$

Proof. By summing up Lemma 11 for all $\pi \in \Pi$ and using the fact that $1 - \exp(-x) \leq x$, we have

$$\sum_{\pi \in \Pi} (p_k(\pi) - p_{k+1}(\pi)) \sum_{h=1}^H \langle \mu_\pi^h, \widehat{\ell}_k^h \rangle \leq \eta \sum_{h'=1}^H \|\widehat{\ell}_k^{h'}\|_1 \cdot \sum_{\pi \in \Pi} p_k(\pi) \sum_{h=1}^H \langle \mu_\pi^h, \widehat{\ell}_k^h \rangle, \quad \forall k \in [K]. \quad (7)$$

To proceed, we need to investigate the Geometric Re-sampling process. Consider the random variable M_k^h whose value is determined in the last line of Algorithm 1. One may view it as a ‘‘truncated’’ geometric random variable, where $\text{Geo}(q)$ is a geometric random variable with parameter q , i.e., $\Pr\{\text{Geo}(q) = n\} = (1 - q)^{n-1}q$. Formally, we have:

$$M_k^h = \min\{\text{Geo}(q_k^h(s_k^h, a_k^h)), L\}, \quad \text{where } q_k^h(s, a) = \mathbb{E}_{\pi \sim p_k} [\mu_\pi^h(s, a)] = \sum_{\pi \in \Pi} p_k(\pi) \mu_\pi^h(s, a). \quad (8)$$

So if we calculate the expectation of $\widehat{\ell}_k^h(s, a)$ only with respect to M_k^h , we will have

$$\mathbb{E} \left[\widehat{\ell}_k^h(s, a) \mid (s_k^h, a_k^h) = (s, a) \right] \leq \frac{\ell_k^h(s, a)}{q_k^h(s, a)}.$$

Let $\mathbb{1}_k^h(s, a)$ be the shorthand notation of $\mathbb{1}[(s_k^h, a_k^h) = (s, a)]$. Then for those $h' \neq h$ in the RHS of Eq. (7), we have

$$\begin{aligned} & \eta \mathbb{E} \left[\sum_{h=1}^H \sum_{s, a} \sum_{\pi \in \Pi} p_k(\pi) \mu_\pi^h(s, a) \widehat{\ell}_k^h(s, a) \sum_{h' \neq h} \|\widehat{\ell}_k^{h'}\|_1 \mid \mathcal{F}_{k-1} \right] \\ & \stackrel{(a)}{\leq} \eta \mathbb{E} \left[\sum_{h=1}^H \sum_{s, a} \mathbb{1}_k^h(s, a) \ell_k^h(s, a) \frac{\sum_{\pi \in \Pi} p_k(\pi) \mu_\pi^h(s, a)}{q_k^h(s, a)} \sum_{h' \neq h} \|\widehat{\ell}_k^{h'}\|_1 \mid \mathcal{F}_{k-1} \right] \\ & \stackrel{(b)}{\leq} \eta H \mathbb{E} \left[\sum_{h' \neq h} \|\widehat{\ell}_k^{h'}\|_1 \mid \mathcal{F}_{k-1} \right] \stackrel{(c)}{\leq} \eta H^2 SA. \end{aligned}$$

where (a) is taking expectation w.r.t. M_k^h , (b) used the definition of q_k^h together with the fact that $\sum_{(s, a)} \mathbb{1}_k^h(s, a) = 1$, and (c) used the fact that $\mathbb{E}[\widehat{\ell}_k^{h'}(s', a') \mid \mathcal{F}_{k-1}] \leq \ell_k^{h'}(s', a') \leq 1$ (Lemma 38).

For those terms with $h = h'$ in Eq. (7), by direct calculation and the fact that $\widehat{\ell}_k^h$ is a one-hot vector, we can write them as

$$\eta \mathbb{E} \left[\sum_{h=1}^H \sum_{s, a} \sum_{\pi \in \Pi} p_k(\pi) \mu_\pi^h(s, a) \left(\widehat{\ell}_k^h(s, a) \right)^2 \mid \mathcal{F}_{k-1} \right] \leq 2\eta \mathbb{E} \left[\sum_{h, s, a} \frac{q_k^h(s, a)}{q_k^h(s, a)} \mid \mathcal{F}_{k-1} \right] \leq 2\eta HSA$$

where we use $\mathbb{E}[(\widehat{\ell}_k^h(s, a))^2 \mid \mathcal{F}_{k-1}] \leq 2(q_k^h(s, a))^{-1}$ (Lemma 39). Combining the terms with $h' \neq h$ and the ones with $h' = h$ gives our conclusion. \square

B.1.4 Proof of Theorem 4

Proof of Theorem 4. From Lemmas 9, 10 and 12, we have

$$\mathcal{R}_K \leq \frac{SAHK}{eL} + \frac{2H}{\eta}(1 + \ln(SA)) + 3\eta H^2 SAK.$$

Therefore, if we pick $\eta^{-1} = \sqrt{HSAK}$ and $L = \sqrt{SAK/H}$,

$$\mathcal{R}_K \leq \frac{H^{3/2}\sqrt{SAK}}{e} + 2H^{3/2}\sqrt{SAK}(1 + \ln(SA)) + 3H^{3/2}\sqrt{SAK} = \tilde{O}\left(H^{3/2}\sqrt{SAK}\right),$$

as desired. \square

Algorithm 3 FTPL for Episodic AMDPs with Bandit Feedback and Unknown Transition

Require: Laplace distribution parameter η . Geometric Re-sampling parameter L .

- 1: Initialize $\mathcal{P}_1 \leftarrow (\Delta(\mathcal{S}))^{[H] \times \mathcal{S} \times \mathcal{A}}$ (the set of all possible transition functions).
 - 2: Sample perturbation $\widehat{\ell}_0 = z$ such that $z^h(s, a)$ is an independent sample of $\text{Laplace}(\eta)$.
 - 3: **for** $k = 1, 2, \dots, K$ **do**
 - 4: Let $(\pi_k, P_k) = \operatorname{argmin}_{(\pi, P) \in \Pi \times \mathcal{P}_k} V(\pi; \widehat{\ell}_{1:k-1} + z, P)$ by Extended Value Iteration (Jaksch et al., 2010). (See also Remark 16 for more details.)
 - 5: **for** $h = 1, 2, \dots, H$ **do**
 - 6: Observe s_k^h , play $a_k^h = \pi_k(s_k^h)$, suffer and observe loss $\ell_k^h(s_k^h, a_k^h)$.
 - 7: **for** $M_k^h = 1, 2, \dots, L$ **do**
 - 8: Sample a fresh perturbation \tilde{z} in the same way as z .
 - 9: Calculate $(\pi'_k, P'_k) = \operatorname{argmin}_{(\pi, P) \in \Pi \times \mathcal{P}_k} V(\pi; \widehat{\ell}_{1:k-1} + \tilde{z}, P)$.
 - 10: Pick the transition $\widehat{P}'_k \in \mathcal{P}_k$ such that $\mu_{\pi'}^h(s_k^h, a_k^h; \widehat{P}'_k)$ is maximized via the COMP-
UOB procedure proposed by Jin et al. (2020).
 - 11: Simulate π'_k for h steps starting from s^1 and following transitions $(\widehat{P}'_k)^1, \dots, (\widehat{P}'_k)^h$.
 - 12: **if** (s_k^h, a_k^h) is visited at step h or $M_k^h = L$ **then**
 - 13: Set $\widehat{\ell}_k^h(s, a) = M_k^h \cdot \ell_k^h(s_k^h, a_k^h) \cdot \mathbb{1}[(s_k^h, a_k^h) = (s, a)]$ and break.
 - 14: Calculate \mathcal{P}_{k+1} according to Eq. (10).
-

B.1.5 Comparison with the CONTEXT-FTPL algorithm

One may think that our algorithm together with its analysis looks quite similar to the CONTEXT-FTPL algorithm (Syrgkanis et al., 2016, Algorithm 2) for adversarial contextual bandits. In fact, we can even convert the episodic AMDP problem with known transition as an instance of their contextual semi-bandit problem: for time slot (k, h) , the “context” is h and the loss vector is $\widehat{\ell}_k^h$. A policy π under context $x = h$ will then give an “action” $\pi(h) = \mu_{\pi}^h$ (the occupancy measure), which means it will suffer loss $\langle \mu_{\pi}^h, \widehat{\ell}_k^h \rangle$. Both algorithms add perturbations to each of the contexts, $1, 2, \dots, H$, denoted by $z^1, z^2, \dots, z^H \in \mathbb{R}^{\mathcal{S}^A}$ respectively.

However, there is a main differences between our setting and theirs: in their setting, the action space (where $\pi(x)$ belongs) is binary. However, in our case, $\mu_{\pi}^h \in [0, 1]^{\mathcal{S}^A}$ is continuous. Though this difference may look tiny, it actually induces extra difficulties: this subtle difference will make their Lemma 10, stated as follows, no longer hold.

Lemma 13 (Syrgkanis et al. (2016, Lemma 10)). *For any contexts x^1, x^2, \dots, x^T and non-negative linear loss functions $\ell^1, \ell^2, \dots, \ell^T$, suppose that $z^h(s, a) \sim \text{Laplace}(\eta)$, CONTEXT-FTPL satisfies*

$$\mathbb{E}_z [\langle \pi^t(x^t), \ell^t \rangle - \langle \pi^{t+1}(x^t), \ell^t \rangle] \leq \eta \cdot \mathbb{E}[\langle \pi^t(x^t), \ell^t \rangle^2], \quad \forall 1 \leq t < T. \quad (9)$$

To see this, consider the simple case that there is only one possible value of the context together with two policies, each associated with action vectors $(0.1, 0.1, 0.2)$ and $(0.2, 0.1, 0.1)$, denoted by π_1 and π_2 , respectively. Set the cumulative (perturbed) loss vector $\ell_{0:t-1}$ as $(0.75, 0.2, 0.6)$ and $\ell_t = (0.1, 0, 0)$ (this is set to be one-hot, so it can be yielded from our Geometric Re-sampling process). Set the Laplace distribution parameter $\eta = 3$. Then, by direct calculation via integration, $p_t(\pi_1) = 0.609453$ and $p_{t+1}(\pi_1) = 0.675248$. As $\langle \pi_1, \ell^t \rangle = 0.01$ and $\langle \pi_2, \ell^t \rangle = 0.02$, the LHS of the Eq. (9) will be 0.00065795 while the RHS will be 0.000651492. Therefore, Eq. (9) simply *does not* hold, even if there are only 2 policies, 3 dimensions and 1 context.

Fortunately, as explained in the main text, though this strong version of “single-step stability lemma” does not hold, we are still able to prove a weaker version, Lemma 3 (which is restated as Lemma 11 in the appendix), to bound the stability term, which is worse only by a factor H , instead of $\|\widehat{\ell}_t\|_{\infty} \leq L$.

B.2 Unknown Transition Case (Theorem 5)

We first present our algorithm for the unknown transition case in Algorithm 3.

B.2.1 Transitions' Confidence Set Construction

We first discuss our construction of transitions' confidence sets. As in [Jin et al. \(2022\)](#), we maintain a confidence set of transitions \mathcal{P}_k for each episode $k \in [K]$ as Eq. (10), where $\mathcal{P}_1 = (\Delta(\mathcal{S}))^{[H] \times \mathcal{S} \times \mathcal{A}}$.

As mentioned in the main text, we also want to ensure that $\mathcal{P}_{k+1} \subseteq \mathcal{P}_k$. Instead of taking $\mathcal{P}_1 \cap \mathcal{P}_2 \cap \dots \cap \mathcal{P}_k$ when doing the optimization, we directly ensure $\mathcal{P}_{k+1} \subseteq \mathcal{P}_k$ when constructing the confidence sets, such that they are always *shrinking*. This is to ensure a well-bounded error term, as we will illustrate in Lemma 18.

$$\mathcal{P}_{k+1} = \mathcal{P}_k \cap \left\{ \widehat{P}: [H] \times \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S}) \mid \left| \widehat{P}^h(s' | s, a) - \overline{P}_k^h(s' | s, a) \right| \leq \varepsilon_i^h(s' | s, a), \forall s, s' \in \mathcal{S}, a \in \mathcal{A} \right\}, \quad (10)$$

$$\text{where } \varepsilon_k^h(s' | s, a) = 4 \sqrt{\frac{\overline{P}_k^h(s' | s, a) \ln(10HSAK/\delta)}{\max\{1, N_k^h(s, a)\}}} + 10 \frac{\ln(10HSAK/\delta)}{\max\{1, N_k^h(s, a)\}}, \quad (11)$$

$$\text{and } \overline{P}_k^h(s' | s, a) = \frac{N_k^h(s' | s, a)}{N_k^h(s, a)},$$

$$N_k^h(s, a) = \sum_{k'=1}^k \mathbb{1}[(s_{k'}^h, a_{k'}^h) = (s, a)], N_k^h(s' | s, a) = \sum_{k'=1}^k \mathbb{1}[s_{k'}^{h+1} = s', (s_{k'}^h, a_{k'}^h) = (s, a)].$$

By the following lemma from [Jin et al. \(2020\)](#), we define K good events, $\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_K$, where \mathcal{E}_k means $\mathbb{P} \in \mathcal{P}_k$. From the following lemma, we can conclude that $\Pr\{\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_K\} \geq 1 - 4\delta$. For simplicity, we also denote $\mathcal{E} = \mathcal{E}_1 \wedge \mathcal{E}_2 \wedge \dots \wedge \mathcal{E}_K$. Hence, $\Pr\{\mathcal{E}\} \geq 1 - 4\delta$ (in fact, we have $\mathcal{E} = \mathcal{E}_K$ as $\mathcal{P}_k \subseteq \mathcal{P}_{k-1}$).

Lemma 14 ([\(Jin et al., 2020, Lemma 2\)](#)). *With probability $1 - 4\delta$, we have $\mathbb{P} \in \mathcal{P}_k$ for all $k \in [K]$.*

Remark 15. *Note that the original definition is slightly different from ours, where there is no intersection operations taken with previous confidence sets. However, as long as \mathbb{P} belongs to all the confidence sets, it clearly belongs to the intersection of them.*

Remark 16. *Note that the Extended Value Iteration ([Jaksch et al., 2010](#)) approach works as long as \mathcal{P}_k has the form $\{P \mid P^h(s' | s, a) \in [L^h(s' | s, a), R^h(s' | s, a)]\}$, but does not require $[L^h(s' | s, a), R^h(s' | s, a)]$ to be centered exactly at $\overline{P}_k^h(s' | s, a)$ (which is indeed the case for our algorithm due to the intersection operations).*

B.2.2 Regret Decomposition

For the unknown-transition cases, we first do the following regret decomposition as [Jin et al. \(2020\)](#):

$$\begin{aligned} \mathcal{R}_K = & \underbrace{\mathbb{E} \left[\sum_{k=1}^K (V(\pi_k; \ell_k, \mathbb{P}) - V(\pi_k; \ell_k, P_k)) \right]}_{\text{ERROR}} + \underbrace{\mathbb{E} \left[\sum_{k=1}^K (V(\pi_k; \ell_k, P_k) - V(\pi_k; \widehat{\ell}_k, P_k)) \right]}_{\text{BIAS1}} + \\ & \underbrace{\mathbb{E} \left[\sum_{k=1}^K (V(\pi_k; \widehat{\ell}_k, P_k) - V(\pi^*; \widehat{\ell}_k, \mathbb{P})) \right]}_{\text{ESTREG}} + \underbrace{\mathbb{E} \left[\sum_{k=1}^K (V(\pi^*; \widehat{\ell}_k, \mathbb{P}) - V(\pi^*; \ell_k, \mathbb{P})) \right]}_{\text{BIAS2}}. \end{aligned}$$

Intuitively, the ERROR term is due to the transition estimation, BIAS1 and BIAS2 terms are due to loss estimation for π_k and π^* , respectively, and ESTREG is the regret of our FTPL algorithm on the estimated transitions P_k and the estimated losses $\widehat{\ell}_k$.

B.2.3 Bounding the ESTREG Term

Theorem 17 (Bounding ESTREG Term). *The ESTREG term is bounded by*

$$\text{ESTREG} = \mathbb{E} \left[\sum_{k=1}^K (V(\pi_k; \widehat{\ell}_k, P_k) - V(\pi^*; \widehat{\ell}_k, \mathbb{P})) \right] \leq \frac{2H}{\eta} (1 + \ln(SA)) + 3\eta H^2 SAK + 8\delta KHL.$$

Proof. For the ESTREG term, we will also decompose it into an error term (not to be confused with the ERROR term which occurs in the decomposition of \mathcal{R}_K ; this error term appears in the decomposition of ESTREG and is related to the ‘be-the-leader’ lemma) and a stability term (as it is defined for the estimated losses, there is no GR error term anymore). However, here we should define our “leader” as

$$(\tilde{\pi}_{k+1}, \tilde{P}_{k+1}) = \operatorname{argmin}_{(\pi, P) \in \Pi \times \mathcal{P}_k} V(\pi; \hat{\ell}_{0:k}, P).$$

Instead of directly using (π_{k+1}, P_{k+1}) as the leader (as we did in the known transition case), we allow the transition \tilde{P}_{k+1} selected from $\mathcal{P}_k \supseteq \mathcal{P}_{k+1}$. This is critical to ensure a low stability term, as we can only derive the “single-step stability lemma” (Lemma 19 in this case) for two probability distributions sharing a same support ($\Pi \times \mathcal{P}_k$ here).

As an analog to the known transition case, we define $p_k(\pi, P)$ as the probability density function (with respect of the perturbation z) of (π_k, P_k) conditioning on $\hat{\ell}_1, \hat{\ell}_2, \dots, \hat{\ell}_{k-1}$. Note that as there are infinitely many transitions, we cannot directly write \Pr_z as in the known-transition setting.

Moreover, as explained before, we allow \tilde{P}_{k+1} to be picked from \mathcal{P}_k instead of \mathcal{P}_{k+1} , so the probability of picking $(\tilde{\pi}_{k+1}, \tilde{P}_{k+1})$ as (π, P) is not simply $p_{k+1}(\pi, P)$. Therefore, we have to define another notation representing the probability density of picking each (π, P) as $(\tilde{\pi}_{k+1}, \tilde{P}_{k+1})$, namely $\tilde{p}_{k+1}(\pi, P)$, which is the probability density of $(\tilde{\pi}_{k+1}, \tilde{P}_{k+1})$ with respect to z , conditioning on $\hat{\ell}_1, \hat{\ell}_2, \dots, \hat{\ell}_k$.

Hence, we can write

$$\begin{aligned} \text{ESTREG} &= \mathbb{E} \left[\underbrace{\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\tilde{\pi}_{k+1}}^h(\tilde{P}_{k+1}) - \mu_{\pi^*}^h(\mathbb{P}), \hat{\ell}_k^h \rangle}_{\text{Error term}} \right] + \\ &\quad \underbrace{\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} (p_k(\pi, P) - \tilde{p}_{k+1}(\pi, P)) \langle \mu_{\pi}^h(P), \hat{\ell}_k^h \rangle dP \right]}_{\text{Stability Term}}. \end{aligned}$$

For the error term, we only need to verify that the “be-the-leader argument” that we used in Lemma 10 still holds. Fortunately, it turns out as long as $\mathcal{P}_1 \supseteq \mathcal{P}_2 \supseteq \dots \supseteq \mathcal{P}_K \supseteq \{\mathbb{P}\}$, we can always conclude the following lemma, whose proof is presented later.

Lemma 18 (Bounding Error Term). *The error term in this case is bounded by*

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\tilde{\pi}_{k+1}}^h(\tilde{P}_{k+1}) - \mu_{\pi^*}^h(\mathbb{P}), \hat{\ell}_k^h \rangle \right] \leq \frac{2H}{\eta} (1 + \ln(SA)) + 4\delta KHL.$$

For the stability term, we need a similar but different single-step stability bound, as

Lemma 19 (Single Step Stability). *For all $k \in [K]$, $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $(\pi, P) \in \Pi \times \mathcal{P}_k$,*

$$\tilde{p}_{k+1}(\pi, P) \geq p_k(\pi, P) \exp \left(-\eta \sum_{h=1}^H \|\hat{\ell}_k^h\|_1 \right).$$

With this lemma, our derivation for the stability term in known-transition cases (Lemma 12) also holds, except that we are using the upper occupancy measures in the Geometric Re-sampling process, instead of the actual occupancy measures. Technically, this means that the event $(s_k^h, a_k^h) = (s, a)$ will happen with a probability

$$\hat{q}_k^h(s, a) = \Pr\{(s_k^h, a_k^h) = (s, a) \mid \mathcal{F}_{k-1}\} = \sum_{\pi \in \Pi} p_k(\pi) \mu_{\pi}^h(s, a; \mathbb{P}), \quad (12)$$

where $p_k(\pi) = \int_{\mathcal{P}_k} p_k(\pi, P) dP$ is the marginal probability of picking π for episode k (with a slight abuse of notation). However, in each execution of the Geometric Re-sampling process, the probability of visiting (s, a) is another probability

$$q_k^h(s, a) = \sum_{\pi \in \Pi} p_k(\pi) \max_{P' \in \mathcal{P}_k} \mu_\pi^h(s, a; P') \neq \hat{q}_k^h(s, a), \quad (13)$$

which means we cannot use Lemmas 38 and 39 anymore.

Fortunately, we are able to derive Corollaries 40 and 41 in such a case, which actually implies the previous two lemmas, given that the actual occupancy measure $\hat{q}_k^h(s, a)$ is bounded by the upper occupancy measure $q_k^h(s, a)$ (which is indeed this case as long as $\mathbb{P} \in \mathcal{P}_k$, i.e., event \mathcal{E}_k holds). However, for the BIAS1 term (Theorem 22), as we will see later, this inconsistency will indeed induce extra difficulties, leading to a $\tilde{O}(H^2 S \sqrt{AK})$ dominating term as in Jin et al. (2020).

The detailed proof of Lemma 20 will be presented after the proof of this theorem.

Lemma 20 (Bounding Stability Term). *The stability term in this case is bounded by*

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} (p_k(\pi, P) - \tilde{p}_{k+1}(\pi, P)) \langle \mu_\pi^h(P), \hat{\ell}_k^h \rangle dP \right] \leq 3\eta H^2 SAK + 4\delta KHL.$$

Combining them together gives

$$\text{ESTREG} \leq \frac{2H}{\eta} (1 + \ln(SA)) + 3\eta H^2 SAK + 8\delta KHL,$$

as claimed. \square

Proof of Lemma 18. The proof still follows the idea of Lemma 10. We rewrite the error term as

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{k=1}^K V(\tilde{\pi}_{k+1}; \hat{\ell}_k, \tilde{P}_{k+1}) - \sum_{k=1}^K V(\pi^*; \hat{\ell}_k, \mathbb{P}) \right) \mathbb{1}[\mathcal{E}] \right] + \\ & \mathbb{E} \left[\left(\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\tilde{\pi}_{k+1}}^h(\tilde{P}_{k+1}) - \mu_{\pi^*}^h(\mathbb{P}), \hat{\ell}_k^h \rangle \right) \mathbb{1}[-\mathcal{E}] \right]. \end{aligned}$$

For the second term, since by definition $0 \leq \hat{\ell}_k^h(s, a) \leq L$ for all k, h, s, a and both $\mu_{\tilde{\pi}_{k+1}}^h(\tilde{P}_{k+1})$ and $\mu_{\pi^*}^h(\mathbb{P})$ are probability distributions, we can bound it as

$$KHL \cdot \Pr\{-\mathcal{E}\} \leq 4\delta KHL.$$

Now consider the summation inside the first expectation. If we add an extra term $V(\tilde{\pi}_1; \hat{\ell}_0, \tilde{P}_1) - V(\pi^*; \hat{\ell}_0, \mathbb{P})$ where $\hat{\ell}_0 = z$ is the perturbation. The following deduction holds under the event \mathcal{E} :

$$\begin{aligned} & \sum_{k=0}^K V(\tilde{\pi}_{k+1}; \hat{\ell}_k, \tilde{P}_{k+1}) - V(\pi^*; \hat{\ell}_{0:K}, \mathbb{P}) \\ & \stackrel{(a)}{\leq} \sum_{k=0}^K V(\tilde{\pi}_{k+1}; \hat{\ell}_k, \tilde{P}_{k+1}) - V(\tilde{\pi}_{K+1}; \hat{\ell}_{0:K}, \tilde{P}_{K+1}) = \sum_{k=0}^{K-1} V(\tilde{\pi}_{k+1}; \hat{\ell}_k, \tilde{P}_{k+1}) - V(\tilde{\pi}_{K+1}; \hat{\ell}_{0:K-1}, \tilde{P}_{K+1}) \\ & \stackrel{(b)}{\leq} \sum_{k=0}^{K-1} V(\tilde{\pi}_{k+1}; \hat{\ell}_k, \tilde{P}_{k+1}) - V(\tilde{\pi}_K; \hat{\ell}_{0:K-1}, \tilde{P}_K) = \sum_{k=0}^{K-2} V(\tilde{\pi}_{k+1}; \hat{\ell}_k, \tilde{P}_{k+1}) - V(\tilde{\pi}_K; \hat{\ell}_{0:K-2}, \tilde{P}_K) \\ & \leq \dots \leq V(\tilde{\pi}_1; \hat{\ell}_0, \tilde{P}_1) - V(\tilde{\pi}_2; \hat{\ell}_0, \tilde{P}_2) \stackrel{(c)}{\leq} 0. \end{aligned}$$

Here, (a) used the optimality of $(\tilde{\pi}_{K+1}, \tilde{P}_{K+1})$ over the set $\Pi \times \mathcal{P}_K$ w.r.t. losses $\hat{\ell}_{0:K}$, which is valid due to \mathcal{E} ; (b) used the optimality of $(\tilde{\pi}_K, \tilde{P}_K)$ over the set $\Pi \times \mathcal{P}_{K-1}$ w.r.t. losses $\hat{\ell}_{0:K-1}$, which is

again valid since $\tilde{P}_{K+1} \in \mathcal{P}_K \subseteq \mathcal{P}_{K-1}$; similarly (c) used the optimality of $(\tilde{\pi}_1, \tilde{P}_1)$ over $\Pi \times \mathcal{P}_0$, which again holds as $\tilde{P}_2 \in \mathcal{P}_0$ (which is the set of all transitions). So we still have the following inequality as Lemma 10:

$$\mathbb{E} \left[\left(\sum_{k=1}^K V(\pi_{k+1}; \hat{\ell}_k, \tilde{P}_{k+1}) - \sum_{k=1}^K V(\pi^*; \hat{\ell}_k, \mathbb{P}) \right) \mathbb{1}[\mathcal{E}] \right] \leq \mathbb{E}[V(\pi^*; \hat{\ell}_0, \mathbb{P}) - V(\tilde{\pi}_1; \hat{\ell}_0, \tilde{P}_1)].$$

By the notation of occupancy measures, we can rewrite the last term as

$$\mathbb{E} \left[\sum_{h=1}^H \langle \mu_{\pi^*}^h(\mathbb{P}), \hat{\ell}_0^h \rangle - \sum_{h=1}^H \langle \mu_{\tilde{\pi}_1}^h(\tilde{P}_1), \hat{\ell}_0^h \rangle \right] \leq 2 \sum_{h=1}^H \mathbb{E}[\|\hat{\ell}_0^h\|_\infty],$$

which is again bounded by $\frac{2H}{\eta}(1 + \ln(SA))$ due to Lemma 44. Combining these two parts (with or without \mathcal{E}) together gives our conclusion. \square

Proof of Lemma 19. We follow the proof of Lemma 11. For a fixed episode $k \in [K]$, we consider any $(\pi, P) \in \Pi \times \mathcal{P}_k$. We use the notation $(\pi, P) = \text{best}(\ell; \mathcal{P})$ to denote $(\pi, P) = \text{argmin}_{(\pi, P) \in \Pi \times \mathcal{P}} V(\pi; \ell, P)$. Then we have

$$\begin{aligned} p_k(\pi, P) &= \int_z \mathbb{1} \left[(\pi, P) = \text{best} \left(\hat{\ell}_{1:k-1} + z; \mathcal{P}_k \right) \right] f(z) \, dz \\ &= \int_z \mathbb{1} \left[(\pi, P) = \text{best} \left(\hat{\ell}_{1:k-1} + (z + \hat{\ell}_k); \mathcal{P}_k \right) \right] f(z + \hat{\ell}_k) \, dz \\ &= \int_z \mathbb{1} \left[(\pi, P) = \text{best} \left(\hat{\ell}_{1:k} + z; \mathcal{P}_k \right) \right] f(z + \hat{\ell}_k) \, dz, \end{aligned}$$

where $f(z)$ is the probability density function of z and the second step made use of the fact that $z + \hat{\ell}_k$ is still linear in z . Moreover,

$$\tilde{p}_{k+1}(\pi, P) = \int_z \mathbb{1} \left[(\pi, P) = \text{best} \left(\hat{\ell}_{1:k} + z; \mathcal{P}_k \right) \right] f(z) \, dz.$$

Again by the fact that $f(z) = \prod_{h=1}^H \exp(-\eta \|z^h\|_1)$, which we used in the proof of Lemma 11, we have

$$f(z + \hat{\ell}_k) = \prod_{h=1}^H \exp \left(-\eta \left(\|z^h + \hat{\ell}_k\|_1 - \|z^h\|_1 \right) \right) f(z),$$

which gives

$$\frac{f(z + \hat{\ell}_k)}{f(z)} \in \left[\exp \left(-\eta \sum_{h=1}^H \|\hat{\ell}_k^h\|_1 \right), \exp \left(\eta \sum_{h=1}^H \|\hat{\ell}_k^h\|_1 \right) \right]$$

by triangle inequality. Therefore, $\tilde{p}_{k+1}(\pi, P)/p_k(\pi, P)$ lies in this interval as well, which is just our claim. \square

Proof of Lemma 20. Let us focus on a single episode, say $k \in [K]$. We should first make sure that $\hat{q}_k^h(s, a) \leq q_k^h(s, a)$ (defined in Equations (12) and (13)), which happens when $\mathbb{P} \in \mathcal{P}_k$, i.e., \mathcal{E}_k holds. Therefore, we rewrite the k -th summand of the stability term as

$$\begin{aligned} &\mathbb{E} \left[\sum_{h=1}^H \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} (p_k(\pi, P) - \tilde{p}_{k+1}(\pi, P)) \langle \mu_\pi^h(P), \hat{\ell}_k^h \rangle \, dP \mathbb{1}[\mathcal{E}_k] \right] + \\ &\mathbb{E} \left[\sum_{h=1}^H \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} (p_k(\pi, P) - \tilde{p}_{k+1}(\pi, P)) \langle \mu_\pi^h(P), \hat{\ell}_k^h \rangle \, dP \mathbb{1}[-\mathcal{E}_k] \right]. \end{aligned} \quad (14)$$

For the second term, we will bound it trivially as $HL \Pr\{-\mathcal{E}_k\} \leq 4\delta HL$ as $p_k, \tilde{p}_{k+1} \in \Delta(\Pi \times \mathcal{P}_k)$ and $\mu_\pi^h \in \Delta(\mathcal{S} \times \mathcal{A})$. For the first term, we will do something similar to Lemma 12, as follows:

Summing up Lemma 19 for all $(\pi, P) \in \Pi \times \mathcal{P}_k$ and using the fact that $\exp(-x) \geq (1-x)$ gives

$$\begin{aligned} & \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} (p_k(\pi, P) - \tilde{p}_{k+1}(\pi, P)) \sum_{h=1}^H \langle \mu_\pi^h(P), \widehat{\ell}_k^h \rangle dP \\ & \leq \eta \sum_{h'=1}^H \|\widehat{\ell}_k^{h'}\|_1 \cdot \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \sum_{h=1}^H \langle \mu_\pi^h(P), \widehat{\ell}_k^h \rangle dP. \end{aligned} \quad (15)$$

By considering the randomness of M_k^h , we will still have the following property, except for a different definition of q_k^h :

$$\mathbb{E} \left[\widehat{\ell}_k^h(s, a) \mid (s_k^h, a_k^h) = (s, a) \right] \leq \frac{\ell_k^h(s, a)}{q_k^h(s, a)}, \quad \text{where } q_k^h(s, a) = \sum_{\pi \in \Pi} p_k(\pi) \max_{P' \in \mathcal{P}_k} \mu_\pi^h(s, a; P'), \quad (16)$$

as when doing the Geometric Re-sampling process, we are picking the transition in \mathcal{P}_k that maximizes the probability of reaching (s, a) . Still use $\mathbb{1}_k^h(s, a)$ as the shorthand notation of $\mathbb{1}[(s_k^h, a_k^h) = (s, a)]$. Then for any history \mathcal{F}_{k-1} and those $h' \neq h$ in Eq. (15),

$$\begin{aligned} & \eta \mathbb{E} \left[\sum_{h=1}^H \sum_{s,a} \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \mu_\pi^h(s, a; P) \widehat{\ell}_k^h(s, a) dP \cdot \sum_{h' \neq h} \|\widehat{\ell}_k^{h'}\|_1 \mathbb{1}[\mathcal{E}_k] \mid \mathcal{F}_{k-1} \right] \\ & \stackrel{(a)}{\leq} \eta \mathbb{E} \left[\sum_{h=1}^H \sum_{s,a} \mathbb{1}_k^h(s, a) \ell_k^h(s, a) \frac{\sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \mu_\pi^h(s, a; P) dP}{q_k^h(s, a)} \sum_{h' \neq h} \|\widehat{\ell}_k^{h'}\|_1 \mathbb{1}[\mathcal{E}_k] \mid \mathcal{F}_{k-1} \right] \\ & \stackrel{(b)}{\leq} \eta H \mathbb{E} \left[\sum_{h' \neq h} \|\widehat{\ell}_k^{h'}\|_1 \mathbb{1}[\mathcal{E}_k] \mid \mathcal{F}_{k-1} \right] \stackrel{(c)}{\leq} \eta H \sum_{s,a} \mathbb{E} \left[\sum_{h=1}^H \frac{\widehat{q}_k^h(s, a)}{q_k^h(s, a)} \mathbb{1}[\mathcal{E}_k] \mid \mathcal{F}_{k-1} \right] \stackrel{(d)}{\leq} \eta H^2 SA, \end{aligned}$$

where (a) is taking expectation w.r.t. M_k^h , (b) used the (new) definition of q_k^h together with the fact that $\sum_{(s,a)} \mathbb{1}_k^h(s, a) = 1$, (c) used Corollary 40 and (d) used $\widehat{q}_k^h(s, a) \leq q_k^h(s, a)$ (which is due to $\mathbb{1}[\mathcal{E}_k]$).

For those terms with $h' = h$ in Eq. (15), by direct calculation and the fact that $\widehat{\ell}_k^h$ is a one-hot vector, we can write them as

$$\begin{aligned} & \eta \mathbb{E} \left[\sum_{h=1}^H \sum_{s,a} \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \mu_\pi^h(s, a; P) \left(\widehat{\ell}_k^h(s, a) \right)^2 dP \mathbb{1}[\mathcal{E}_k] \mid \mathcal{F}_{k-1} \right] \\ & \leq 2\eta \mathbb{E} \left[\sum_{h,s,a} \frac{\sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \mu_\pi^h(s, a; P) dP}{q_k^h(s, a)} \frac{\widehat{q}_k^h(s, a)}{q_k^h(s, a)} \mathbb{1}[\mathcal{E}_k] \mid \mathcal{F}_{k-1} \right] \\ & \leq 2\eta \sum_{s,a} \mathbb{E} \left[\sum_{h=1}^H \frac{\widehat{q}_k^h(s, a)}{q_k^h(s, a)} \mathbb{1}[\mathcal{E}_k] \mid \mathcal{F}_{k-1} \right] \leq 2\eta H SA, \end{aligned}$$

by applying Corollary 41 together with the fact that $\widehat{q}_k^h(s, a) \leq q_k^h(s, a)$ when \mathcal{E}_k happens. Combining the terms with $h' \neq h$ and the ones with $h' = h$ gives

$$\text{Eq. (14)} \leq 3\eta H^2 SA + 4\delta HL.$$

Therefore, the stability term is bounded by $3\eta H^2 SAK + 4\delta KHL$, as claimed. \square

B.2.4 Bounding Other Terms

The terms other than ESTREG can be bounded similarly to Jin et al. (2020), as follows:

Theorem 21 (Bounding ERROR Term). *The ERROR term is bounded by*

$$\text{ERROR} = \mathbb{E} \left[\sum_{k=1}^K (V(\pi_k; \ell_k, \mathbb{P}) - V(\pi_k; \ell_k, P_k)) \right] \leq \tilde{\mathcal{O}} \left(H^2 S \sqrt{AK} + \delta KH \right).$$

Theorem 22 (Bounding BIAS1 Term). *The BIAS term is bounded by*

$$\begin{aligned} \text{BIAS1} &= \mathbb{E} \left[\sum_{k=1}^K (V(\pi_k; \ell_k, P_k) - V(\pi_k; \hat{\ell}_k, P_k)) \right] \\ &\leq \tilde{\mathcal{O}} \left(\frac{HKSA}{L} + H^2 S \sqrt{AK} + H^3 S^3 A + \delta KH \right). \end{aligned}$$

Remark 23. *This term looks quite similar to the GR error term (Lemma 9). However, they are in fact different as we will have some extra terms due to the UOB technique. In other words, we are having different probabilities when reaching (s_k^h, a_k^h) and when doing Geometric Re-sampling (c.f. Lemma 38 and corollary 40). Therefore, this term will be further decomposed into two parts, where the first one is due to bias of the GR estimator and the second one is due to the UOB technique and can be bounded similar to Jin et al. (2022, Lemma A.3). Check the proof below for more details.*

Theorem 24 (Bounding BIAS2 Term). *The BIAS2 term is bounded by*

$$\text{BIAS2} = \mathbb{E} \left[\sum_{k=1}^K (V(\pi^*; \hat{\ell}_k, \mathbb{P}) - V(\pi^*; \ell_k, \mathbb{P})) \right] = \tilde{\mathcal{O}}(\delta KHL).$$

Proof of Theorem 5. By combining Theorems 17, 21, 22 and 24 together, we will have

$$\mathcal{R}_T \leq \tilde{\mathcal{O}} \left(H^2 S \sqrt{AK} + \frac{H}{\eta} + \eta H^2 SAK + \frac{HSAK}{L} + H^2 S \sqrt{AK} + H^3 S^3 A + \delta KHL \right).$$

Picking $\eta = \left(\sqrt{HSAK} \right)^{-1}$, $L = \sqrt{SAK/H}$ and $\delta = 1/K$ gives

$$\mathcal{R}_T \leq \tilde{\mathcal{O}} \left(H^2 S \sqrt{AK} + H^{3/2} \sqrt{SAK} + H \sqrt{K} + H^3 S^3 A \right) = \tilde{\mathcal{O}} \left(H^2 S \sqrt{AK} + H^3 S^3 A \right),$$

which finishes the proof. \square

Proof of Theorem 21. We need the following key lemma from Jin et al. (2020):⁴

Lemma 25 (Jin et al. (2020, Lemma 4)). *Conditioning on \mathcal{E} , for any set of policies $\{\pi_k \in \Pi\}_{k \in [K]}$ and any collection of transitions $\{P_k^{s,h}\}_{s \in \mathcal{S}, h \in [H]}$ such that $P_k^{s,h} \in \mathcal{P}_k$, with probability $1 - 2\delta$,*

$$\sum_{k=1}^K \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} |\mu_{\pi_k}^h(s, a; P_k^{s,h}) - \mu_{\pi_k}^h(s, a; \mathbb{P})| \leq \tilde{\mathcal{O}} \left(H^2 S \sqrt{AK} \right).$$

As all losses are in $[0, 1]$ (note that in the ERROR term we are considering true losses), we have

$$\sum_{k=1}^K (V(\pi_k; \ell_k, \mathbb{P}) - V(\pi_k; \ell_k, P_k)) \leq \sum_{k=1}^K \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} |\mu_{\pi_k}^h(s, a; P_k) - \mu_{\pi_k}^h(s, a; \mathbb{P})|,$$

which is bounded by $\tilde{\mathcal{O}}(H^2 S \sqrt{AK})$ with probability $1 - 2\delta$ by the previous lemma. Let the event (i.e., it is bounded by $\tilde{\mathcal{O}}(H^2 S \sqrt{AK})$) be \mathcal{E}' . Then

$$\Pr\{\mathcal{E}' \wedge \mathcal{E}\} = \Pr\{\mathcal{E}' \mid \mathcal{E}\} \Pr\{\mathcal{E}\} \geq 1 - 6\delta.$$

⁴The original paper has a slightly different notation as they assumed the states to be ‘layered’, i.e., $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2 \cup \dots \cup \mathcal{S}_H$ such that the states in \mathcal{S}_h can only transit to \mathcal{S}_{h+1} , $\forall 1 \leq h < H$. Therefore, their \mathcal{S} should be H times larger than ours. They also used T for our K , L for our H and X for our \mathcal{S} .

Therefore, we write

$$\begin{aligned} \mathbb{E} \left[\sum_{k=1}^K (V(\pi_k; \ell_k, \mathbb{P}) - V(\pi_k; \ell_k, P_k)) \right] &= \mathbb{E} \left[\sum_{k=1}^K (V(\pi_k; \ell_k, \mathbb{P}) - V(\pi_k; \ell_k, P_k)) \mathbb{1}[\mathcal{E} \wedge \mathcal{E}'] \right] + \\ &\quad \mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi_k}^h(\mathbb{P}) - \mu_{\pi_k}^h(P_k), \ell_k^h \rangle \mathbb{1}[-\mathcal{E} \vee -\mathcal{E}'] \right] \\ &= \tilde{\mathcal{O}} \left(H^2 S \sqrt{AK} + \delta KH \right), \end{aligned}$$

where the last step used the fact that $\mu_{\pi_k}^h(\mathbb{P})$ and $\mu_{\pi_k}^h(P_k)$ are both probability distributions and $0 \leq \ell_k^h(s, a) \leq 1$. \square

Proof of Theorem 22. Write our BIAS1 term in terms of occupancy measures:

$$\text{BIAS1} = \mathbb{E} \left[\sum_{k=1}^K \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \sum_{h=1}^H \langle \mu_{\pi}^h(P), \ell_k^h - \hat{\ell}_k^h \rangle dP \right].$$

Consider the k -th summand of it, denoted as BIAS1_k . We decompose it into two parts, depending on whether \mathcal{E}_k holds:

$$\begin{aligned} \text{BIAS1}_k &\leq \mathbb{E} \left[\sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \sum_{h=1}^H \mathbb{E} \left[\langle \mu_{\pi}^h(P), \ell_k^h - \hat{\ell}_k^h \rangle \mid \mathcal{F}_{k-1} \right] dP \mathbb{1}[\mathcal{E}_k] \right] + \\ &\quad \mathbb{E} \left[\sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \sum_{h=1}^H \langle \mu_{\pi}^h(P), \ell_k^h \rangle dP \mathbb{1}[-\mathcal{E}_k] \right] \\ &\triangleq \text{BIAS1}_k^{\mathcal{E}} + \text{BIAS1}_k^{-\mathcal{E}}. \end{aligned}$$

For $\text{BIAS1}_k^{-\mathcal{E}}$, we bound it trivially as $H \Pr\{-\mathcal{E}_k\} \leq 4\delta H$ as $p_k \in \Delta(\pi \times \mathcal{P}_k)$, $\mu_{\pi}^h(P) \in \Delta(\mathcal{S} \times \mathcal{A})$ and $\ell_k^h(s, a) \in [0, 1]$. For $\text{BIAS1}_k^{\mathcal{E}}$, we still adopt the notations of $\hat{q}_k^h(s, a)$ and $q_k^h(s, a)$, which are defined as

$$\begin{aligned} \hat{q}_k^h(s, a) &= \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \mu_{\pi}^h(s, a; \mathbb{P}) dP, \\ q_k^h(s, a) &= \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \max_{P' \in \mathcal{P}_k} \mu_{\pi}^h(s, a; P') dP. \end{aligned}$$

Applying Corollary 40 to $\mathbb{E}[\hat{\ell}_k^h(s, a) \mid \mathcal{F}_{k-1}]$, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$ then gives

$$\text{BIAS1}_k^{\mathcal{E}} = \mathbb{E} \left[\sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \sum_{h=1}^H \left\langle \mu_{\pi}^h(P), \left(1 - \frac{\hat{q}_k^h}{q_k^h} + \frac{\hat{q}_k^h}{q_k^h} (1 - q_k^h)^L \right) \ell_k^h \right\rangle dP \mathbb{1}[\mathcal{E}_k] \right]$$

(every operation for the second term of the inner product is element-wise). As $\mathbb{1}[\mathcal{E}_k]$ implies $\hat{q}_k^h(s, a) \leq q_k^h(s, a)$, we can further bound $\text{BIAS1}_k^{\mathcal{E}}$ as

$$\begin{aligned} \text{BIAS1}_k^{\mathcal{E}} &\leq \mathbb{E} \left[\sum_{h=1}^H \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \mu_{\pi}^h(s, a; P) \frac{q_k^h(s, a) - \hat{q}_k^h(s, a)}{q_k^h(s, a)} \ell_k^h(s, a) dP \mathbb{1}[\mathcal{E}_k] \right] + \\ &\quad \mathbb{E} \left[\sum_{h=1}^H \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \mu_{\pi}^h(s, a; P) (1 - q_k^h(s, a))^L dP \mathbb{1}[\mathcal{E}_k] \right]. \end{aligned}$$

For the second term, we can simply make use of the fact that

$$\mathbb{1}[\mathcal{E}_k] = 1 \implies \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \mu_{\pi}^h(s, a; P) dP \leq q_k^h(s, a) \quad (17)$$

together with the condition that $\ell_k^h(s, a) \in [0, 1]$ and consequently bound it by

$$\mathbb{E} \left[\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} q_k^h(s, a) (1 - q_k^h(s, a))^L \mathbb{1}[\mathcal{E}_k] \right] \stackrel{(c)}{\leq} \frac{HSA}{eL},$$

where (c) used the fact that $q(1 - q)^L \leq qe^{-Lq} \leq \frac{1}{eL}$, just as what we did in Lemma 9. For the first term, with a slight abuse of notations, we still use $p_k(\pi)$ to denote the probability of playing π at episode k , i.e., $p_k(\pi) = \int_{P \in \mathcal{P}_k} p_k(\pi, P) dP$. Then again by Eq. (17), we are actually facing

$$\mathbb{E} \left[\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} (q_k^h(s, a) - \hat{q}_k^h(s, a)) \mathbb{1}[\mathcal{E}_k] \right] \quad (18)$$

$$= \mathbb{E} \left[\sum_{\pi \in \Pi} p_k(\pi) \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left(\max_{P' \in \mathcal{P}_k} \mu_\pi^h(s, a; P') - \mu_\pi^h(s, a; \mathbb{P}) \right) \mathbb{1}[\mathcal{E}_k] \right]. \quad (19)$$

Then we follow the idea of Jin et al. (2022, Lemma A.3). We fix the step $h \in [H]$ and the state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$. Therefore, for each policy $\pi \in \Pi$, we can define $\hat{P}_\pi \in \mathcal{P}_k$ to be transition corresponding to the upper-occupancy bound, i.e., it maximizes $\mu_\pi^h(s, a; P)$ over all transitions $P \in \mathcal{P}_k$. Therefore, with the help of the so-called ‘‘occupancy difference lemma’’ (Jin et al., 2021, Lemma D.3.1), we can write the summand in Eq. (19) corresponding to k, h, s, a as

$$\begin{aligned} q_k^h(s, a) - \hat{q}_k^h(s, a) &= \sum_{\pi \in \Pi} p_k(\pi) \left(\mu_\pi^h(s, a; \hat{P}_{\pi, h, s, a}) - \mu_\pi^h(s, a; \mathbb{P}) \right) \\ &= \sum_{\pi \in \Pi} p_k(\pi) \sum_{h'=0}^{h-1} \sum_{x \in \mathcal{S}, y \in \mathcal{A}, z \in \mathcal{S}} \mu_\pi^{h'}(x, y; \mathbb{P}) \left(\mathbb{P}^{h'}(z | x, y) - \hat{P}_\pi^{h'}(z | x, y) \right) \mu_\pi^{h|h'+1}(s, a | z; \hat{P}_\pi), \end{aligned}$$

where $\mu_\pi^{h|h'+1}(s, a | z; P)$ is the so-called ‘‘conditional occupancy measure’’, which is defined as the conditional probability of reaching the state-action pair (s, a) at step h from state z at step $h' + 1$ with policy π and transition P . By \mathcal{E}_k , we have $\mathbb{P} \in \mathcal{P}_k$. Therefore, by the definition of confidence radii, we can further bound

$$|q_k^h(s, a) - \hat{q}_k^h(s, a)| \leq \sum_{\pi \in \Pi} p_k(\pi) \sum_{h'=0}^{h-1} \sum_{x, y, z} \mu_\pi^{h'}(x, y; \mathbb{P}) \epsilon_k^{h'}(z | x, y) \mu_\pi^{h|h'+1}(s, a | z; \hat{P}_\pi),$$

where ϵ_k^h is defined as in Eq. (11).

Then, we consider the conditional occupancy measure $\mu_\pi^{h|h'+1}$ w.r.t. \hat{P}_π . We can still use occupancy difference lemmas (but now we only consider steps between $h' + 1$ and h) to write its difference with the conditional occupancy measure w.r.t. \mathbb{P} as

$$\begin{aligned} &\mu_\pi^{h|h'+1}(s, a | z; \hat{P}_\pi) - \mu_\pi^{h|h'+1}(s, a | z; \mathbb{P}) \\ &\leq \sum_{h''=h'+1}^{h-1} \sum_{u \in \mathcal{S}, v \in \mathcal{A}, w \in \mathcal{S}} \mu_\pi^{h''|h'+1}(u, v | z; \mathbb{P}) \epsilon_k^{h''}(w | u, v) \mu_\pi^{h|h''+1}(s, a | w; \hat{P}_\pi) \\ &\leq \sum_{h''=h'+1}^{h-1} \sum_{u \in \mathcal{S}, v \in \mathcal{A}, w \in \mathcal{S}} \mu_\pi^{h''|h'+1}(u, v | z; \mathbb{P}) \epsilon_k^{h''}(w | u, v) \pi^h(a | s), \end{aligned}$$

where the first step follows from the same reasoning as the unconditioned ones and the second step,

$$\mu_\pi^{h|h''+1}(s, a | w; \hat{P}_\pi) = \pi^h(a | s) \Pr\{s^h = s | s^{h''+1} = w, \pi, \hat{P}_\pi\} \leq \pi^h(a | s).$$

Hence, plugging back into Eq. (19) gives its bound as

$$\mathbb{E} \left[\sum_{h, s, a} \sum_{\pi \in \Pi} p_k(\pi) \sum_{h'=0}^{h-1} \sum_{x, y, z} \mu_\pi^{h'}(x, y; \mathbb{P}) \epsilon_k^{h'}(z | x, y) \mu_\pi^{h|h'+1}(s, a | z; \hat{P}_\pi) \right]$$

$$\leq \sum_{h,s,a} \mathbb{E} \left[\sum_{\pi \in \Pi} p_k(\pi) \sum_{h'=0}^{h-1} \sum_{x,y,z} \mu_\pi^{h'}(x,y;\mathbb{P}) \epsilon_k^{h'}(z|x,y) \mu_\pi^{h|h'+1}(s,a|z;\mathbb{P}) \right] + \\ \sum_{h,s,a} \mathbb{E} \left[\sum_{\pi \in \Pi} p_k(\pi) \sum_{h'=0}^{h-1} \sum_{x,y,z} \mu_\pi^{h'}(x,y;\mathbb{P}) \epsilon_k^{h'}(z|x,y) \sum_{h''=h'+1}^{h-1} \sum_{u,v,w} \mu_\pi^{h''|h'+1}(u,v|z;\mathbb{P}) \epsilon_k^{h''}(w|u,v) \pi^h(a|s) \right].$$

The remaining part of the proof is exactly the same as that for Lemma A.3 of [Jin et al. \(2022\)](#), which eventually shows,

$$\sum_{k=1}^K \text{BIAS1}_k^\varepsilon = \tilde{O} \left(\frac{HSAK}{L} + H^2 S \sqrt{AK} + H^3 S^3 A + \delta HK \right). \quad (20)$$

Combining the two parts together (with or without \mathcal{E}_k) gives,

$$\text{BIAS1} = \tilde{O} \left(\frac{HSAK}{L} + H^2 S \sqrt{AK} + H^3 S^3 A + \delta HK \right),$$

as claimed. \square

Proof of Theorem 24. This proof is quite simple. We still decompose BIAS2 into two parts:

$$\text{BIAS2} = \mathbb{E} \left[\sum_{k=1}^K \left(V(\pi^*; \widehat{\ell}_k, \mathbb{P}) - V(\pi^*; \ell_k, \mathbb{P}) \right) \right] \\ = \sum_{k=1}^K \mathbb{E} \left[\left(V(\pi^*; \widehat{\ell}_k, \mathbb{P}) - V(\pi^*; \ell_k, \mathbb{P}) \right) \mathbb{1}[\mathcal{E}_k] \right] + \sum_{k=1}^K \mathbb{E} \left[\left(V(\pi^*; \widehat{\ell}_k, \mathbb{P}) - V(\pi^*; \ell_k, \mathbb{P}) \right) \mathbb{1}[\neg \mathcal{E}_k] \right].$$

For the first term, as $\mathbb{1}[\mathcal{E}_k]$ infers $\widehat{q}_k^h(s,a) \leq q_k^h(s,a)$, from [Corollary 40](#), we have

$$\mathbb{E}[\widehat{\ell}_k(s,a) | \mathcal{F}_{k-1}] \leq \ell_k(s,a), \quad \forall k \in [K], (s,a) \in \mathcal{S} \times \mathcal{A}.$$

Therefore, as both π^* and \mathbb{P} are deterministic, this term is upper bounded by 0. For the second term, we trivially bound each of the summand by $HL \Pr\{\neg \mathcal{E}_k\} \leq 4\delta HL$ as $|\widehat{\ell}_k^h(s,a)| \leq L$. Therefore, combining two terms together completes the proof. \square

C Analysis of Episodic AMDP Algorithms with Delayed Feedback (Theorem 6)

In this section, we consider episodic AMDPs with delayed bandit feedback and unknown transitions. The algorithm is presented in [Algorithm 4](#), which is very similar to [Algorithm 3](#) except for the part on handling delayed feedback, highlighted in violet.

C.1 Regret Decomposition

Proof of Theorem 6. For this case, we still use the regret decomposition as [Theorem 5](#), as follows:

$$\mathcal{R}_K = \underbrace{\mathbb{E} \left[\sum_{k=1}^K \left(V(\pi_k; \ell_k, \mathbb{P}) - V(\pi_k; \ell_k, P_k) \right) \right]}_{\text{ERROR}} + \underbrace{\mathbb{E} \left[\sum_{k=1}^K \left(V(\pi_k; \ell_k, P_k) - V(\pi_k; \widehat{\ell}_k, P_k) \right) \right]}_{\text{BIAS1}} + \\ \underbrace{\mathbb{E} \left[\sum_{k=1}^K \left(V(\pi_k; \widehat{\ell}_k, P_k) - V(\pi^*; \widehat{\ell}_k, \mathbb{P}) \right) \right]}_{\text{ESTREG}} + \underbrace{\mathbb{E} \left[\sum_{k=1}^K \left(V(\pi^*; \widehat{\ell}_k, \mathbb{P}) - V(\pi^*; \ell_k, \mathbb{P}) \right) \right]}_{\text{BIAS2}}.$$

Algorithm 4 FTPL for Episodic AMDPs with Delayed Bandit Feedback and Unknown Transition

Require: Laplace distribution parameter η . Geometric Re-sampling parameter L .

- 1: Initialize $\mathcal{P}_1 \leftarrow (\Delta(\mathcal{S}))^{[H] \times \mathcal{S} \times \mathcal{A}}$.
 - 2: Sample perturbation $\widehat{\ell}_0 = z$ such that $z^h(s, a)$ is an independent sample of $\text{Laplace}(\eta)$.
 - 3: **for** $k = 1, 2, \dots, K$ **do**
 - 4: Let $(\pi_k, P_k) = \operatorname{argmin}_{(\pi, P) \in \Pi \times \mathcal{P}_k} V(\pi; \sum_{k' \in \Omega_k} \widehat{\ell}_{k'} + z, P)$ by Extended Value Iteration (Jaksch et al., 2010), where $\Omega_k = \{k' \mid k' + d_{k'} < k\}$. (See also Remark 16 for more details.)
 - 5: **for** $h = 1, 2, \dots, H$ **do**
 - 6: Observe s_k^h , play $a_k^h = \pi_k(s_k^h)$, suffer loss $\ell_k^h(s_k^h, a_k^h)$.
 - 7: **for All** $k' < k$ **such that** $k' + d_{k'} = k$ **do**
 - 8: **for** $h = 1, 2, \dots, H$ **do**
 - 9: **for** $M_{k'}^h = 1, 2, \dots, L$ **do**
 - 10: Sample a fresh perturbation \tilde{z} in the same way as z .
 - 11: Calculate $(\pi_{k'}^h, P_{k'}^h) = \operatorname{argmin}_{(\pi, P) \in \Pi \times \mathcal{P}_{k'}} V(\pi; \sum_{j \in \Omega_{k'}} \widehat{\ell}_j + \tilde{z})$.
 - 12: Pick the transition $\widehat{P}_{k'}^h \in \mathcal{P}_{k'}$ such that $\mu_\pi^h(s_{k'}^h, a_{k'}^h; \widehat{P}_{k'}^h)$ is maximized via the COMP-UOB procedure proposed by Jin et al. (2020).
 - 13: Simulate $\pi_{k'}^h$ for h steps starting from s^1 and following transitions $(\widehat{P}_{k'}^h)^1, \dots, (\widehat{P}_{k'}^h)^h$.
 - 14: **if** $(s_{k'}^h, a_{k'}^h)$ is visited at step h or $M_{k'}^h = L$ **then**
 - 15: Set $\widehat{\ell}_{k'}^h(s, a) = M_{k'}^h \cdot \ell_{k'}^h(s_{k'}^h, a_{k'}^h) \cdot \mathbb{1}[(s_{k'}^h, a_{k'}^h) = (s, a)]$ and break.
 - 16: Calculate \mathcal{P}_{k+1} according to Eq. (10).
-

Note that as delays will not affect transitions as well as the loss estimators (viewed in hindsight, i.e., the sequence $\{\widehat{\ell}_k\}_{k \in [K]}$ will be the same as if there is no delays), so the ERROR, BIAS1 and BIAS2 can still be bounded by Theorems 21, 22 and 24, respectively. The only difference occurs when bounding ESTREG, which we show as follows.

Lemma 26 (Bounding ESTREG Term with Delayed Feedback). *The ESTREG term is bounded by*

$$\mathbb{E} \left[\sum_{k=1}^K \left(V(\pi_k; \widehat{\ell}_k, P_k) - V(\pi^*; \widehat{\ell}_k, \mathbb{P}) \right) \right] \leq \frac{2H}{\eta} (1 + \ln(SA)) + 5\eta H^2 SAK + \eta H^2 SA\mathfrak{D} + 12\delta KHL.$$

As mentioned in the main body, the key difference is that, we will compete a learner that is not only cheating but also stepping one episode further. However, as it is still using FTPL, we can still bound the stability term as in Lemma 20. Therefore, the proof is postponed to the end of this section.

Combining the bounds for the four terms together, we will have

$$\mathcal{R}_T \leq \tilde{\mathcal{O}} \left(H^2 SA\sqrt{K} + \frac{H}{\eta} + \eta H^2 SA(K + \mathfrak{D}) + \frac{SAHK}{L} + \delta HKL \right).$$

Therefore, picking $\eta = \left(\sqrt{HSA(K + \mathfrak{D})} \right)$, $L = \sqrt{SAK/H}$ and $\delta = 1/K$ gives

$$\mathcal{R}_T \leq \tilde{\mathcal{O}} \left(H^2 SA\sqrt{K} + H^{3/2} \sqrt{SA\mathfrak{D}} \right),$$

as claimed. □

Proof of Lemma 26. Slightly different from the main text, we now consider the following *two* learners, where the first one is a “cheating learner” that does not suffer any delays, and the second one is a “cheating leader” that not only does not suffer any delays, but also looks one step further.

$$(\widehat{\pi}_k, \widehat{P}_k) \triangleq \operatorname{argmin}_{(\pi, P) \in \Pi \times \mathcal{P}_k} V(\pi; \widehat{\ell}_{0:k-1}, P), \quad (\widetilde{\pi}_{k+1}, \widetilde{P}_{k+1}) \triangleq \operatorname{argmin}_{(\pi, P) \in \Pi \times \mathcal{P}_k} V(\pi; \widehat{\ell}_{0:k}, P).$$

Note that both of them are defined w.r.t. transitions in \mathcal{P}_k instead of the subset \mathcal{P}_{k+1} , which is the same as Appendix B.2. We also define the following three density functions with respect to

the perturbation $z: p_k(\pi, P)$ for (π_k, P_k) conditioning on $\widehat{\ell}_1, \widehat{\ell}_2, \dots, \widehat{\ell}_{k-1}, \widehat{p}_k(\pi, P)$ for $(\widehat{\pi}_k, \widehat{P}_k)$ conditioning on $\widehat{\ell}_1, \widehat{\ell}_2, \dots, \widehat{\ell}_{k-1}$ and $\widetilde{p}_{k+1}(\pi, P)$ for $(\widetilde{\pi}_{k+1}, \widetilde{P}_{k+1})$ conditioning on $\widehat{\ell}_1, \widehat{\ell}_2, \dots, \widehat{\ell}_k$.

The purpose of defining two learners is to decouple the effects from delays and the inherent FTPL regret. One can see that our $(\widehat{\pi}_k, \widehat{P}_k)$ is equivalent to (π_k, P_k) in Appendix B.2 while $(\widetilde{\pi}_{k+1}, \widetilde{P}_{k+1})$ remains the same. Therefore, the difference between $(\widehat{\pi}_k, \widehat{P}_k)$ and $(\widetilde{\pi}_k, \widetilde{P}_k)$ can be bounded exactly the same as Appendix B.2 and we only need to care about delays, i.e., the difference between (π_k, P_k) and $(\widehat{\pi}_k, \widehat{P}_k)$. Formally, we decompose the ESTREG into three terms:

$$\begin{aligned} \mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi_k}^h(P_k) - \mu_{\pi^*}^h(\mathbb{P}), \widehat{\ell}_k^h \rangle \right] &= \underbrace{\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi_k}^h(P_k) - \mu_{\widehat{\pi}_k}^h(\widehat{P}_k), \widehat{\ell}_k^h \rangle \right]}_{\text{Cheating regret}} + \\ &\underbrace{\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\widehat{\pi}_k}^h(\widehat{P}_k) - \mu_{\widetilde{\pi}_{k+1}}^h(\widetilde{P}_{k+1}), \widehat{\ell}_k^h \rangle \right]}_{\text{Stability term}} + \underbrace{\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\widetilde{\pi}_{k+1}}^h(\widetilde{P}_{k+1}) - \mu_{\pi^*}^h(\mathbb{P}), \widehat{\ell}_k^h \rangle \right]}_{\text{Error term}}. \end{aligned}$$

Note that the error term and the stability term are exactly the same as Appendix B.2, so we can directly make use of Lemmas 18 and 20 and bound them by $\frac{2H}{\eta}(1 + \ln(SA)) + 4\delta KHL$ and $3\eta H^2SAK + 4\delta KHL$, respectively. Now consider the cheating regret. Similar to the stability term, we will have the following single-step stability lemma:

Lemma 27. *For any $k \in [K]$, $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $(\pi, P) \in \Pi \times \mathcal{P}_k$, we have*

$$\widehat{p}_k(\pi, P) \geq p_k(\pi, P) \exp \left(-\eta \sum_{k' \in \widetilde{\Omega}_k} \sum_{h=1}^H \|\widehat{\ell}_{k'}^h\|_1 \right),$$

where $\widetilde{\Omega}_k \triangleq [k-1] \setminus \Omega_k = \{k' < k \mid k' + d_{k'} \geq k\}$, i.e., the first $k-1$ rounds excluding those where the feedback is available before round k .

Proof. Note that the proof of Lemma 19 does not rely on the concrete choice of π_k and $\widetilde{\pi}_{k+1}$. Therefore, adopting the proof of Lemma 19 with $\widehat{\pi}_k$ as $\widetilde{\pi}_{k+1}$ will complete our proof. \square

With the help of Lemma 27, we can bound the cheating regret similar to the stability term. To see this, consider a fixed $k \in [K]$, we have

$$\begin{aligned} &\mathbb{E} \left[\sum_{\pi \in \Pi} \int_{\mathcal{P}_k} (p_k(\pi, P) - \widehat{p}_k(\pi, P)) \sum_{h=1}^H \langle \mu_{\pi}^h(P), \widehat{\ell}_k^h \rangle dP \right] \\ &= \mathbb{E} \left[\sum_{\pi \in \Pi} \int_{\mathcal{P}_k} (p_k(\pi, P) - \widehat{p}_k(\pi, P)) \sum_{h=1}^H \langle \mu_{\pi}^h(P), \widehat{\ell}_k^h \rangle dP \mathbb{1}[\mathcal{E}_k] \right] + \\ &\mathbb{E} \left[\sum_{\pi \in \Pi} \int_{\mathcal{P}_k} (p_k(\pi, P) - \widehat{p}_k(\pi, P)) \sum_{h=1}^H \langle \mu_{\pi}^h(P), \widehat{\ell}_k^h \rangle d\mathbb{1}[\neg \mathcal{E}_k] \right] \\ &\leq \mathbb{E} \left[\eta \sum_{k' \in \widetilde{\Omega}_k} \sum_{h'=1}^H \|\widehat{\ell}_{k'}^{h'}\|_1 \cdot \sum_{\pi \in \Pi} \int_{\mathcal{P}_k} p_k(\pi, P) \sum_{h=1}^H \langle \mu_{\pi}^h(P), \widehat{\ell}_k^h \rangle dP \mathbb{1}[\mathcal{E}_k] \right] + 4\delta KHL \\ &\leq \eta \mathbb{E} \left[\sum_{h=1}^H \sum_{(k', h') \in \widetilde{\Omega}_k \times [H]} \mathbb{1}[(k', h') \neq (k, h)] \|\widehat{\ell}_{k'}^{h'}\|_1 \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \widehat{q}_k^h(s, a) \widehat{\ell}_k^h(s, a) \mathbb{1}[\mathcal{E}_k] \right] + \\ &\eta \mathbb{E} \left[\sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} \sum_{h=1}^H \widehat{q}_k^h(s, a) \left(\widehat{\ell}_k^h(s, a) \right)^2 \mathbb{1}[\mathcal{E}_k] \right] + 4\delta KHL, \end{aligned}$$

where $\hat{q}_k^h(s, a)$ is the actual probability of reaching (s, a) and $q_k^h(s, a)$ is the probability of reaching (s, a) in a single Geometric Re-sampling trial, as defined in Equations (12) and (13). Note that $\mathbb{1}[\mathcal{E}_k]$ implies $\hat{q}_k^h(s, a) \leq q_k^h(s, a)$.

For the second term, using Lemma 39 and $\hat{q}_k^h(s, a) \leq q_k^h(s, a)$ gives $2\eta SAH$. For the first one, taking expectation w.r.t. M_k^h in $\hat{\ell}_k^h(s, a) = \mathbb{1}[(s_k^h, a_k^h) = (s, a)]\ell_k^h(s, a)M_k^h(s, a)$ and then w.r.t. $\|\hat{\ell}_{k'}^h\|_1$ as in Lemmas 12 and 20 gives $\eta H^2 SA|\tilde{\Omega}_k|$. Further noticing that

$$\sum_{k=1}^K |\tilde{\Omega}_k| = \sum_{k=1}^K \sum_{k'=1}^{k-1} \mathbb{1}[k' + d_{k'} \geq k] = \sum_{k'=1}^{K-1} \sum_{k=k'+1}^K \mathbb{1}[k' + d_{k'} \geq k] = \sum_{k'=1}^{K-1} d_{k'} = \mathfrak{D},$$

we have the cheating regret is bounded by

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \langle \mu_{\pi_k}^h(P_k) - \mu_{\hat{\pi}_k}^h(\hat{P}_k), \hat{\ell}_k^h \rangle \right] \leq \eta H^2 SA \mathfrak{D} + 2\eta HSAK + 4\delta KHL.$$

The ESTREG term is then consequently bounded by

$$\text{ESTREG} \leq \frac{2H}{\eta} (1 + \ln(SA)) + 3\eta H^2 SAK + \eta H^2 SA \mathfrak{D} + 2\eta HSAK + 12\delta KHL,$$

which is at most $\frac{2H}{\eta} (1 + \ln(SA)) + 5\eta H^2 SAK + \eta H^2 SA \mathfrak{D} + 12\delta KHL$, as claimed. \square

D Analysis of Infinite-horizon AMDP Algorithms

D.1 FTPL-Based Efficient Algorithm (Theorem 7)

In this section present our Algorithm 6 together with its analysis. As described in the main body, we will divide the time horizon $[T]$ into J epochs and fix a policy π_j for the j -th epoch, namely $\mathcal{T}_j = \{(j-1)H+1, (j-1)H+2, \dots, jH\}$ where $H = \frac{T}{J}$ is the length of each epoch (overloading the notation H from the episodic setting since they have a similar meaning).

D.1.1 Switching Procedure

The most significant difference between infinite-horizon AMDPs and episodic AMDPs is that the agent will not be reset to s^1 at the beginning of an ‘‘epoch’’. To formalize our problem as an online linear optimization problem (i.e., the total loss represented as $\sum_{t=1}^T \langle \mu_{\pi_t}^t, \ell_t \rangle$), we have to ensure the distribution over all states is exactly $\mu_{\pi_j}^t$ for most $t \in \mathcal{T}_j$. Before presenting the switching procedure from Chandrasekaran and Tewari (2021), we first restate the assumption together with several properties that they used. For the sake of completeness, we also include their proofs here.

Assumption 28 (Existence of a Staying State, Restatement of Assumption 2 and Chandrasekaran and Tewari (2021, Assumption 5.1)). *The MDP \mathcal{M} has a state s^* and an action a^* such that $\mathbb{P}(s^* | s^*, a^*) = 1$.*

Lemma 29 (Chandrasekaran and Tewari (2021, Lemma 5.2)). *For any two distinct states $s, s' \in \mathcal{S}$, there exists a policy $\pi_{s,s'}$ and $l_{s,s'} \leq 2D$ such that*

$$\Pr\{T(s' | \pi_{s,s'}, s) = l_{s,s'}\} \geq \frac{1}{4D}.$$

Proof. By definition of diameter (as in Definition 1), there exists a policy $\pi_{s,s'}$ such that $\mathbb{E}[T(s' | \mathcal{M}, \pi_{s,s'}, s)] \leq D$. By Markov’s inequality, this implies $\Pr\{T(s' | \mathcal{M}, \pi_{s,s'}, s) \leq 2D\} \geq \frac{1}{2}$. By pigeonhole principle, there consequently exists $l_{s,s'} \leq 2D$ such that $\Pr\{T(s' | \mathcal{M}, \pi_{s,s'}, s) = l_{s,s'}\} \geq \frac{1}{2} \cdot \frac{1}{2D} = \frac{1}{4D}$. \square

Lemma 30 (Chandrasekaran and Tewari (2021, Theorem 5.3)). *For an MDP that satisfies Assumption 28, there exists $l^* \leq 2D$ such that for all states $s' \neq s^*$, there exists policy $\pi_{s'}$ such that*

$$\Pr\{T(s' | \mathcal{M}, \pi_{s'}, s^*) = l^*\} \geq \frac{1}{4D}.$$

Algorithm 5 Policy Switching in Infinite-Horizon AMDP (Chandrasekaran and Tewari, 2021)

Require: Current state $s \in \mathcal{S}$. Goal policy $\pi \in \Pi$. Current time t .

- 1: **while true do**
- 2: Move to state s^* using policy π_{s,s^*} as defined by Lemma 29 and update s, t concurrently.
- 3: Sample the target state $g \sim \mu_\pi^t$.
- 4: Use policy π_g from Lemma 30 to move l^* steps from s^* and update s, t concurrently.
- 5: **if** $s = g$ **then**
- 6: Sample a Bernoulli random variable $I \sim \text{Ber}(\frac{p^*}{p_g})$.
- 7: **if** $I = 1$ **then**
- 8: **return**

▷ The **while** loop will repeat if $s \neq g$ or $I \neq 1$.

Furthermore, denote $p_{s'}$ as the probability above. Let $p^* = \min_{s \in \mathcal{S}} p_s$. Then $p^* \geq \frac{1}{4D}$.

Proof. From the previous lemma, there exists an $l_{s'} \leq 4D$ for all $s' \neq s^*$ such that there is a policy $\pi_{s^*,s'}$ hitting s' from s^* in time exactly $l_{s'}$ with probability at least $\frac{1}{4D}$. Let $l^* = \max_{s' \neq s^*} l_{s'}$ and π_{s^*} be the policy that first stays at s^* for $(l^* - l_{s'})$ steps and then follows $\pi_{s^*,s'}$ for $l_{s'}$ steps suffices. \square

Now we are able to present the switching procedure from Chandrasekaran and Tewari (2021), as in Algorithm 5.

Theorem 31 (Correctness of Algorithm 5, Chandrasekaran and Tewari (2021, Lemma 5.6)). *Let the random variable denoting the time that Algorithm 5 terminates be t_{switch} . Then for any state $s \in \mathcal{S}$*

$$\Pr\{s_t = s \mid t_{\text{switch}} = t\} = \mu_\pi^t(s), \quad \forall t \in [T].$$

Proof. The key idea is to write

$$\Pr\{s_t = s \mid t_{\text{switch}} = t\} = \frac{\Pr\{s_t = s, g = s, t_{\text{switch}} = t\}}{\Pr\{t_{\text{switch}} = t\}}$$

and then bound the numerator and denominator separately. For the denominator,

$$\begin{aligned} & \Pr\{t_{\text{switch}} = t\} \\ &= \sum_{s \in \mathcal{S}} \Pr\{s_t = s, g = s, s_{t-l^*} = s^*\} \times \Pr\{t_{\text{switch}} = t \mid s_t = s, g = s, s_{t-l^*} = s^*\} \\ &= \sum_{s \in \mathcal{S}} \Pr\{s_t = s \mid g = s, s_{t-l^*} = s^*\} \times \Pr\{g = s, s_{t-l^*} = s^*\} \times \\ & \quad \Pr\{t_{\text{switch}} = t \mid s_t = s, g = s, s_{t-l^*} = s^*\} \\ &= \sum_{s \in \mathcal{S}} p_s \times \Pr\{g = s, s_{t-l^*} = s^*\} \times \frac{p^*}{p_s} = p^* \times \Pr\{s_{t-l^*} = s^*\}, \end{aligned}$$

where the last step used definition of p_s and I . For the numerator,

$$\begin{aligned} & \Pr\{g = s, s_t = s, t_{\text{switch}} = t\} \\ &= \Pr\{g = s, s_t = s, s_{t-l^*} = s^*, t_{\text{switch}} = t\} \\ &= \Pr\{t_{\text{switch}} = t, s_t = s \mid g = s, s_{t-l^*} = s^*\} \times \Pr\{g = s, s_{t-l^*} = s^*\} \\ &= \Pr\{s_t = s \mid g = s, s_{t-l^*} = s^*\} \times \Pr\{t_{\text{switch}} = t \mid s_t = s, g = s, s_{t-l^*} = s^*\} \times \\ & \quad \Pr\{s_{t-l^*} = s^*\} \times \Pr\{g = s \mid s_{t-l^*} = s^*\} \\ &= p_s \times \frac{p^*}{p_s} \times \Pr\{s_{t-l^*} = s^*\} \times \mu_{\pi'}^t(s). \end{aligned}$$

Plugging them back gives our desired result. \square

Theorem 32 (Efficiency of Algorithm 5, Chandrasekaran and Tewari (2021, Lemma 5.7)). *The expected time spent on Algorithm 5 is bounded by $12D^2$ for each execution.*

Algorithm 6 FTPL for Infinite-horizon AMDPs with Bandit Feedback and Known Transition

Require: Laplace distribution parameter η . Geometric Re-sampling parameter L .

- 1: Sample perturbations $\{z^t \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}\}_{t \in [T]}$ where $z^t(s, a) \sim \text{Laplace}(\eta)$.
- 2: **for** $j = 1, 2, \dots, J$ **do**
- 3: Calculate the policy π_j for this epoch as

$$\pi_j = \underset{\pi \in \Pi}{\operatorname{argmin}} \left(\sum_{j'=1}^{j-1} \sum_{t \in \mathcal{T}_{j'}} \langle \mu_{\pi}^t, \hat{\ell}^t \rangle + \sum_{t=1}^T \langle \mu_{\pi}^t, z^t \rangle \right). \quad (21)$$

- 4: Execute Algorithm 5 with parameters s^t, π_j, t (note that Algorithm 5 will update t internally).
- 5: **for** All remaining time slots in \mathcal{T}_j , i.e., $\mathcal{T}_j \cap [t, T]$ **do**
- 6: Play $a^t = \pi_j(s^t)$, observe the loss $\ell^t(s^t, a^t)$ and the next state $s^{t+1} \in \mathcal{S}$.
- 7: **for** $M^t = 1, 2, \dots, L$ **do**
- 8: Resample a fresh perturbation and get new policy π'_j from Eq. (21).
- 9: Draw a sample from $\text{Ber}(\mu_{\pi'_j}^t(s^t, a^t; \mathbb{P}))$. If it is 1 or $M^t = L$, terminate and set

$$\hat{\ell}^t(s, a) = \mathbb{1}[(s, a) = (s^t, a^t)] \ell^t(s^t, a^t) M^t, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

Proof. Every time we try to catch the policy from s^* , we succeed with probability $p^* \geq \frac{1}{4D}$. Thus, the expected number of times we try is $4D$ and each attempt takes $l^* \leq 2D$ steps. Between each of these attempts, we move at most D steps in expectation to reach s^* again. Thus, in total, we have

$$\mathbb{E}[t_{\text{switch}} - t_0] \leq 4D(2D + D) \leq 12D^2,$$

as claimed. \square

D.1.2 The Algorithm

With the help of Algorithm 5, we now present our algorithm, Algorithm 6. As mentioned in the main text, another important difference due to the “non-resetting” nature of an infinite-horizon AMDP is that, we have to generate T perturbations z^1, z^2, \dots, z^T , whereas only H perturbations is needed in the episodic settings. For each FTPL update, we will include all of them in the argmin operation, as in Eq. (21).

This difference can be explained from the contextual bandits’ point of view (c.f. Appendix B.1.5). In infinite-horizon AMDPs, the possible number of “contexts” is now T , as for each policy π , it will have T distinct features $\mu_{\pi}^1, \mu_{\pi}^2, \dots, \mu_{\pi}^T$. In contrast, for episodic AMDPs, there are only H different contexts as only $\{\mu_{\pi}^h\}_{h=1}^H$ can appear. Therefore, as noticed by Syrgkanis et al. (2016), we have to add perturbations to each of the contexts, which are in total T of them.

D.1.3 Proof of Main Theorem

Proof of Theorem 7. To calculate the regret guarantee of Algorithm 6, we consider the following quantity $\widetilde{\mathcal{R}}_T$ defined as if there is no cost for a policy switching. By Theorem 32, there can be at most JD^2 time slots spent on executing Algorithm 5. Henceforth, the difference between \mathcal{R}_T and $\widetilde{\mathcal{R}}_T$ is at most JD^2 .

$$\widetilde{\mathcal{R}}_T \triangleq \mathbb{E} \left[\sum_{t=1}^T \langle \mu_{\pi_{j(t)}}^t, \ell^t \rangle - \langle \mu_{\pi^*}^t, \ell^t \rangle \right] = \mathbb{E} \left[\sum_{j=1}^J \sum_{\pi \in \Pi} p_j(\pi) \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t, \ell^t \rangle - \sum_{t=1}^T \langle \mu_{\pi^*}^t, \ell^t \rangle \right], \quad (22)$$

where $p_j(\pi)$ is the probability of picking π w.r.t. z , conditioning on $\mathcal{F}_{(j-1)H}$ and $j(t)$ is the epoch that t belongs to, namely $j(t) = \lceil \frac{t}{H} \rceil$. Then, we can decompose $\widetilde{\mathcal{R}}_T$ into three terms exactly the same as what we did in Appendix B.1:

$$\widetilde{\mathcal{R}}_T = \mathbb{E} \left[\underbrace{\sum_{t=1}^T \langle \mu_{\pi_{j(t)}}^t, \ell^t - \hat{\ell}^t \rangle + \sum_{t=1}^T \langle \mu_{\pi^*}^t, \hat{\ell}^t - \ell^t \rangle}_{\text{GR error term}} \right] +$$

$$\underbrace{\mathbb{E} \left[\sum_{j=1}^J \sum_{\pi \in \Pi} p_{j+1}(\pi) \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t - \mu_{\pi^*}^t, \widehat{\ell}^t \rangle \right]}_{\text{Error term}} + \underbrace{\mathbb{E} \left[\sum_{j=1}^J \sum_{\pi \in \Pi} (p_j(\pi) - p_{j+1}(\pi)) \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t, \widehat{\ell}^t \rangle \right]}_{\text{Stability term}}.$$

The GR error term is quite similar to Appendix B.1:

Lemma 33. *The GR error term is bounded by*

$$\mathbb{E} \left[\sum_{t=1}^T \langle \mu_{\pi_{j(t)}}^t, \ell^t - \widehat{\ell}^t \rangle + \sum_{t=1}^T \langle \mu_{\pi^*}^t, \widehat{\ell}^t - \ell^t \rangle \right] \leq \frac{SAT}{eL}.$$

For the error term, we still use the similar “be-the-leader” analysis as Lemma 10, except for we are now facing a slightly different V -function (which is defined for infinite-horizon). Moreover, as mentioned in the main text, we are using a different bound when facing T different perturbations. As a result, we will have worse dependency on S and A , but with better dependency on the number of contexts, which is T here (and is H in episodic settings). The result is stated as follows:

Lemma 34. *The error term is bounded by*

$$\mathbb{E} \left[\sum_{j=1}^J \sum_{\pi \in \Pi} p_{j+1}(\pi) \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t - \mu_{\pi^*}^t, \widehat{\ell}^t \rangle \right] \leq \frac{10}{\eta} S \sqrt{AT \ln A}.$$

For the stability term, again much similar to Appendix B.1, we have

Lemma 35. *The stability term is bounded by*

$$\mathbb{E} \left[\sum_{j=1}^J \sum_{\pi \in \Pi} (p_{j+1}(\pi) - p_j(\pi)) \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t, \widehat{\ell}^t \rangle \right] \leq 2\eta H^2 SAJ.$$

Therefore, our regret is bounded by

$$\mathcal{R}_T \leq \widetilde{\mathcal{R}}_T + JD^2 \leq \frac{SAT}{eL} + \frac{10}{\eta} S \sqrt{AT \ln A} + 2\eta H^2 SAJ + JD^2.$$

Picking $\eta = S^{1/3} D^{-2/3} T^{-1/3}$, $J = S^{2/3} A^{1/2} D^{-4/3} T^{5/6}$ and $L = S^{1/3} A^{1/2} D^{-2/3} T^{1/6}$ gives $\mathcal{R}_T = \widetilde{\mathcal{O}}(S^{2/3} A^{1/2} D^{2/3} T^{5/6})$. \square

Proof of Lemma 33. We follow the proof of Lemma 9 by replacing K with J and the GR estimator $\widehat{\ell}_k^h$ with $\widehat{\ell}_t$. First notice that, from Lemma 38, $\mathbb{E}[\widehat{\ell}_t(s, a) \mid \mathcal{F}_{(j-1)H}] \leq \ell^t(s, a)$ for all $t \in \mathcal{T}_j$. Moreover, as π^* is deterministic (i.e., it does not depend on the randomness from the algorithm), the term related to $\mu_{\pi^*}^t$ is bounded by

$$\mathbb{E} \left[\sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi^*}^t, \widehat{\ell}^t - \ell^t \rangle \right] = \mathbb{E} \left[\sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi^*}^h, \mathbb{E}[\widehat{\ell}^t \mid \mathcal{F}_{(j-1)H}] - \ell^t \rangle \right] \leq 0.$$

For the first term, again by Lemma 38, we have

$$\mathbb{E} \left[\sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi_j}^t, \ell^t - \widehat{\ell}^t \rangle \right] = \sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbb{E} \left[\mu_{\pi_j}^t(s, a) \cdot (1 - q^t(s, a))^L \ell^t(s, a) \right],$$

where $q^t(s, a)$ is the probability of visiting (s, a) in a single execution of the Geometric Re-sampling process, which is just

$$q^t(s, a) = \mathbb{E}[\mu_{\pi_j}^t(s, a)] = \sum_{\pi \in \Pi} p_j(\pi) \mu_{\pi}^t(s, a)$$

in our case. By noticing that $q(1-q)^L \leq qe^{-Lq} \leq \frac{1}{eL}$ (Neu and Bartók, 2013), we have

$$\mathbb{E} \left[\sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi_j}^t, \ell^t - \widehat{\ell}^t \rangle \right] \leq HJSA \frac{1}{eL} = \frac{SAT}{eL},$$

as claimed. \square

Proof of Lemma 34. The proof still uses the standard “be-the-leader” technique, but in a slightly different manner as we are adding perturbations to all time indices. Instead, we follow the idea of Syrgkanis et al. (2016, Lemma 7) and prove by induction that the following inequality holds for all J and any policy $\pi \in \Pi$:

$$\sum_{t=1}^T \langle \mu_{\pi_1}^t, z^t \rangle + \sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi_{j+1}}^t, \ell^t \rangle \leq \sum_{t=1}^T \langle \mu_{\pi}^t, z^t \rangle + \sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t, \ell^t \rangle.$$

Obviously, for $J = 0$, this inequality holds. Suppose that this inequality holds for J , then we consider $J + 1$. Let $\pi = \pi_{J+2}$. Adding $\sum_{t \in \mathcal{T}_{J+1}} \langle \mu_{\pi_{J+2}}^t, \widehat{\ell}^t \rangle$ to both sides gives

$$\sum_{t=1}^T \langle \mu_{\pi_1}^t, z^t \rangle + \sum_{j=1}^{J+1} \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi_{j+1}}^t, \ell^t \rangle \leq \sum_{t=1}^T \langle \mu_{\pi_{J+2}}^t, z^t \rangle + \sum_{j=1}^{J+1} \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi_{J+2}}^t, \ell^t \rangle.$$

However, by definition of π_{J+2} (which is the argmin of the right-handed-side for all policies), it is further bounded by

$$\sum_{t=1}^T \langle \mu_{\pi_1}^t, z^t \rangle + \sum_{j=1}^{J+1} \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi_{j+1}}^t, \ell^t \rangle \leq \sum_{t=1}^T \langle \mu_{\pi}^t, z^t \rangle + \sum_{j=1}^{J+1} \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t, \ell^t \rangle$$

for any policy $\pi \in \Pi$, which means that the induction hypothesis for $J + 1$. Therefore, by picking $\pi = \pi^*$ for the real J , we can conclude that

$$\sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi_{j+1}}^t, \ell^t \rangle - \sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi^*}^t, \ell^t \rangle \leq \sum_{t=1}^T \langle \mu_{\pi^*}^t, z^t \rangle - \sum_{t=1}^T \langle \mu_{\pi_1}^t, z^t \rangle.$$

Then taking expectation on both sides gives the error term is bounded by

$$\mathbb{E}_z \left[\max_{\pi \in \Pi} \sum_{t=1}^T \langle \mu_{\pi}^t, z^t \rangle - \min_{\pi \in \Pi} \sum_{t=1}^T \langle \mu_{\pi}^t, z^t \rangle \right],$$

which is bounded by $\frac{10}{\eta} \sqrt{TSA \cdot \ln|\Pi|} = \frac{10}{\eta} S \sqrt{AT \ln A}$ by Lemma 45 (note that as $\ln|\Pi| = S \ln A < SAT$, the condition of applying Lemma 45 indeed holds). \square

Proof of Lemma 35. This follows directly from Lemma 12 with some slight modifications as well. For clarity, we rewrite the full proof here.

We first give the single-step stability lemma for infinite-horizon AMDPs, whose proof will be presented later:

Lemma 36. For all $j \in [J]$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$p_{j+1}(\pi) \geq p_j(\pi) \exp \left(-\eta \sum_{t \in \mathcal{T}_j} \|\widehat{\ell}^t\|_1 \right), \quad \forall \pi \in \Pi.$$

By summing up Lemma 36 for all $\pi \in \Pi$ and using the fact that $1 - \exp(-x) \leq x$, we have

$$\sum_{\pi \in \Pi} (p_j(\pi) - p_{j+1}(\pi)) \sum_{t \in \mathcal{T}_j} \langle \mu_\pi^t, \widehat{\ell}^t \rangle \leq \eta \sum_{t' \in \mathcal{T}_j} \|\widehat{\ell}^{t'}\|_1 \cdot \sum_{\pi \in \Pi} p_j(\pi) \sum_{t \in \mathcal{T}_j} \langle \mu_\pi^t, \widehat{\ell}^t \rangle, \quad \forall j \in [J]. \quad (23)$$

Again noticing that $M^t = \min\{\text{Geo}(q^t(s^t, a^t)), L\}$ where $q^t(s, a) = \sum_{\pi \in \Pi} p_j(\pi) \mu_\pi^t(s, a)$ if $t \in \mathcal{T}_j$. Then calculate the expectation of $\widehat{\ell}^t(s, a)$ only with respect to M^t , we will have

$$\mathbb{E} \left[\widehat{\ell}^t(s, a) \mid (s^t, a^t) = (s, a) \right] \leq \frac{\ell^t(s, a)}{q^t(s, a)}.$$

Let $\mathbb{1}^t(s, a)$ be the shorthand notation of $\mathbb{1}[(s^t, a^t) = (s, a)]$. Then for those $t' \neq t$ in Eq. (23),

$$\begin{aligned} & \eta \mathbb{E} \left[\sum_{t \in \mathcal{T}_j} \sum_{s, a} \sum_{\pi \in \Pi} p_j(\pi) \mu_\pi^t(s, a) \widehat{\ell}^t(s, a) \sum_{t' \neq t} \|\widehat{\ell}^{t'}\|_1 \mid \mathcal{F}_{(j-1)H} \right] \\ & \stackrel{(a)}{\leq} \eta \mathbb{E} \left[\sum_{h=1}^H \sum_{s, a} \mathbb{1}^t(s, a) \ell^t(s, a) \frac{\sum_{\pi \in \Pi} p_j(\pi) \mu_\pi^t(s, a)}{q^t(s, a)} \sum_{t' \neq t} \|\widehat{\ell}^{t'}\|_1 \mid \mathcal{F}_{(j-1)H} \right] \\ & \stackrel{(b)}{\leq} \eta H \mathbb{E} \left[\sum_{t' \in \mathcal{T}_j} \sum_{s, a} \|\widehat{\ell}^{t'}\|_1 \mid \mathcal{F}_{(j-1)H} \right] \stackrel{(c)}{\leq} \eta H^2 S A. \end{aligned}$$

where (a) is taking expectation w.r.t. M^t , (b) used the definition of q^t together with the fact that $\sum_{(s, a)} \mathbb{1}^t(s, a) = 1$, and (c) used the fact that $\mathbb{E}[\widehat{\ell}^{t'}(s, a) \mid \mathcal{F}_{(j-1)H}] \leq \ell^{t'}(s, a) \leq 1$ (Lemma 38).

For those terms with $t' = t$ in Eq. (23), by direct calculation and the fact that $\widehat{\ell}^t$ is a one-hot vector, we can bound them as

$$\eta \mathbb{E} \left[\sum_{t \in \mathcal{T}_j} \sum_{s, a} \sum_{\pi \in \Pi} p_j(\pi) \mu_\pi^t(s, a) \left(\widehat{\ell}^t(s, a) \right)^2 \mid \mathcal{F}_{(j-1)H} \right] \leq 2\eta \mathbb{E} \left[\sum_{h, s, a} \frac{q^t(s, a)}{q^t(s, a)} \mid \mathcal{F}_{(j-1)H} \right] \leq 2\eta H S A$$

by noticing $\mathbb{E}[(\widehat{\ell}^t(s, a))^2 \mid \mathcal{F}_{(j-1)H}] \leq 2(q^t(s, a))^{-1}$ (Lemma 39). Combining the terms with $t' \neq t$ and the ones with $t' = t$ gives our conclusion. \square

Proof of Lemma 36. The proof will be similar to, but different from Lemma 11, as we are now adding different perturbations. We now use a slightly different definition of the best-function. Let $\pi = \text{best}(\ell, z)$ where $\ell = \{\ell^1, \ell^2, \dots, \ell^m\}$ and $z = \{z^1, z^2, \dots, z^T\}$ to denote

$$\pi = \underset{\pi \in \Pi}{\text{argmin}} \left(\sum_{t=1}^m \langle \mu_\pi^t, \ell^t \rangle + \sum_{t=1}^T \langle \mu_\pi^t, z^t \rangle \right).$$

Then we have

$$\begin{aligned} p_j(\pi) &= \int_z \mathbb{1} \left[\pi = \text{best} \left(\{\widehat{\ell}^1, \dots, \widehat{\ell}^{(j-1)H}\}, \{z^1, z^2, \dots, z^T\} \right) \right] f(z) \, dz \\ &= \int_z \mathbb{1} \left[\pi = \text{best} \left(\{\widehat{\ell}^1, \dots, \widehat{\ell}^{(j-1)H}\}, \{z^1, \dots, z^{(j-1)H}, z^{(j-1)H+1} + \widehat{\ell}^{(j-1)H+1}, \dots, z^{jH} + \widehat{\ell}^{jH}, z^{jH+1}, \dots, z^T\} \right) \right] \\ & \quad f \left(z + \{0, \dots, 0, \widehat{\ell}^{(j-1)H+1}, \dots, \widehat{\ell}^{jH}, 0, 0, \dots, 0\} \right) \, dz \\ &= \int_z \mathbb{1} \left[\pi = \text{best} \left(\{\widehat{\ell}^1, \dots, \widehat{\ell}^{(j-1)H}, \widehat{\ell}^{(j-1)H+1}, \dots, \widehat{\ell}^{jH}\}, z \right) \right] \\ & \quad f \left(z + \{0, \dots, 0, \widehat{\ell}^{(j-1)H+1}, \dots, \widehat{\ell}^{jH}, 0, 0, \dots, 0\} \right) \, dz, \end{aligned}$$

Algorithm 7 Hedge for Infinite-horizon AMDPs with Bandit Feedback and Known Transition

Require: Learning rate η . Number of epochs J .

- 1: **for** $j = 1, 2, \dots, J$ **do**
- 2: Calculate the distribution of policies for the j -th epoch as

$$p_j(\pi) \propto \exp \left(-\eta \sum_{t=1}^{(j-1)H} \langle \mu_\pi^t, \widehat{\ell}^t \rangle \right). \quad (24)$$

- 3: Sample the policy $\pi_j \sim p_j$ for this epoch.
- 4: Execute Algorithm 5 with parameters s^t, π_j, t (note that Algorithm 5 will update t internally).
- 5: **for** All remaining $\mathcal{T}_j \cap [t, T]$ **do**
- 6: Play $a^t = \pi_j(s^t)$, observe loss $\ell^t(s^t, a^t)$ and the next state s^{t+1} . Set

$$\widehat{\ell}^t(s, a) = \mathbb{1}[(s, a) = (s^t, a^t)] \frac{\ell^t(s^t, a^t)}{\sum_{\pi \in \Pi} p_j(\pi) \mu_\pi^t(s, a)}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

where $f(z)$ is the probability density function of z and the second step makes use of the fact that $z + \widehat{\ell}^j$ is still linear in z . Moreover,

$$p_{j+1}(\pi) = \int_z \mathbb{1} \left[\pi = \text{best} \left(\{ \widehat{\ell}^1, \dots, \widehat{\ell}^{jH} \}, z \right) \right] f(z) \, dz.$$

For simplicity, denote $\widetilde{\ell}^j = \{0, \dots, 0, \widehat{\ell}^{(j-1)H+1}, \dots, \widehat{\ell}^{jH}, 0, 0, \dots, 0\} = \{ \mathbb{1}[t \in \mathcal{T}_j] \widehat{\ell}^t \}_{t=1}^T$. Again using the fact that $f(z) = \prod_{h=1}^H \exp(-\eta \|z^h\|_1)$, we have

$$f(z + \widetilde{\ell}^j) = \prod_{t \in \mathcal{T}_j} \exp \left(-\eta \left(\|z^t + \widehat{\ell}^t\|_1 - \|z^t\|_1 \right) \right) f(z),$$

which gives

$$\frac{f(z + \widetilde{\ell}^j)}{f(z)} \in \left[\exp \left(-\eta \sum_{t \in \mathcal{T}_j} \|\widehat{\ell}^t\|_1 \right), \exp \left(\eta \sum_{t \in \mathcal{T}_j} \|\widehat{\ell}^t\|_1 \right) \right]$$

by triangle inequality. Therefore, $p_{j+1}(\pi)/p_j(\pi)$ lies in this interval as well, which is just our claim. \square

D.2 Hedge-Based Inefficient Algorithm (Theorem 8)

In this section, we present our Hedge-based inefficient algorithm for infinite-horizon AMDPs with bandit feedback and known transitions. We still use the same epoching mechanism as Algorithm 6.

For Hedge, which is different from FTPL, we will *explicitly* maintain a distribution $p_j \in \Delta(\Pi)$ over all policies for each epoch, and randomly draw one $\pi_j \sim p_j$ for the j -th epoch. As the distribution p_j can be directly calculated (we do not care about computational efficiency now), we can use importance weighting estimator to estimate the losses. The algorithm is presented in Algorithm 7.

Proof of Theorem 8. As Appendix D.1, we still define $\widetilde{\mathcal{R}}_T$ as Eq. (22). We can still conclude that $\mathcal{R}_T \leq \widetilde{\mathcal{R}}_T + JD^2$. We first show that the importance weighting estimator is indeed unbiased. Notice that the probability of visiting (s, a) at some slot $t \in \mathcal{T}_j$ is exactly $\sum_{\pi \in \Pi} p_j(\pi) \mu_\pi^t(s, a)$, which means, by Lemma 42, we have

$$\mathbb{E} \left[\widehat{\ell}^t(s, a) \middle| \mathcal{F}_{(j-1)H} \right] = \ell^t(s, a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, t \in \mathcal{T}_j, j \in [J].$$

Let $\widetilde{\ell}_j(\pi)$ be the random variable denoting the total loss of policy π for epoch j :

$$\widetilde{\ell}_j(\pi) = \sum_{t \in \mathcal{T}_j} \langle \mu_\pi^t, \widehat{\ell}^t \rangle.$$

So Eq. (24) is just $p_j(\pi) \propto \exp(-\eta \sum_{j'=1}^{j-1} \tilde{\ell}_{j'}(\pi))$. Therefore, by standard properties of Hedge (Lemma 37), for any realization of $\{\tilde{\ell}_j\}_{j \in [J]}$ (and also $\{\hat{\ell}^t\}_{t \in [T]}$), we will have

$$\sum_{j=1}^J \langle p_j, \tilde{\ell}_j \rangle - \sum_{j=1}^J \tilde{\ell}_j(\pi^*) \leq \frac{\ln |\Pi|}{\eta} + \eta \sum_{j=1}^J \sum_{\pi \in \Pi} p_j(\pi) \tilde{\ell}_j^2(\pi). \quad (25)$$

Consider the second term of the right-handed-side. For a fixed $j \in [J]$, it becomes

$$\begin{aligned} \sum_{\pi \in \Pi} p_j(\pi) \tilde{\ell}_j^2(\pi) &= \sum_{\pi \in \Pi} p_j(\pi) \left(\sum_{t \in \mathcal{T}_j} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mu_{\pi}^t(s,a) \hat{\ell}^t(s,a) \right)^2 \\ &\leq H \sum_{\pi \in \Pi} p_j(\pi) \sum_{t \in \mathcal{T}_j} \left(\sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mu_{\pi}^t(s,a) \hat{\ell}^t(s,a) \right)^2 \\ &= H \sum_{\pi \in \Pi} p_j(\pi) \sum_{t \in \mathcal{T}_j} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left(\mu_{\pi}^t(s,a) \hat{\ell}^t(s,a) \right)^2, \end{aligned}$$

where the first inequality made use of Cauchy-Schwartz inequality while the second equality used the fact that $\hat{\ell}^t$ is one-hot. Plugging back into Eq. (25) and taking expectation on both sides,

$$\begin{aligned} &\mathbb{E} \left[\sum_{j=1}^J \langle p_j, \tilde{\ell}_j \rangle - \sum_{j=1}^J \tilde{\ell}_j(\pi^*) \right] \\ &\stackrel{(a)}{\leq} \frac{S \ln A}{\eta} + H \sum_{j=1}^J \mathbb{E} \left[\sum_{\pi \in \Pi} p_j(\pi) \sum_{t \in \mathcal{T}_j} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mu_{\pi}^t(s,a) \left(\hat{\ell}^t(s,a) \right)^2 \middle| \mathcal{F}_{(j-1)H} \right] \\ &\stackrel{(b)}{\leq} \frac{S \ln A}{\eta} + \eta H \sum_{j=1}^J \mathbb{E} \left[\sum_{\pi \in \Pi} p_j(\pi) \sum_{t \in \mathcal{T}_j} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mu_{\pi}^t(s,a) \cdot \frac{1}{\sum_{\pi \in \Pi} p_j(\pi) \mu_{\pi}^t(s,a)} \middle| \mathcal{F}_{(j-1)H} \right] \\ &= \frac{S \ln A}{\eta} + \eta H^2 J S A = \frac{S \ln A}{\eta} + \eta H S A T, \end{aligned}$$

where (a) used $\mu_{\pi}^t(s,a) \leq 1$ and (b) used Lemma 43. Moreover, for the left-hand side, we have

$$\mathbb{E} \left[\sum_{j=1}^J \langle p_j, \tilde{\ell}_j \rangle - \sum_{j=1}^J \tilde{\ell}_j(\pi^*) \right] = \sum_{j=1}^J \mathbb{E} \left[\sum_{\pi \in \Pi} p_j(\pi) \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t - \mu_{\pi^*}^t, \hat{\ell}^t \rangle \middle| \mathcal{F}_{(j-1)H} \right].$$

By using Lemma 42, this is exactly

$$\sum_{j=1}^J \mathbb{E} \left[\sum_{\pi \in \Pi} p_j(\pi) \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi}^t - \mu_{\pi^*}^t, \hat{\ell}^t \rangle \middle| \mathcal{F}_{(j-1)H} \right] = \mathbb{E} \left[\sum_{j=1}^J \sum_{t \in \mathcal{T}_j} \langle \mu_{\pi_j}^t - \mu_{\pi^*}^t, \ell^t \rangle \right] = \widetilde{\mathcal{R}}_T.$$

Therefore, we will have

$$\mathcal{R}_T \leq \widetilde{\mathcal{R}}_T + J D^2 \leq \frac{S \ln A}{\eta} + \eta H S A T + J D^2,$$

which gives $\mathcal{R}_T = \widetilde{\mathcal{O}}(S^{2/3} A^{1/3} D^{2/3} T^{2/3})$ when picking $J = S^{2/3} A^{1/3} D^{-4/3} T^{2/3}$ and $\eta = S^{1/3} A^{-1/3} D^{-2/3} T^{-2/3}$. \square

Lemma 37 (Property of Hedge). *Suppose that we are using Hedge for T -round online learning problem that has K actions, i.e., at time slot $t \in [T]$, picking i_t according to the probability distribution $p_t \in \Delta([K])$ which is defined as:*

$$p_t(i) \propto \exp \left(-\eta \sum_{\tau=1}^{t-1} \ell_{\tau}(i) \right), \quad \forall i \in [K], t \in [T],$$

where $\ell_t(i) \geq 0$ is the non-negative loss associated with action i at time slot t . Then, for all $i^* \in [K]$, we have

$$\sum_{t=1}^T (\langle p_t, \ell_t \rangle - \ell_t(i^*)) \leq \frac{\ln K}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^K p_t(i) \ell_t^2(i).$$

Note that here we are considering non-randomized loss functions here.

Proof. For simplicity, define $L_t(i)$ as $\sum_{\tau=1}^t \ell_\tau(i)$. Let

$$\Phi_t = \frac{1}{\eta} \ln \left(\sum_{i=1}^K \exp(-\eta L_t(i)) \right),$$

then

$$\begin{aligned} \Phi_t - \Phi_{t-1} &= \frac{1}{\eta} \ln \left(\frac{\sum_{i=1}^K \exp(-\eta L_t(i))}{\sum_{i=1}^K \exp(-\eta L_{t-1}(i))} \right) \\ &= \frac{1}{\eta} \ln \left(\sum_{i=1}^K p_t(i) \exp(-\eta \ell_t(i)) \right) \\ &\stackrel{(a)}{\leq} \frac{1}{\eta} \ln \left(\sum_{i=1}^K p_t(i) (1 - \eta \ell_t(i) + \eta^2 \ell_t^2(i)) \right) \\ &= \frac{1}{\eta} \ln \left(1 - \eta \langle p_t, \ell_t \rangle + \eta^2 \sum_{i=1}^K p_t(i) \ell_t^2(i) \right) \\ &\stackrel{(b)}{\leq} -\langle p_t, \ell_t \rangle + \eta \sum_{i=1}^K p_t(i) \ell_t^2(i), \end{aligned}$$

where (a) used $\exp(-x) \leq 1 - x + x^2$ for all $x \geq 0$ and (b) used $\ln(1+x) \leq x$. Therefore, summing over t gives

$$\begin{aligned} \sum_{t=1}^T \langle p_t, \ell_t \rangle &\leq \Phi_0 - \Phi_T + \eta \sum_{t=1}^T \sum_{i=1}^N p_t(i) \ell_t^2(i) \\ &\leq \frac{\ln N}{\eta} - \frac{1}{\eta} \ln(\exp(-\eta L_T(i^*))) + \eta \sum_{t=1}^T \sum_{i=1}^N p_t(i) \ell_t^2(i) \\ &\leq \frac{\ln N}{\eta} + L_T(i^*) + \eta \sum_{t=1}^T \sum_{i=1}^N p_t(i) \ell_t^2(i). \end{aligned}$$

Moving $L_T(i^*)$ to the left-handed-side then gives our conclusion. \square

E Auxiliary Lemmas

E.1 Geometric Re-sampling Properties

In this section, we list two properties of the Geometric Re-sampling estimator (Neu and Bartók, 2013) that we used in the analysis. For the sake of completeness, we also include their proofs here.

Lemma 38 (Neu and Bartók (2013, Lemma 1)). *Consider the Geometric Re-sampling estimator*

$$\hat{\ell}_k^h(s, a) = \mathbb{1}[(s_k^h, a_k^h) = (s, a)] M_k^h(s, a) \ell_k^h(s, a). \quad (26)$$

Let $\Pr\{(s_k^h, a_k^h) = (s, a) \mid \mathcal{F}_{k-1}\} = q_k^h(s, a)$. Suppose that the probability of visiting (s, a) in the re-sampling process is also $q_k^h(s, a)$, then we have

$$\mathbb{E} \left[\hat{\ell}_k^h(s, a) \mid \mathcal{F}_{k-1} \right] = (1 - (1 - q_k^h(s, a))^L) \ell_k^h(s, a).$$

Proof. By direct calculation, we have

$$\begin{aligned}\mathbb{E} [M_k^h(s, a) | \mathcal{F}_{k-1}, (s_k^h, a_k^h) = (s, a)] &= \sum_{n=1}^{\infty} n(1-q)^{n-1}q - \sum_{n=L}^{\infty} (n-L)(1-q)^{n-1}q \\ &= (1 - (1-q)^L) \sum_{n=1}^{\infty} n(1-q)^{n-1}q = \frac{1 - (1-q)^L}{q}.\end{aligned}$$

So we have

$$\begin{aligned}\mathbb{E} [\widehat{\ell}_k^h(s, a) | \mathcal{F}_{k-1}] &= \Pr\{(s_k^h, a_k^h) = (s, a) | \mathcal{F}_{k-1}\} \ell_k^h(s, a) \mathbb{E} [M_k^h(s, a) | \mathcal{F}_{k-1}, (s_k^h, a_k^h) = (s, a)] \\ &= (1 - (1 - q_k^h(s, a))^L) \ell_k^h(s, a),\end{aligned}$$

as desired. \square

Lemma 39. *For the Geometric Re-sampling estimator as defined in the previous lemma, we have*

$$\mathbb{E} [(\widehat{\ell}_k^h(s, a))^2 | \mathcal{F}_{k-1}] \leq \frac{2(\ell_k^h(s, a))^2}{q_k^h(s, a)}.$$

Proof. By definition, write

$$\begin{aligned}\mathbb{E} [(\widehat{\ell}_k^h(s, a))^2 | \mathcal{F}_{k-1}] &= \mathbb{E} [\mathbb{1}\{(s_k^h, a_k^h) = (s, a)\}^2 (\ell_k^h(s, a))^2 (M_k^h(s, a))^2] \\ &= \Pr\{(s_k^h, a_k^h) = (s, a)\} (\ell_k^h(s, a))^2 \mathbb{E} [(M_k^h(s, a))^2 | \mathcal{F}_{k-1}, (s_k^h, a_k^h) = (s, a)].\end{aligned}\quad (27)$$

Simply write $q_k^h(s, a)$ as q . Note that $M_k^h(s, a) = \min\{L, \text{Geo}(q)\}$, it is stochastically dominated by the geometric distribution with parameter q , whose second moment is bounded by

$$\begin{aligned}\mathbb{E} [(M_k^h(s, a))^2 | \mathcal{F}_{k-1}, (s_k^h, a_k^h) = (s, a)] \\ \leq \mathbb{E}[(\text{Geo}(q))^2] = \text{Var}(\text{Geo}(q)) + (\mathbb{E}[\text{Geo}(q)])^2 = \frac{1-q}{q^2} + \frac{1}{q} \leq \frac{2}{q^2},\end{aligned}\quad (28)$$

which means

$$\mathbb{E} [(\widehat{\ell}_k^h(s, a))^2 | \mathcal{F}_{k-1}] \leq q(\ell_k^h(s, a))^2 \frac{2}{q^2} = \frac{2(\ell_k^h(s, a))^2}{q},$$

as claimed. \square

Corollary 40. *Still consider the GR estimator defined in Eq. (26). Suppose that $\Pr\{(s_k^h, a_k^h) = (s, a) | \mathcal{F}_{k-1}\} = \widehat{q}_k^h(s, a)$ and the probability of visiting (s, a) in each re-sampling procedure is $q_k^h(s, a)$ (where $\widehat{q}_k^h(s, a) \neq q_k^h(s, a)$). We then have*

$$\mathbb{E} [\widehat{\ell}_k^h(s, a) | \mathcal{F}_{k-1}] = \frac{\widehat{q}_k^h(s, a)}{q_k^h(s, a)} (1 - (1 - q_k^h(s, a))^L) \ell_k^h(s, a).$$

Proof. The calculation of $\mathbb{E}[M_k^h(s, a) | \mathcal{F}_{k-1}, (s_k^h, a_k^h) = (s, a)]$ is the same as the one in Lemma 38. Therefore,

$$\begin{aligned}\mathbb{E} [\widehat{\ell}_k^h(s, a) | \mathcal{F}_{k-1}] &= \Pr\{(s_k^h, a_k^h) = (s, a) | \mathcal{F}_{k-1}\} \ell_k^h(s, a) \mathbb{E} [M_k^h(s, a) | \mathcal{F}_{k-1}, (s_k^h, a_k^h) = (s, a)] \\ &= \frac{\widehat{q}_k^h(s, a)}{q_k^h(s, a)} (1 - (1 - q_k^h(s, a))^L) \ell_k^h(s, a),\end{aligned}$$

as claimed. \square

Corollary 41. *Suppose the same condition as the previous corollary, i.e., still considering the GR estimator defined in Eq. (26) where $\Pr\{(s_k^h, a_k^h) = (s, a) | \mathcal{F}_{k-1}\} = \widehat{q}_k^h(s, a)$ and the probability of visiting (s, a) in each re-sampling procedure is $q_k^h(s, a)$. We have*

$$\mathbb{E} [(\widehat{\ell}_k^h(s, a))^2 | \mathcal{F}_{k-1}] \leq \frac{2(\ell_k^h(s, a))^2 \widehat{q}_k^h(s, a)}{q_k^h(s, a) q_k^h(s, a)}.$$

Proof. Still decompose the variance as Eq. (27). Still write $\hat{q}_k^h(s, a)$ as \hat{q} and $q_k^h(s, a)$ as q . Then we still have $M_k^h(s, a) = \min\{L, \text{Geo}(q)\}$, which gives $\mathbb{E}[(M_k^h(s, a))^2 \mid \mathcal{F}_{k-1}, (s_k^h, a_k^h) = (s, a)] \leq \frac{2}{q^2}$ by Eq. (28). Therefore,

$$\mathbb{E} \left[(\hat{\ell}_k^h(s, a))^2 \mid \mathcal{F}_{k-1} \right] \leq \hat{q} (\ell_k^h(s, a))^2 \frac{2}{q^2} = \frac{2(\ell_k^h(s, a))^2 \hat{q}_k^h(s, a)}{q_k^h(s, a) \hat{q}_k^h(s, a)},$$

as claimed. \square

E.2 Importance Weighting Properties

Lemma 42. *For the Importance Weighting estimator*

$$\hat{\ell}^t(s, a) = \mathbb{1}[(s^t, a^t) = (s, a)] \frac{\ell^t(s^t, a^t)}{\Pr\{(s^t, a^t) = (s, a) \mid \mathcal{F}\}}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

where \mathcal{F} is a filtration, we will have

$$\mathbb{E}[\hat{\ell}^t(s, a) \mid \mathcal{F}] = \ell^t(s, a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

Proof. For simplicity, denote $q^t(s, a) = \Pr\{(s^t, a^t) = (s, a) \mid \mathcal{F}\}$. Then

$$\mathbb{E}[\hat{\ell}^t(s, a) \mid \mathcal{F}] = q^t(s, a) \cdot \frac{\ell^t(s, a)}{q^t(s, a)} = \ell^t(s, a)$$

for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. \square

Lemma 43. *For the same Importance Weighting Estimator, we will have*

$$\mathbb{E} \left[(\hat{\ell}^t(s, a))^2 \mid \mathcal{F} \right] = \frac{(\ell^t(s, a))^2}{q^t(s, a)}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

where $q^t(s, a) \triangleq \Pr\{(s^t, a^t) = (s, a) \mid \mathcal{F}\}$.

Proof. Direct calculation gives $\mathbb{E} \left[(\hat{\ell}^t(s, a))^2 \mid \mathcal{F} \right] = q^t(s, a) \cdot \left(\frac{\ell^t(s, a)}{q^t(s, a)} \right)^2 = \frac{(\ell^t(s, a))^2}{q^t(s, a)}, \forall (s, a)$. \square

E.3 Auxiliary Lemmas for Error Terms

In this section, we present two lemmas that will play an important role when bounding the error terms (as used in Lemmas 10, 18 and 34).

Lemma 44 (Wang and Dong (2020, Fact 2)). *Let X_1, X_2, \dots, X_n be i.i.d. random variables drawn from $\text{Exp}(\eta)$ which is the exponential distribution, then*

$$\mathbb{E} \left[\max_{1 \leq i \leq n} X_i \right] \leq \frac{1 + \ln n}{\eta}.$$

Lemma 45 (Generalization of Syrgkanis et al. (2016, Lemma 8)). *Let $\{z^t \in \mathbb{R}^d\}_{t=1}^T$ be a sequence of d -dimensional random variable such that $z_i^t \sim \text{Laplace}(\eta)$ for all $i \in [m]$ and $t \in [T]$. Let X be a set of sequences of the form $\{x^t \in [0, 1]^d\}_{t=1}^T$. As long as $\ln|X| < dT$, we have*

$$\mathbb{E}_z \left[\max_{x \in X} \sum_{t=1}^T \langle x^t, z^t \rangle \right] - \mathbb{E}_z \left[\min_{x \in X} \sum_{t=1}^T \langle x^t, z^t \rangle \right] \leq \frac{10}{\eta} \sqrt{dT \ln|X|}.$$

Proof. Note that the key difference between this theorem and Syrgkanis et al. (2016, Lemma 8) is that, their theorem assumed a binary decision set, i.e., $x_i^t \in \{0, 1\}$ instead of $[0, 1]$. However, their proof still holds with only a little modification. The first step is still noticing that the distribution of Laplace

random variables is symmetric around 0, so we only need to bound $2 \mathbb{E}_z \left[\max_{x \in X} \sum_{t=1}^T \langle x^t, z^t \rangle \right]$, which is bounded by, for any $\lambda \geq 0$,

$$\begin{aligned}
\mathbb{E}_z \left[\max_{x \in X} \sum_{t=1}^T \langle x^t, z^t \rangle \right] &= \frac{1}{\lambda} \ln \left(\exp \left(\mathbb{E}_z \left[\max_{x \in X} \lambda \sum_{t=1}^T \langle x^t, z^t \rangle \right] \right) \right) \\
&\leq \frac{1}{\lambda} \ln \left(\mathbb{E}_z \left[\max_{x \in X} \exp \left(\lambda \sum_{t=1}^T \langle x^t, z^t \rangle \right) \right] \right) \\
&\leq \frac{1}{\lambda} \ln \left(\sum_{x \in X} \mathbb{E}_z \left[\exp \left(\lambda \sum_{t=1}^T \langle x^t, z^t \rangle \right) \right] \right) \\
&\leq \frac{1}{\lambda} \ln \left(\sum_{x \in X} \prod_{t=1}^T \mathbb{E}_z \left[\exp (\lambda \langle x^t, z^t \rangle) \right] \right) \\
&= \frac{1}{\lambda} \ln \left(\sum_{x \in X} \prod_{t=1}^T \mathbb{E}_z \left[\exp \left(\lambda \sum_{i=1}^d x_i^t z_i^t \right) \right] \right) \\
&\leq \frac{1}{\lambda} \ln \left(\sum_{x \in X} \prod_{t=1}^T \prod_{i=1}^d (\mathbb{E}_z [\exp (\lambda z_i^t)])^{x_i^t} \right),
\end{aligned}$$

where the last step used the fact that $x_i^t \leq 1$ (and thus $y^{x_i^t}$ is a concave function in y). Furthermore, by using the fact that $\mathbb{E}_z [\exp (\lambda z_i^t)]$ is just the moment generating function of Laplace random variables evaluated at λ , it is just $(1 - \frac{\lambda^2}{\eta^2})^{-1}$ as long as $\lambda < \eta$. As it is always larger than 1, we can directly bound

$$\begin{aligned}
\mathbb{E}_z \left[\max_{x \in X} \sum_{t=1}^T \langle x^t, z^t \rangle \right] &\leq \frac{1}{\lambda} \ln \left(\sum_{x \in X} \prod_{t=1}^T \sum_{i=1}^d (\mathbb{E}_z [\exp (\lambda z_i^t)])^{x_i^t} \right) \\
&\leq \frac{1}{\lambda} \ln \left(\sum_{x \in X} \prod_{t=1}^T \prod_{i=1}^d \left(\frac{1}{1 - \frac{\lambda^2}{\eta^2}} \right)^{x_i^t} \right) \leq \frac{1}{\lambda} \ln \left(\sum_{x \in X} \prod_{t=1}^T \prod_{i=1}^d \frac{1}{1 - \frac{\lambda^2}{\eta^2}} \right) \\
&= \frac{1}{\lambda} \ln \left(|X| \left(\frac{1}{1 - \frac{\lambda^2}{\eta^2}} \right)^{dT} \right) = \frac{1}{\lambda} \ln |X| + \frac{dT}{\lambda} \ln \left(\frac{1}{1 - \frac{\lambda^2}{\eta^2}} \right).
\end{aligned}$$

By using the fact that $\frac{1}{1-x} \leq \exp(2x)$ for all $x \leq \frac{1}{4}$, as long as $\lambda \leq \frac{\eta}{2}$, we will have

$$\mathbb{E}_z \left[\max_{x \in X} \sum_{t=1}^T \langle x^t, z^t \rangle \right] \leq \frac{1}{\lambda} \ln \left(|X| \left(\frac{1}{1 - \frac{\lambda^2}{\eta^2}} \right)^{dT} \right) = \frac{1}{\lambda} \ln |X| + \frac{2dT}{\lambda} \frac{\lambda^2}{\eta^2}.$$

By picking $\lambda = \frac{\eta \sqrt{\ln |X|}}{2\sqrt{dT}} < \frac{\eta}{2}$ (according to the assumption that $\ln |X| < dT$) gives the bound $\frac{5}{\eta} \sqrt{dT \ln |X|}$, which is what we want. \square