
Supplement: Non-Asymptotic Gap-Dependent Regret Bounds for Tabular MDPs

Max Simchowitz
UC Berkeley
msimchow@berkeley.edu

Kevin Jamieson
University of Washington
jamieson@cs.washington.edu

Abstract

This paper establishes that optimistic algorithms attain gap-dependent and non-asymptotic logarithmic regret for episodic MDPs. In contrast to prior work, our bounds do not suffer a dependence on diameter-like quantities or ergodicity, and smoothly interpolate between the gap dependent logarithmic-regret, and the $\tilde{O}(\sqrt{HSAT})$ -minimax rate. The key technique in our analysis is a novel “clipped” regret decomposition which applies to a broad family of recent optimistic algorithms for episodic MDPs.

1 Introduction

Reinforcement learning (RL) is a powerful paradigm for modeling a learning agent’s interactions with an unknown environment, in an attempt to accumulate as much reward as possible. Because of its flexibility, RL can encode such a vast array of different problem settings - many of which are entirely intractable. Therefore, it is crucial to understand what conditions enable an RL agent to effectively learn about its environment, and to account for the success of RL methods in practice.

In this paper, we consider tabular Markov decision processes (MDPs), a canonical RL setting where the agent seeks to learn a *policy* mapping discrete states $x \in \mathcal{S}$ to one of finitely many actions $a \in \mathcal{A}$, in an attempt to maximize cumulative reward over an episode horizon H . We shall study the *regret* setting, where the learner plays a policy π_k for a sequence of episodes $k = 1, \dots, K$, and suffers a regret proportional to the average sub-optimality of the policies π_1, \dots, π_K .

In recent years, the vast majority of literature has focused on obtaining *minimax* regret bounds that match the worst-case dependence on the number states $|\mathcal{S}|$, actions $|\mathcal{A}|$, and horizon length H ; namely, a cumulative regret of $\sqrt{H|\mathcal{S}||\mathcal{A}|T}$, where $T = KH$ denotes the total number of rounds of the game [1]. While these bounds are succinct and easy to interpret, they paint an overly pessimistic account of the complexity of these problems, and do not elucidate the favorable structural properties of which a learning agent can hope to take advantage.

The earlier literature, on the other hand, establishes a considerable more favorable regret of the form $C \log T$, where C is an instance-dependent constant given in terms of the *sub-optimality gaps* associated with each action at a given state, defined as

$$\text{gap}_\infty(x, a) = \mathbf{V}^{\pi^*}(x) - \mathbf{Q}^{\pi^*}(x, a), \quad (1)$$

where \mathbf{V}^{π^*} and \mathbf{Q}^{π^*} denote the value and Q functions for an optimal policy π^* , and the subscript- ∞ denotes these bounds hold for a non-episodic, infinite horizon setting. Depending on the constant C , the regret $C \log T$ can yield a major improvement over the \sqrt{T} minimax scaling. Unfortunately, these analyses are asymptotic in nature, and only take effect after a large number of rounds, depending on other potentially-large, highly-conservative, or difficult-to-verify problem-dependent quantities such as hitting times or measures of uniform ergodicity [8, 15, 11].

To fully account for the empirical performance of RL algorithms, we seek regret bounds which take advantage of favorable problem instances, but apply in *finite time* and for practically realistic numbers of rounds T .

Contributions: As a first step in this direction, [16] introduced a novel algorithm called EULER, which enjoys reduced dependence on the episode horizon H for favorable instances, while maintaining the same worst-case dependence for other parameters in their analysis as in [1].

In this paper, we take the next step by demonstrating that a common class of algorithms for solving MDPs, based on the *optimism* principle, attains gap-dependent, problem-specific bounds similar to those previously found only in the asymptotic regime. For concreteness, we specialize our analysis to a minor modification of the EULER algorithm we call StrongEuler; as we explain in Section 3, our analysis extends more broadly to other optimistic algorithms as well. We show that

- For any episodic MDP \mathcal{M} , StrongEuler enjoys a high probability regret bound of $C_{\mathcal{M}} \log(1/\delta)$ for all rounds $T \geq 1$, where the constant $C_{\mathcal{M}}$ depends on the sub-optimality gaps between actions at different states, as well as the horizon length, and contains an additive almost-gap-independent term that scales as $AS^2 \text{poly}(H)$ (Corollary 2.1).

Unlike previous gap-dependent regret bounds,

- The constant $C_{\mathcal{M}}$ does not suffer worst-case dependencies on other problem dependent quantities such as mixing times, hitting times or measures of ergodicity. However, the constant $C_{\mathcal{M}}$ *does* take advantage of *benign* problem instances (Definition 2.2).
- The regret bound of $C_{\mathcal{M}} \log(1/\delta)$ is valid for any total number of rounds $T \geq 1$. Selecting $\delta = 1/T$, this implies a *non-asymptotic* expected regret bound of $C_{\mathcal{M}} \log T^1$.
- The regret of StrongEuler interpolates between instance-dependent regret $C_{\mathcal{M}} \log T$ and mini-max regret $\tilde{O}(\sqrt{H|\mathcal{S}||\mathcal{A}|T})$, the latter of which may be sharper for smaller T (Theorem 2.4). Following [16], this dependence on H may also be refined for benign instances.

Lastly, while the StrongEuler algorithm affords sharper regret bounds than past algorithms, our analysis techniques extend more generally to other optimism based algorithms:

- We introduce a novel “clipped” regret decomposition (Proposition 3.1) which applies to a broad family of optimistic algorithms, including the algorithms analyzed in [16, 6, 5, 9, 1].
- Following our analysis of StrongEuler, the clipped regret decomposition can establish analogous gap-dependent $\log T$ -regret bounds for many of the algorithms mentioned above.

What is $C_{\mathcal{M}}$? In many settings, we show that $C_{\mathcal{M}}$ is dominated by an analogue to the sum over the reciprocals of the gaps defined in (1). This is known to be optimal for non-dynamic MDP settings like contextual bandits, and we prove a lower bound (Proposition 2.2) which shows that this is unimprovable for general MDPs as well. Furthermore, building on [16], we show this adapts to problems with additional structure, yielding, e.g., a horizon H -free bound for contextual bandits.

However, our gap-dependent bound also suffers from a certain dependence on the *smallest nonzero gap* gap_{\min} (see Definition 2.1), which may dominate in some settings. We prove a lower bound (Theorem 2.3) which shows that optimistic algorithms in the recent literature - including StrongEuler - necessarily suffer a similar term in their regret. We believe this insight will motivate new algorithms for which this dependence can be removed, leading to new design principles and actionable insights for practitioners. Finally, our regret bound incurs an (almost) gap-independent burn-in term, which is standard for optimistic algorithms, and which we believe is an exciting direction of research to remove.

Altogether, we believe that the results in our paper serve as a preliminary but significant step to attaining sharp, instance-dependent, and *non-asymptotic* bounds for tabular MDPs, and hope that our analysis will guide the design of future algorithms that attain these bounds.

1.1 Related Work

Like the multi-armed bandit setting, regret bounds for MDP algorithms have been characterized both in *gap-independent* forms that rely solely on $S := |\mathcal{S}|$, $A := |\mathcal{A}|$, H , T , and in *gap-dependent* forms

¹By this, we mean that for any *fixed* $T \geq 1$, one can attain $C_{\mathcal{M}} \log T$ regret. Extending the bound to anytime regret is left to future work

which take into account the gaps (1), as well as other instance-specific properties of the rewards and transition probabilities.

Finite Sample Bounds, Gap-Independent Bounds: A number of notable recent works give undiscounted regret bounds for finite-horizon, tabular MDPs, nearly all of them relying on the principle of optimism which we describe in Section 3 [4, 1, 5, 9, 16]. Many of the more recent works [1, 16, 6] attain a regret of \sqrt{HSAT} , matching the known lower bound of \sqrt{HSAT} established in [12, 8, 4]. As mentioned above, the EULER algorithm of [16] attains the minimax rates and simultaneously enjoys a reduced dependence on H in benign problem instances, such as the contextual bandits setting where the transition probabilities do not depend on the current state or learners actions, or when the total cumulative rewards over any roll-out are bounded by 1 in magnitude.

Diameter Dependent Bounds: In the setting of infinite horizon MDPs with discounted regret, many previous works have established logarithmic regret bounds of the form $C(\mathcal{M}) \log T$, where $C(\mathcal{M})$ is a constant depending on the underlying MDP. Notably, [8] give an algorithm which attains a $\tilde{\mathcal{O}}(\sqrt{D^2 S^2 AT})$ gap-independent regret, and an $\tilde{\mathcal{O}}(\frac{D^2 S^2 A}{\text{gap}_*} \log(T))$ gap-dependent regret bound, where gap_* is the difference between the mean infinite-horizon reward of π_* and the next-best stationary policy, and where D denotes the maximum expected traversal time between any two states x, x' , under the policy which attains the minimal traversal time between those two states. We note that if $\text{gap}_\infty(x, a)$ denotes the sub-optimality of any action a at state x as in (1), then $\text{gap}_* \leq \min_{x,a} \text{gap}_\infty(x, a)$. The bounds in this work, on the other hand, depend on an average over inverse gaps, rather than a worst case. Moreover, the diameter D can be quite large when there exist difficult-to-access states. We stress that the bound due to [8] is non-asymptotic, but the bound in terms of gap_* dependences other worst-case quantities measuring ergodicity.

Asymptotic Bounds: Prior to [8], and building on the bounds of [3], [15] presented bounds in terms of a diameter-related quantity $\bar{D} \geq D$, which captures the minimal hitting time between states when restricted to optimal policies. [15] prove that their algorithm enjoys a regret² of $\sum_{(s,a) \in \text{CRIT}} \frac{\bar{D}^2}{\text{gap}_\infty(x,a)} \log(T)$ asymptotically in T where CRIT contains those sub-optimal state-action pairs (x, a) such that a can be made to be the unique, optimal action at x by replacing $p(s'|s, a)$ with some other vector on the \mathcal{S} -simplex. Recently, [11] present per-instance lower bounds for both structured and unstructured MDPs, which apply to any algorithm which enjoys sub-linear regret on any problem instance, and an algorithm which matches these bounds asymptotically. This bound replaces \bar{D}^2 with \bar{H}^2 , where \bar{H} denotes the range of the bias functions, an analogue of H for the non-episodic setting [2]. We further stress that whereas the logarithmic regret bounds of [8] hold for finite time with polynomial dependence on the problem parameters, the number of episodes needed for the bounds of [3, 15, 11] to hold may be exponentially large, and depend on additional, pessimistic problem-dependent quantities (e.g. a uniform hitting time in Proposition 29 in [14]).

Novelty of this work: The major contribution of our work is showing problem-dependent $\log(T)$ regret bounds which i) attain a refined dependence on the gaps, as in [15], ii) apply in finite time after a burn-in time only polynomial in S, A, H and the gaps, iii) depend only on H and not on the diameter D (and thus, are not adversely affected by difficult to access states), and iv) smoothly interpolate between $\log T$ regret and the minimax \sqrt{HSAT} rate attained by [1] et seq.

1.2 Problem Setting, Notation, and Organization

Episodic MDP: A *stationary*, episodic MDP is a tuple $\mathcal{M} := (\mathcal{S}, \mathcal{A}, H, r, p, p_0, R)$, where for each $x \in \mathcal{S}, a \in \mathcal{A}$ we have that $R(x, a) \in [0, 1]$ is a random reward with expectation $r(x, a)$, $p : \mathcal{S} \times \mathcal{A} \rightarrow \Delta^{\mathcal{S}}$ denotes transition probabilities, $p_0 \in \Delta^{\mathcal{S}}$ is an initial distribution over states, and H is the horizon, or length of the episode. A policy π is a sequence of mappings $\pi_h : \mathcal{S} \rightarrow \mathcal{A}$. For our given MDP \mathcal{M} , we let \mathbb{E}^π and \mathbb{P}^π denote the expectation and probability operator with respect to the law of sequence $(x_1, a_1), \dots, (x_H, a_H)$, where $x_1 \sim p_0, a_h = \pi_h(x_h), x_{h+1} \sim p(x_h, a_h)$. We define the *value* of π as $\mathbf{V}_0^\pi := \mathbb{E}^\pi \left[\sum_{h=1}^H r(x_h, a_h) \right]$ and for $h \in [H]$ and $x \in \mathcal{S}$, $\mathbf{V}_h^\pi(x) := \mathbb{E}^\pi \left[\sum_{h' \geq h} r(x_{h'}, a_{h'}) \mid x_h = x \right]$, which we identify with a vector in $\mathbb{R}^{\mathcal{S}}$. We define the associated Q-function $\mathbf{Q}^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, $\mathbf{Q}_h^\pi(x, a) := r(x, a) + p(x, a)^\top \mathbf{V}_{h+1}^\pi$, so that

²[15] actually presents a bound of the form $\frac{\bar{D}^2 S A}{\min_{(s,a) \in \text{CRIT}} \text{gap}_\infty(x,a)} \log(T)$ but it is straightforward to extract the claimed form from the proof.

$\mathbf{Q}_h^\pi(x, \pi_h(x)) = \mathbf{V}_h^\pi(x)$. We denote the *set* of optimal policies $\pi^* := \arg \max_\pi \mathbf{V}_0^\pi$, and let $\pi_h^*(x) := \{a : \pi_h(x) = a, \pi \in \pi^*\}$ denote the set of optimal actions. Lastly, given any optimal $\pi \in \pi^*$, we introduce the shorthand $\mathbf{V}_h^* = \mathbf{V}_h^\pi$ and $\mathbf{Q}_h^* = \mathbf{Q}_h^\pi$, where we note that even when π is not unique, \mathbf{V}_h^* and \mathbf{Q}_h^* do not depend on the choice of optimal policy.

Episodic Regret: We consider a game that proceeds in rounds $k = 1, \dots, K$, where at each state an algorithm Alg selects a policy π_k , and observes a roll out $(x_1, a_1), \dots, (x_H, a_H) \sim \mathbb{P}^{\pi_k}$. The goal is to minimize the cumulative simple regret, defined as $\text{Regret}_K := \sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}$.

Notation and Organization: For $n \in \mathbb{N}$, we define $[n] = \{1, \dots, n\}$. For two expressions f, g that are functions of any problem-dependent variables of \mathcal{M} , we say $f \lesssim g$ ($f \gtrsim g$, respectively) if there exists a universal constant $c > 0$ independent of \mathcal{M} such that $f \leq cg$ ($f \geq cg$, respectively). \lesssim will denote an informal, approximate inequality. Section 2 presents our main results, and Section 3 sketches the proof and highlights the novelty of our techniques. All references to the appendix refer to the appendix of the supplement. All formal proofs, and many rigorous statement of results, are deferred to the appendix, whose organization and notation are described at length in Appendix A.

1.3 Optimistic Algorithms

Lastly, we introduce *optimistic algorithms* which select a policy which is optimal for an overestimated, or *optimistic*, estimate of the true Q -function, \mathbf{Q}^* .

Definition 1.1 (Optimistic Algorithm). We say that an algorithm Alg satisfies *optimism* if, for each round $k \in [K]$ and stage $h \in [H]$, it constructs an *optimistic* Q -function $\overline{\mathbf{Q}}_{k,h}(x, a)$ and policy $\pi_k = (\pi_{k,h})$ satisfying $\forall x, a : \overline{\mathbf{Q}}_{k,H+1}(x, a) = 0, \overline{\mathbf{Q}}_{k,h}(x, a) \geq \mathbf{Q}_h^*(x, a)$, and $\pi_{k,h}(x) \in \arg \max_a \overline{\mathbf{Q}}_{k,h}(x, a)$. The associated *optimistic value function* is $\overline{\mathbf{V}}_{k,h}(x) := \overline{\mathbf{Q}}_{k,h}(x, \pi_{k,h}(x))$.

We shall colloquially refer to an algorithm as *optimistic* if it satisfies optimism with high probability. Optimism has become the dominant approach for learning finite-horizon MDPs, and all recent low-regret algorithms are optimistic [5, 6, 1, 16, 9]. In *model-based* algorithms, the overestimates $\overline{\mathbf{Q}}_{k,h}$ are constructed recursively as $\overline{\mathbf{Q}}_{k,h}(x, a) = \widehat{r}_k(x, a) + \widehat{p}_k(x, a)^\top \overline{\mathbf{V}}_{k,h+1} + \mathbf{b}_{k,h}(x, a)$, where $\widehat{r}_k(x, a)$ and $\widehat{p}_k(x, a)$ are empirical estimates of the mean rewards and transition probabilities, and $\mathbf{b}_{k,h}(x, a) \geq 0$ is a confidence *bonus* designed to ensure that $\overline{\mathbf{Q}}_{k,h}(x, a) \geq \mathbf{Q}^*(x, a)$. Letting $n_k(x, a)$ denote the total number of times a given state-action pair is visited, a simple bonus $\mathbf{b}_{k,h}(x, a) \asymp \sqrt{\frac{H \log(SAHK/\delta)}{n_k(x, a)}}$ suffices to induce optimism, yielding the UCBVI-CH algorithm of [1]. This leads to an episodic regret bound of $\sqrt{H^2 SAT}$, a factor of \sqrt{H} greater than the minimax rate. More refined bonuses based on the ‘‘Bernstein trick’’ achieve the optimal H -dependence [1], and the EULER algorithm of [16] adopts further refinements to replace worst-case H dependence with more adaptive quantities. The StrongEuler algorithm considered in this work applies similarly adaptive bonuses, but our analysis extends to all aforementioned bonus configurations. We remark that there are also *model-free* optimistic algorithms based on Q -learning (see, e.g. [9]) that construct overestimates in a slightly different fashion. While our main technical contribution, the clipped regret decomposition (Proposition 3.1), applies to *all* optimistic algorithms, our subsequent analysis is tailored to model-based approaches, and may not extend straightforwardly to Q -learning methods.

2 Main Results

Logarithmic Regret for Optimistic Algorithms: We now state regret bounds that describe the performance of StrongEuler, an instance of the model-based, optimistic algorithms described above. StrongEuler is based on carefully selected bonuses from [16], and formally instantiated in Algorithm 1 in Appendix E. We emphasize that other optimistic algorithms enjoy similar regret bounds, but we restrict our analysis to StrongEuler to attain the sharpest H -dependence. The key quantities at play are the *suboptimality-gaps* between the Q -functions:

Definition 2.1 (Suboptimality Gaps). For $h \in [H]$, define the stage-dependent suboptimality gap $\text{gap}_h(x, a) := \mathbf{V}_h^*(x) - \mathbf{Q}_h^*(x, a)$, as well as the minimal stage-independent gap $\text{gap}(x, a) := \min_h \text{gap}_h(x, a)$, and the minimal gap $\text{gap}_{\min} := \min_{x,a,h} \{\text{gap}_h(x, a) : \text{gap}_h(x, a) > 0\}$.

Note that any optimal $a^* \in \pi_h^*(x)$ satisfies the Bellman equation $\mathbf{Q}_h^*(x, a^*) = \max_a \mathbf{Q}_h^*(x, a) = \mathbf{V}_h^*(x)$, and thus $\text{gap}_h(x, a^*) = 0$ iff $a^* \in \pi_h^*(x)$. Following [16], we consider two illustrative benign problem settings which afford an improved dependence on the horizon H :

Definition 2.2 (Benign Settings). We say that an MDP \mathcal{M} is a *contextual bandit instance* if $p(x'|x, a)$ does not depend on x or a . An MDP \mathcal{M} has \mathcal{G} -bounded rewards if, for any policy π , $\sum_{h=1}^H R(x_h, a_h) \leq \mathcal{G}$ holds with probability 1 over trajectories $((x_h, a_h)) \sim \mathbb{P}^\pi$.

Lastly, we define \mathcal{Z}_{opt} as the set of pairs (x, a) for which a is optimal at x for some stage $h \in [H]$: $\mathcal{Z}_{\text{opt}} := \{(x, a) : \exists h \in [H] \text{ with } a \in \pi_h^*(x)\}$ and its complement $\mathcal{Z}_{\text{sub}} := \mathcal{S} \times \mathcal{A} - \mathcal{Z}_{\text{opt}}$. Note that typically $|\mathcal{Z}_{\text{opt}}| \lesssim H|\mathcal{S}|$ or even $|\mathcal{Z}_{\text{opt}}| \lesssim |\mathcal{S}|$ (see Remark B.2 in the appendix). We now state our first result, which gives a gap-dependent regret bound that scales as $\log(1/\delta)$ with probability at least $1 - \delta$. The result is a consequence of a more general result stated as Theorem 2.4, itself a simplified version of more precise bounds stated in Appendix B.1.

Corollary 2.1. Fix $\delta \in (0, 1/2)$, and let $A = |\mathcal{A}|$, $S = |\mathcal{S}|$, $M = (SAH)^2$. Then with probability at least $1 - \delta$, StrongEuler run with confidence parameter δ enjoys the following regret bound for all $K \geq 1$:

$$\begin{aligned} \text{Regret}_K \lesssim & \left(\sum_{(x,a) \in \mathcal{Z}_{\text{sub}}} \frac{H^3}{\text{gap}(x,a)} \log \frac{MT}{\delta} \right) + \frac{H^3 |\mathcal{Z}_{\text{opt}}|}{\text{gap}_{\min}} \log \frac{MT}{\delta} \\ & + H^4 SA(S \vee H) \log \frac{MH}{\text{gap}_{\min}} \log \frac{MT}{\delta}. \end{aligned} \quad (2)$$

Moreover, if \mathcal{M} is either a contextual bandits instance, or has \mathcal{G} -bounded rewards for $\mathcal{G} \lesssim 1$, then the factors of H^3 on the first line can be sharpened to H . In addition, if \mathcal{M} is a contextual bandits instance, the factor of H^3 in the first term (summing over $(x, a) \in \mathcal{Z}_{\text{sub}}$) can be sharpened to 1.

Setting $\delta = 1/T$ and noting that $\sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq KH = T$ with probability 1 (recall $R(x, a) \in [0, 1]$), we see that the expected regret $\mathbb{E}[\sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}]$ can be bounded by replacing $1/\delta$ with T in right hand side of the inequality (2); this yields an expected regret that scales as $\log T$.

Three regret terms: The first term in Corollary 2.1 reflects the sum over sub-optimal state-action pairs, which a lower bound (Proposition 2.2) shows is unimprovable in general. In the infinite horizon setting, [11] gives an algorithm whose regret is asymptotically bounded by an analogue of this term. The third term characterizes the burn-in time suffered by nearly all model-based finite-time analyses and is the number of rounds necessary before standard concentration of measure arguments kick in. The second term is less familiar and is addressed in Section 2.2 below.

H dependence: Comparing to known results from the infinite-horizon setting, one expects the optimal dependence of the first term on the horizon to be H^2 . However, we cannot rule out that the optimal dependence is H^3 for the following three reasons: (i) the infinite-horizon analogues D, \bar{D}, \bar{H} (Section 1.1) are not directly comparable to the horizon H ; (ii) in the episodic setting, we have a potentially different value function \mathbf{V}_h^* for each $h \in [H]$, whereas the value functions of the infinite horizon setting are constant across time; (iii) the H^3 may be unavoidable for non-asymptotic (in T) bounds, even if H^2 is the optimal asymptotic dependence after sufficient burn-in (possibly depending on diameter-like quantities). Resolving the optimal H dependence is left as future work. We also note that for contextual bandits, we incur *no* H dependence on the first term; and thus the first term coincides with the known asymptotically optimal (in T), instance-specific regret [7].

Guarantees for other optimistic algorithms: To make the exposition concrete, we only provide regret bounds for the StrongEuler algorithm. However, the “gap-clipping” trick (Proposition 3.1) and subsequent analysis template described in Section 3.1 can be applied to obtain similar bounds for other recent optimistic algorithms, as in [1, 5, 16, 6].³

2.1 Sub-optimality Gap Lower Bound

Our first lower bound shows that when the total number of rounds $T = KH$ is large, the first term of Corollary 2.1 is unavoidable in terms of regret. Specifically, for every possible choice of gaps, there exists an instance whose regret scales on the order of the first term in (2).

Following standard convention in the literature, the lower bound is stated for algorithms which have sublinear worst case regret. Namely, we say that an algorithm Alg is α -uniformly good if, for any MDP instance \mathcal{M} , there exists a constant $C_{\mathcal{M}} > 0$ such that $\mathbb{E}^{\mathcal{M}}[\text{Regret}_K] \leq C_{\mathcal{M}} K^\alpha$ for all K .⁴

³To achieve logarithmic regret, some of these algorithms require a minor modification to their confidence intervals; otherwise, the gap-dependent regret scales as $\log^2 T$. See Appendix E for details.

⁴We may assume as well that Alg is allowed to take the number of episodes K as a parameter.

Proposition 2.2 (Regret Lower Bound). *Let $S \geq 2$, and $A \geq 2$, and let $\{\Delta_{x,a}\}_{x,a \in [S] \times [A]} \subset (0, H/8)$ denote a set of gaps. Then, for any $H \geq 1$, there exists an MDP \mathcal{M} with states $\mathcal{S} = [S+2]$, actions $\mathcal{A} = [A]$, and H stages, such that,*

$$\begin{aligned} \text{gap}_1(x, a) &= \Delta_{x,a}, & \forall x \in [S], a \in \mathcal{A} \\ \text{gap}_h(x, a) &\geq 1/2, & \forall x \in \{S+1, S+2\}, a \in \mathcal{A} - \{1\}, \end{aligned}$$

and any α -uniformly good algorithm satisfies

$$\lim_{K \rightarrow \infty} \frac{\mathbb{E}^{\mathcal{M}}[\text{Regret}_K]}{\log T} \gtrsim (1 - \alpha) \sum_{x,a: \text{gap}_1(x,a) > 0} \frac{H^2}{\text{gap}_1(x, a)}$$

The above proposition is proven in Appendix H, using a construction based on [4]. For simplicity, we stated an asymptotic lower bound. We remark that if the constant $C_{\mathcal{M}}$ is $\text{poly}(|\mathcal{S}|, |\mathcal{A}|, H)$, then one can show that the above asymptotic bound holds as soon as $K \geq (|\mathcal{S}||\mathcal{A}|H/\text{gap}_*)^{\mathcal{O}(1/(1-\alpha))}$, where $\text{gap}_* := \{\min \text{gap}_1(x, a) : \text{gap}_1(x, a) > 0\}$. More refined non-asymptotic regret bounds can be obtained by following [7].

2.2 Why the dependence on gap_{\min} ?

Without the second term, Corollary 2.1 would only suffer one factor of $1/\text{gap}_{\min}$ due to the sum over state-actions pairs $(x, a) \in \mathcal{Z}_{\text{sub}}$ (when the minimum is achieved by a single pair). However, as remarked above, $|\mathcal{Z}_{\text{opt}}|$ typically scales like $|\mathcal{S}|$ and therefore the second term scales like $|\mathcal{S}|/\text{gap}_{\min}$, with a dependence on $1/\text{gap}_{\min}$ that is at least a factor of $|\mathcal{S}|$ more than we would expect. Here, we show that $|\mathcal{S}|/\text{gap}_{\min}$ is unavoidable for the sorts of optimistic algorithms that we typically see in the literature; a rigorous proof is deferred to Appendix G.

Theorem 2.3 (Informal Lower Bound). *Fix $\delta \in (0, 1/8)$. For universal constants c_1, c_2, c_3, c_4 , if $\epsilon \in (0, c_1)$, and S satisfies $c_2 \log(\epsilon^{-1}/\delta) \leq S \leq c_3 \epsilon^{-1}/\log(\epsilon^{-1}/\delta)$, there exists an MDP with $|\mathcal{S}| = S$, $|\mathcal{A}| = 2$ and horizon $H = 2$, such that exactly one state has a sub-optimality gap of $\text{gap}_{\min} = \epsilon$ and all other states have a minimum sub-optimality gap $\text{gap}_h(x, a) \geq 1/2$. For this MDP, $\sum_{h,x,a: \text{gap}_h(x,a) > 0} \frac{1}{\text{gap}_h(x,a)} \lesssim S + \frac{1}{\text{gap}_{\min}}$ but all existing optimistic algorithms for finite-horizon MDPs which are δ -correct suffer a regret of at least $\frac{S}{\text{gap}_{\min}} \log(1/\delta) \gtrsim \sum_{h,x,a: \text{gap}_h(x,a) > 0} \frac{\log(1/\delta)}{\text{gap}_h(x,a)} + \frac{S \log(1/\delta)}{\text{gap}_{\min}}$ with probability at least $1 - c_4 \delta$.*

The particular instance described in Appendix G that witnesses this lower bound is instructive because it demonstrates a case where optimism results in *over-exploration*.

2.3 Interpolating with Minimax Regret for Small T

We remark that while the logarithmic regret in Corollary 2.1 is non-asymptotic, the expression can be loose for a number of rounds T that is small relative to the sum of the inverse gaps. Our more general result interpolates between the $\log T$ gap-dependent and \sqrt{T} gap-independent regret regimes.

Theorem 2.4 (Main Regret Bound for StrongEuler). *Fix $\delta \in (0, 1/2)$, and let $A = |\mathcal{A}|$, $S = |\mathcal{S}|$, $M = (SAH)^2$. Further, define for all $\epsilon > 0$ the set $\mathcal{Z}_{\text{sub}}(\epsilon) := \{(x, a) \in \mathcal{Z}_{\text{sub}} : \text{gap}(x, a) < \epsilon\}$. Then with probability at least $1 - \delta$, StrongEuler run with confidence parameter δ enjoys the following regret bound for all $K \geq 2$:*

$$\begin{aligned} \text{Regret}_K &\lesssim \min_{\epsilon > 0} \left\{ \sqrt{|\mathcal{Z}_{\text{sub}}(\epsilon)| H T (\log T) \log \frac{MT}{\delta}} \right\} + \sum_{(x,a) \in \mathcal{Z}_{\text{sub}} \setminus \mathcal{Z}_{\text{sub}}(\epsilon)} \frac{H^3}{\text{gap}(x, a)} \log \left(\frac{MT}{\delta} \right) \\ &\quad + \min \left\{ \sqrt{|\mathcal{Z}_{\text{opt}}| H T (\log T) \log \frac{MT}{\delta}}, |\mathcal{Z}_{\text{opt}}| \frac{H^3}{\text{gap}_{\min}} \log \left(\frac{MT}{\delta} \right) \right\} \\ &\quad + H^4 SA(S \vee H) \min \log \frac{MT}{\delta} \left\{ \log \frac{MT}{\delta}, \log \frac{MH}{\text{gap}_{\min}} \right\} \\ &\lesssim \sqrt{HSAT \log(T) \log \left(\frac{MT}{\delta} \right)} + H^4 SA(S \vee H) \log^2 \frac{TM}{\delta}, \end{aligned}$$

where the second inequality follows from the first with $\max\{\max_{\epsilon} |\mathcal{Z}_{\text{sub}}(\epsilon)|, |\mathcal{Z}_{\text{opt}}|\} \leq SA$. Moreover, if \mathcal{M} is an instance of contextual bandits, then the factors of H under the square roots can be refined to a 1, and if \mathcal{M} has $\lesssim 1$ -bounded rewards, then these same factors of H can be replaced by a $1/H$. In both settings, logarithmic terms can be refined as in Corollary 2.1.

By the same argument as above, Theorem 2.4 with $\delta = 1/T$ implies an expected regret scaling like gap-dependent $\log T$ or worst-case \sqrt{HSAT} . In Appendix B.1, we state a more refined bound given in terms of the reward bound \mathcal{G} , and the maximal variance of any state-action pair (Theorem B.2).

3 Gap-Dependent bounds via ‘clipping’

In this section, we (i) introduce the key properties of optimistic algorithms, (ii) explain existing approaches to the analysis of such algorithms, and (iii) introduce the “clipping trick”, and sketch how this technique yields gap-dependent, non-asymptotic bounds.

Definition 3.1 (Optimistic Surplus). Given an optimistic algorithm Alg, we define the (optimistic) *surplus* $\mathbf{E}_{k,h}(x, a) := \overline{\mathbf{Q}}_{k,h}(x, a) - r(x, a) - p(x, a)^\top \overline{\mathbf{V}}_{k,h+1}$. Alg is *strongly optimistic* if $\mathbf{E}_{k,h}(x, a) \geq 0$ for all $k \geq 1$, and $(x, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$, which implies that Alg is also optimistic.

While the nomenclature “suplus” is unique to our work, surplus-like terms arise in many prior regret analyses [5, 16]. The notion of *strong optimism* is novel to this work, and facilitates a sharper H -dependence in contextual bandit setting of Definition 2.2; intuitively, strong optimism means that the Q-function $\overline{\mathbf{Q}}_{k,h}$ at stage h over-estimates \mathbf{Q}_h^* more than $\overline{\mathbf{Q}}_{k,h+1}$ does $\mathbf{Q}_{k,h+1}^*$.

The Regret Decomposition for Optimistic Algorithms: Under optimism alone, we can see that for any h and any $a^* \in \pi^*(x)$,

$$\overline{\mathbf{V}}_{k,h}(x) = \max_a \overline{\mathbf{Q}}_{k,h}(x, a) \geq \overline{\mathbf{Q}}_{k,h}(x, a^*) \geq \mathbf{Q}_h^*(x, a^*) = \mathbf{V}_h^*(x),$$

and therefore, we can bound the sub-optimality of π_k as $\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq \overline{\mathbf{V}}_{k,0} - \mathbf{V}_0^{\pi_k}$.

We can decompose the regret further by introducing the following notation: we let $\omega_{k,h}(x, a) := \mathbb{P}^{\pi_k}[(x_h, a_h) = (x, a)]$ denote the probability of visiting x and playing a at time h in episode k . A standard regret decomposition (see e.g. Lemma E.15 [5]) then shows that for a trajectory $(x_h, a_h)_{h=1}^H$, $\overline{\mathbf{V}}_{k,0} - \mathbf{V}_0^{\pi_k} = \mathbb{E}^{\pi_k}[\sum_{h=1}^H \mathbf{E}_{k,h}(x_h, a_h)] = \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \mathbf{E}_{k,h}(x, a)$, yielding a regret bound of

$$\sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq \sum_{k=1}^K \overline{\mathbf{V}}_{k,0} - \mathbf{V}_0^{\pi_k} \leq \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \mathbf{E}_{k,h}(x, a).$$

Existing Analysis of MDPs: We begin by sketching the flavor of minimax analyses. Introducing the notation $n_k(x, a) := \{\#\text{times } (x, a) \text{ is visited before episode } k\}$, existing analyses carefully manipulate the surpluses $\mathbf{E}_{k,h}(x, a)$ to show that $\sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \mathbf{E}_{k,h}(x, a) \lesssim \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \frac{\overline{\mathcal{C}}_{\mathcal{M}}}{\sqrt{n_k(x, a)}} + \text{lower order terms}$, where typically $\overline{\mathcal{C}}_{\mathcal{M}} = \text{poly}(H, \log(T/\delta))$.

Finally, they replace $n_k(x, a)$ with an “idealized analogue”, $\bar{n}_k(x, a) := \sum_{j=1}^k \sum_{h=1}^H \omega_{j,h}(x, a) := \sum_{j=1}^k \omega_j(x, a)$, where we introduce $\omega_j(x, a) := \sum_{h=1}^H \omega_{j,h}(x, a)$ denote the expected number of visits of (x, a) at episode j . Letting $\{\mathcal{F}_k\}$ denote the filtration capturing all events up to the end episode k , we see that $\mathbb{E}[\bar{n}_k(x, a) - \bar{n}_{k-1} | \mathcal{F}_{k-1}] = \omega_k(x, a)$, and thus by standard concentration arguments (see Lemma B.7, or Lemma 6 in [6]), $\bar{n}_k(x, a)$ and $n_k(x, a)$ are within a constant factor of each other for all k such that $\bar{n}_k(x, a)$ is sufficiently large. Hence, by replacing $n_k(x, a)$ with $\bar{n}_k(x, a)$, we have (up to lower order terms)

$$\sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \lesssim \sum_{x,a} \sum_{k=1}^K \omega_k(x, a) \frac{\overline{\mathcal{C}}_{\mathcal{M}}}{\sqrt{\bar{n}_k(x, a)}} + \text{lower order terms.} \quad (3)$$

A $\sqrt{SAK \text{ poly}(H)}$ bound is typically concluded using a careful application of Cauchy-Schwartz, and an integration-type lemma (e.g., Lemma C.1). An analysis of this flavor is used in Appendix B.4.

On the other hand, one can *exactly* establish the identity $\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} = \sum_{x,a} \sum_{h=1}^H \omega_{k,h}(x, a) \text{gap}_h(x, a)$. Then one can achieve a gap dependent bound as soon as one can show that the algorithm ceases to select suboptimal actions a at (x, h) after sufficiently large T . Crucially, determining if action a is (sub)optimal at (x, h) requires precise knowledge about the value function at other states in the MDP at future stages $h' > h$. This difficulty is why previous gap-dependent analyses appeal to diameter or ergodicity assumptions, which ensure sufficient uniform exploration of the MDP to reason about the value function at subsequent stages.

3.1 The Clipping Trick

We now introduce the “clipping trick”, a technique which merges both the minimax analysis in terms of the surpluses $\mathbf{E}_{k,h}(x, a)$, and the gap-dependent strategy, which attempts to control how many times a given suboptimal action is selected. Core to our analysis, define the *clipping* operator

$$\text{clip}[x \mid \epsilon] = x\mathbb{I}\{x \geq \epsilon\},$$

for all $x, \epsilon > 0$. We can now state our first main technical result, which states that the sub-optimality $\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}$ can be controlled by a sum over surpluses which have been *clipped* to zero whenever they are sufficiently small.

Proposition 3.1. *Let $\mathbf{g}\ddot{\mathbf{a}}\mathbf{p}_h(x, a) := \frac{\mathbf{g}\mathbf{a}\mathbf{p}_{\min}}{2H} \vee \frac{\mathbf{g}\mathbf{a}\mathbf{p}_h(x, a)}{4H}$. Then, if π_k is induced by an optimistic algorithm with surpluses $\mathbf{E}_{k,h}(x, a)$,*

$$\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq 2e \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \text{clip}[\mathbf{E}_{k,h}(x, a) \mid \mathbf{g}\ddot{\mathbf{a}}\mathbf{p}_h(x, a)].$$

If the algorithm is strongly optimistic, and \mathcal{M} is a contextual bandits instance, we can replace $\mathbf{g}\ddot{\mathbf{a}}\mathbf{p}_h(x, a)$ with $\mathbf{g}\ddot{\mathbf{a}}\mathbf{p}_h(x, a) := \frac{\mathbf{g}\mathbf{a}\mathbf{p}_{\min}}{2H} \vee \frac{\mathbf{g}\mathbf{a}\mathbf{p}_h(x, a)}{4}$.

The above proposition is a consequence of a more general bound, Theorem B.3, given in Appendix B. Unlike gap-dependent bounds that appeal to hitting-time arguments, we *do not* reason about when a suboptimal action a will cease to be taken. Indeed, an algorithm may still choose a suboptimal action a *even if* the surplus $\mathbf{E}_{k,h}(x, a)$ is small, because *future* surpluses may be large. Instead, we argue in two parts:

1. A sub-optimal action $a \notin \pi_h^*(x)$ is taken only if $\bar{\mathbf{Q}}_{k,h}(x, a) \geq \mathbf{Q}_h^*(x, a^*)$ for some $a^* \in \pi_h^*(x)$, or equivalently in terms of the surplus, only if $\mathbf{E}_{k,h}(x, a) + p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}^*) > \mathbf{g}\mathbf{a}\mathbf{p}_h(x, a)$. Thus if Alg selects a suboptimal action, then this is because either the *current* surplus $\mathbf{E}_{k,h}(x, a)$ is larger than $\Omega(\frac{\mathbf{g}\mathbf{a}\mathbf{p}_h(x, a)}{H})$, or the expectation over *future* surpluses, captured by $p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}^*)$ is larger than $(1 - \mathcal{O}(\frac{1}{H}))\mathbf{g}\mathbf{a}\mathbf{p}_h(x, a)$. Intuitively, the first case occurs when (x, a) has not been visited enough times, and the second when the future state/action pairs have not experienced sufficient visitation. In the first case, we can clip the surplus at $\Omega(\frac{\mathbf{g}\mathbf{a}\mathbf{p}_h(x, a)}{H})$; in the second, $\mathbf{E}_{k,h}(x, a) + p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}^*) \leq (1 + \mathcal{O}(\frac{1}{H}))p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}^*)$, and push the contribution of $\mathbf{E}_{k,h}(x, a)$ into the contribution of future surpluses. This incurs a factor of at most $(1 + \mathcal{O}(\frac{1}{H}))^H \lesssim 1$, avoiding an exponential dependence on H .
2. Clipping surpluses for pairs (x, a) for optimal $a \in \pi_h^*(x)$ requires more care. We introduce “half-clipped” surpluses $\ddot{\mathbf{E}}_{k,h}(x, a) := \text{clip}[\mathbf{E}_{k,h}(x, a) \mid \frac{\mathbf{g}\mathbf{a}\mathbf{p}_{\min}}{2H}]$ where *all* actions are clipped at $\mathbf{g}\mathbf{a}\mathbf{p}_{\min}/2H$, and recursively define value functions $\ddot{\mathbf{V}}_h^{\pi_k}(\cdot)$ corresponding to these clipped surpluses (see Definition D.1). We then show that, for $\dot{\mathbf{V}}_0^{\pi_k} := \mathbb{E}_{x \sim p_0}[\dot{\mathbf{V}}_1(x)]$, we have (Lemma D.2)

$$\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq 2(\ddot{\mathbf{V}}_0^{\pi_k} - \mathbf{V}_0^{\pi_k}).$$

This argument is based on carefully analyzing when $\pi_{k,h}$ first recommends a suboptimal action $\pi_{k,h}(x) \notin \pi^*(x)$, and showing that when this occurs, $\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}$ is roughly lower bounded by $\frac{\mathbf{g}\mathbf{a}\mathbf{p}_{\min}}{H}$ times the probability of visiting a state x where $\pi_{k,h}(x)$ plays suboptimally. We can then subtract off $\frac{\mathbf{g}\mathbf{a}\mathbf{p}_{\min}}{2H}$ from all the surplus terms at the expense of at most halving the suboptimality, and using the fact $\mathbf{E}_{k,h} - \frac{\mathbf{g}\mathbf{a}\mathbf{p}_{\min}}{2H} \leq \text{clip}[\mathbf{E}_{k,h} \mid \frac{\mathbf{g}\mathbf{a}\mathbf{p}_{\min}}{2H}]$ concludes the bound. This step is crucial, because it allows us to clip the surpluses even at pairs (x, a) where $a \in \pi_h^*(x)$ is the optimal action. We note that in the formal proof of Proposition 3.1, this half-clipping precedes the clipping of suboptimal actions described above.

Unfortunately, the first step involving the half-clipping is rather coarse, and leads to $S/\mathbf{g}\mathbf{a}\mathbf{p}_{\min}$ term in the final regret bound. As argued in Theorem 2.3, this is unavoidable for existing optimistic algorithms, and suggests that Proposition 3.1 cannot be significantly improved in general.

3.2 Analysis of StrongEuler

Recall that StrongEuler is precisely described by Definition 1.1 up to our particular choice of confidence intervals defined (see Algorithm 1 in Appendix E). We now state a surplus bound (proved in Appendix F) that holds for these particular choice of confidence intervals, and which ensures that the strong optimism criterion of Definition 1.1 is satisfied:

Proposition 3.2 (Surplus Bound for Strong Euler (Informal)). *Let $M = SAH$, and define the variances $\text{Var}_{h,x,a}^* := \text{Var}[R(x,a)] + \text{Var}_{x' \sim p(x,a)}[\mathbf{V}_{h+1}^*(x')]$. Then, with probability at least $1 - \delta/2$, the following holds for all $(x,a) \in \mathcal{S} \times \mathcal{A}$, $h \in [H]$ and $k \geq 1$,*

$$0 \leq \mathbf{E}_{k,h}(x,a) \lesssim \underbrace{\sqrt{\frac{\text{Var}_{h,x,a}^* \log(Mn_k(x,a)/\delta)}{n_k(x,a)}}}_{\mathbf{B}_{k,h}^{\text{lead}}(x,a)} + \text{lower order terms.}$$

We emphasize that Proposition 3.2, and its formal analogue Proposition B.4 in Appendix B.2, are the *only* part of the analysis that relies upon the particular form of the StrongEuler confidence intervals; to analyze other model-based optimistic algorithms, one would simply establish an analogue of this proposition, and continue the analysis in much the same fashion. While Q-learning [9] also satisfies optimism, it induces a more intricate surplus structure, which may require a different analysis.

Recalling the clipping from Proposition 3.1, we begin the gap-dependent bound with $\sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \lesssim \sum_{x,a,k,h} \omega_{k,h}(x,a) \text{clip}[\mathbf{E}_{k,h}(x,a) | \text{g}\ddot{\text{a}}\text{p}_h(x,a)]$. Neglecting lower order terms, Proposition 3.2 ensures that this is approximately less than $\sum_{x,a,k,h} \omega_{k,h}(x,a) \text{clip}[\mathbf{B}_{k,h}^{\text{lead}}(x,a) | \text{g}\ddot{\text{a}}\text{p}_h(x,a)]$. Introduce the minimal (over h) clipping-gaps $\text{g}\ddot{\text{a}}\text{p}(x,a) := \min_h \text{g}\ddot{\text{a}}\text{p}_h(x,a) \geq \frac{\text{gap}(x,a) \vee \text{gap}_{\min}}{4H}$ and maximal variances $\text{Var}_{x,a}^* := \max_h \text{Var}_{h,x,a}^*$. We can then render $\mathbf{B}_{k,h}^{\text{lead}}(x,a) \leq f(n_k(x,a))$, where $f(u) \lesssim \text{clip}\left[\sqrt{\frac{1}{u} \text{Var}_{x,a}^* \log(Mu/\delta)} | \text{g}\ddot{\text{a}}\text{p}(x,a)\right]$. Recalling the approximation $n_k(x,a) \approx \bar{n}_k(x,a)$ described above, we have, to first order,

$$\begin{aligned} \sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} &\lesssim \sum_{x,a,k,h} \omega_{k,h}(x,a) \text{clip}[\mathbf{B}_{k,h}^{\text{lead}}(x,a) | \text{g}\ddot{\text{a}}\text{p}_h(x,a)] \\ &\lesssim \sum_{x,a,k} \omega_k(x,a) f(n_k(x,a)) \lesssim \sum_{x,a,k} \omega_k(x,a) f(\bar{n}_k(x,a)), \end{aligned}$$

where we recall the expected visitations $\omega_k(x,a) := \sum_{h=1}^H \omega_{k,h}(x,a)$. Since $\bar{n}_k(x,a) := \sum_{j=1}^k \omega_j(x,a)$, we can regard the above as an *integral* of the function $f(u)$ (see Lemma C.1), with respect to the *density* $\omega_k(x,a)$. Evaluating this integral (Lemma B.9) yields (up to lower order terms)

$$\sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \lesssim \sum_{x,a} \frac{H \text{Var}_{x,a}^* \log \frac{MT}{\delta}}{\min_h \text{g}\ddot{\text{a}}\text{p}_h(x,a)} \lesssim \sum_{x,a} \frac{H \text{Var}_{x,a}^* \log \frac{MT}{\delta}}{\text{gap}(x,a) \vee \text{gap}_{\min}}.$$

Finally, bounding $\text{Var}_{x,a}^* \leq H^2$ and splitting the bound into the states $\mathcal{Z}_{\text{sub}} := \{(x,a) : \text{gap}(x,a) > 0\}$ and $\mathcal{Z}_{\text{opt}} := \{(x,a) : \text{gap}(x,a) = 0\}$ recovers the first two terms in Corollary 2.1. In benign instances (Definition 2.2), we can bound $\text{Var}_{h,x,a}^* \lesssim 1$, improving the H -dependence. In contextual bandits, we save an additional H factor via $\text{g}\ddot{\text{a}}\text{p}_h(x,a) \gtrsim (\text{gap}_{\min}/H) \vee \text{gap}(x,a)$. The interpolation with the minimax rate in Theorem 2.4 is described in greater detail in Appendix B.4.

References

- [1] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 263–272. JMLR. org, 2017.
- [2] Peter L Bartlett and Ambuj Tewari. Regal: A regularization based algorithm for reinforcement learning in weakly communicating mdps. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 35–42. AUAI Press, 2009.
- [3] Apostolos N Burnetas and Michael N Katehakis. Optimal adaptive policies for markov decision processes. *Mathematics of Operations Research*, 22(1):222–255, 1997.
- [4] Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2818–2826, 2015.
- [5] Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying pac and regret: Uniform pac bounds for episodic reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 5713–5723, 2017.
- [6] Christoph Dann, Lihong Li, Wei Wei, and Emma Brunskill. Policy certificates: Towards accountable reinforcement learning. *arXiv preprint arXiv:1811.03056*, 2018.
- [7] Aurélien Garivier, Pierre Ménard, and Gilles Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 2018.
- [8] Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.
- [9] Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4868–4878, 2018.
- [10] Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- [11] Jungseul Ok, Alexandre Proutiere, and Damianos Tranos. Exploration in structured reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 8888–8896, 2018.
- [12] Ian Osband and Benjamin Van Roy. On lower bounds for regret in reinforcement learning. *stat*, 1050:9, 2016.
- [13] Max Simchowitz, Kevin Jamieson, and Benjamin Recht. Best-of-k-bandits. In *Conference on Learning Theory*, pages 1440–1489, 2016.
- [14] Ambuj Tewari. *Reinforcement learning in large or unknown MDPs*. University of California, Berkeley, 2007.
- [15] Ambuj Tewari and Peter L Bartlett. Optimistic linear programming gives logarithmic regret for irreducible mdps. In *Advances in Neural Information Processing Systems*, pages 1505–1512, 2008.
- [16] Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. *arXiv preprint arXiv:1901.00210*, 2019.

Contents

1	Introduction	1
1.1	Related Work	2
1.2	Problem Setting, Notation, and Organization	3
1.3	Optimistic Algorithms	4
2	Main Results	4
2.1	Sub-optimality Gap Lower Bound	5
2.2	Why the dependence on gap_{\min} ?	6
2.3	Interpolating with Minimax Regret for Small T	6
3	Gap-Dependent bounds via ‘clipping’	7
3.1	The Clipping Trick	8
3.2	Analysis of StrongEuler	9
A	Notation and Organization	13
	Notation Table	13
I	Precise Results and Analysis	15
B	Precise Statement and Rigorous Proof Sketch of Main Regret Bounds	15
B.1	More Precise Statement of Regret Bound Theorem 2.4	15
B.2	Rigorous proof of upper bounds: Preliminaries	17
B.3	Proof of Corollary B.1: A proof via integration	18
B.4	Proof of Theorem B.2	21
C	Proof of Technical Lemmas	22
C.1	Proof of clipping with future bounds, Lemma B.6	22
C.2	Proof of sampling lemma (Lemma B.7)	24
C.3	Proof of integral conversion, Lemma B.8	24
C.4	Proof of interal conversion for \mathcal{G} -bounds, Lemma B.10	25
C.5	General Integral computations (Lemma B.9)	25
D	Proof of ‘clipping’ bound: Proposition 3.1 / Theorem B.3	29
D.1	Proof of Lemma D.2	30
D.2	Proof of Lemma D.1	32
D.3	Proof of Lemma D.3	33
II	StrongEuler and its surpluses	34

E	The StrongEuler Algorithm	34
F	Analysis of StrongEuler: Proof of Proposition B.4	35
F.1	Proof of Optimism	35
F.2	Proof of Surplus Bound Upper Bound	37
F.3	Definition of $\mathcal{A}^{\text{conc}}$, and proofs of supporting lemmas	40
III	Lower Bounds	44
G	Min-Gap Lower Bound for Optimistic Algorithms (Theorem 2.3)	44
G.1	Formal Statement	44
G.2	Algorithm Class	44
G.3	Formal Lower Bound Instance	45
G.4	The Lower Bound:	45
G.4.1	Proof of Proposition G.3	46
G.5	Proof of Claim G.5	48
H	Information Theoretic Lower Bound (Proposition 2.2)	51
H.1	Construction of the hard instance	51
H.2	Regret Lower Bound Decomposition	52
H.3	Proof of Equation (29)	52

General Notation
\lesssim denotes inequality up to a universal constant.
$f \approx g$ denotes $f \lesssim g \lesssim f$. $\log_+(x) := \log \max\{x, 1\}$.
\mathbb{I} denotes an indicator function
$\mathcal{M} = (\mathcal{S}, \mathcal{A}, [H], p_0, p, R)$ denotes an MDP
H denote the horizon,
\mathcal{A} and \mathcal{S} denotes the space of actions and states
$A := \mathcal{A} $ and $S := \mathcal{S} $
$h \in [H], a \in \mathcal{A}, x \in \mathcal{S}$ are used for stages, actions, and states
$R(x, a) \in [0, 1]$ denotes the R.V. with reward distribution at (x, a) .
$r(x, a) := \mathbb{E}[R(x, a)]$ denotes expected reward
$p_0(x)$ denotes initial distribution of x_1
$p(x' x, a)$ denotes transition probability.
$M = SAH$
K denotes number of episodes, indexed with $k \in [K]$
$T = KH$ denotes total length of game.

A Notation and Organization

Organization: This section describes the organization of the appendix, and clarifying our notation. The remainder of the appendix is divided into three parts:

Part I presents more detailed statements of the regret upper bounds obtained by StrongEuler, and their complete proofs. Section B.2 introduces Corollary B.1 and Theorem B.2, refining Corollary 2.1 and Theorem 2.4, from the main text. The section continues to prove both results. In addition, we introduce Theorem B.3, which refines the clipped regret decomposition Proposition 3.1. The proofs in this section rely on numerous technical lemmas, whose proofs are deferred to Section C. Finally, this section states Proposition B.4, which ensures that StrongEuler is optimistic and provides a precise bound on the surpluses $\mathbf{E}_{k,h}(x, a)$, described informally in Proposition 3.2.

In Part II, we present the StrongEuler algorithm and its guarantees. Section E describes how StrongEuler instantiates the *model-based* examples of optimistic algorithms described in Section 1.3; the algorithm and choice of confidence bonuses are specified in pseudocode. In Section F, we prove the surplus bound Proposition B.4, and verify that StrongEuler is strongly optimistic

Lastly, Part III contains the proofs of our lower bounds. Section G proves the $\Omega(S/\text{gap}_{\min})$ lower bound described in Theorem 2.3, and rigorously describes the class of algorithms to which it applies. Finally, Section H proves the information theoretic lower bound, Proposition 2.2.

Notational Rationale: Unfortunately, the regret analysis of tabular MDPs requires significant notational overhead. Here we take a moment to highlight some notational conventions that we shall use throughout. The superscript $(\cdot)^*$ denotes “optimal” quantities, i.e. the optimal policy π^* , the optimal value \mathbf{V}^* , and variances of the optimal policy $\text{Var}_{h,x,a}^*$. The accents $\overline{(\cdot)}$ will be used to denote upper bounds on quantities, e.g. an optimistic Q-function $\overline{\mathbf{Q}}_{k,h}$ is an upper bound on $\mathbf{Q}_{k,h}^*$, and $\overline{\text{var}}$ is an upper bound on the variance, and so on. $\underline{(\cdot)}$ will denote lower bounds on quantities.

For example, StrongEuler will maintain lower bounds on the values $\underline{\mathbf{V}}_{k,h} \leq \mathbf{V}^*$. The accent $\check{(\cdot)}$ will pertain to clipped quantities; e.g. $\check{\text{gap}}$ is the gap-value at which surpluses are clipped. Many quantities, like $\text{gap}_h(x, a)$ (gaps) and $\text{Var}_{h,x,a}^*$ (variances) depend on the triples (x, a, h) . The quantities $\text{gap}(x, a)$ and $\text{Var}_{x,a}^*$ with h suppressed to denote worse-case bounds on these term over $h \in [H]$; e.g. $\text{gap}(x, a) := \min_{h \in [H]} \text{gap}_h(x, a)$ and $\text{Var}_{x,a}^* := \max_{h \in [H]} \text{Var}_{h,x,a}^*$.

Policies, Value Functions, Q-functions
$\pi = (\pi_h)_{h=1}^H$ denotes a policy with $\pi_h : \mathcal{S} \rightarrow \mathcal{A}$
\mathbf{V}_0^π denotes the value of π
$\mathbf{V}_h^\pi(x)$ denotes the value of π at $h \in [H]$ and $x \in \mathcal{S}$
$\mathbf{Q}_h^\pi(x, a)$ denotes Q-function of π
$\mathbf{V}_0^*, \mathbf{V}_h^*(x), \mathbf{Q}^*(x, a)$ denote optimal value, value function, Q-function
$\pi_h^*(x)$ denotes the set of optimal actions at $h \in [H], x \in \mathcal{S}$.
$\overline{\mathbf{V}}_{k,h}(x)/\overline{\mathbf{Q}}_{k,h}(x, a)$ denotes optimistic value/Q function
$\pi_{k,h}(x) = \arg \max_a \overline{\mathbf{Q}}_{k,h}(x, a)$ denotes optimistic policy

Problem Dependent Quantities
$\text{gap}_h(x, a) := \mathbf{V}_h^*(x) - \mathbf{Q}_h^*(x, a)$.
$\text{gap}(x, a) := \min_h \text{gap}_h(x, a)$
$\text{gap}_{\min} := \min_{x,a} \{ \text{gap}_h(x, a) : \text{gap}_h(x, a) > 0 \}$
$\alpha_{x,a,h} \in [0, 1]$ denotes transition suboptimality (Definition B.3)
$\text{Var}_{h,x,a}^* := \text{Var}[R(x, a)] + \text{Var}_{x' \sim p(x,a)}[\mathbf{V}_{h+1}^*(x')]$
$\text{Var}_{x,a}^* := \max_h \text{Var}_{h,x,a}^*$
$\overline{\text{Var}} := \max_{x,a,h} \text{Var}_{h,x,a}^*$
$\mathcal{G} \leq H$: upper bound on $\sum_{h=1}^H R(x, \pi_h(x))$ (Definition 2.2)
$\overline{H}_T := \min\{\overline{\text{Var}}, \mathcal{G}^2/H\}$

Quantities for Analysis
$\mathbf{L}(u) := \sqrt{2 \log(10M^2 \max\{u, 1\}^2/\delta)}$
$\omega_{k,h}(x, a) := \mathbb{P}^{\pi_k}[(x_h, a_h) = (x, a)]$ denotes the surplus
$\omega_k(x, a) := \sum_{h=1}^H \omega_{k,h}(x, a)$ denotes the surplus
$n_k(x, a)$ denotes the number of times (x, a) is observed up to time $k - 1$
$\overline{n}_k(x, a) := \sum_{t=1}^k \omega_t(x, a)$.
$\tau(x, a)$ denotes time after which $\overline{n}_k(x, a)$ is sufficiently large
$\mathcal{A}^{\text{conc}}$ (good concentration event)
$\mathcal{E}^{\text{samp}}$ (good sampling event, Lemma B.7)
$H_{\text{sample}} \lesssim H \log \frac{M}{\delta}$ number of samplers for $\mathcal{E}^{\text{samp}}$ to apply
$\text{g}\check{\text{a}}\text{p}_h(x, a) := \frac{\text{gap}_{\min}}{2H} \vee \frac{\text{gap}_h(x, a)}{4(H\alpha_{x,a,h} \vee 1)}$ (clipped gap)
$\text{g}\check{\text{a}}\text{p}_{\min} := \min_{x,a,h} \text{g}\check{\text{a}}\text{p}_h(x, a)$ (clipped gap)
$\text{Var}_{h,x,a}^\pi := \text{Var}[R(x, a)] + \text{Var}_{x' \sim p(x,a)}[\mathbf{V}_{h+1}^\pi(x')]$
$\text{Var}_{h,x,a}^{(k)} = \min \left\{ \text{Var}_{h,x,a}^*, \text{Var}_{h,x,a}^{\pi_k} \right\}$
$\mathbf{E}_{k,h}(x, a) := \mathbf{Q}_{k,h}(x, a) - r(x, a) - p(x, a)^\top \overline{\mathbf{V}}_{k,h+1}$ denotes the surplus
$\mathbf{E}_{k,h}(x, a) \lesssim \mathbf{B}_{k,h}^{\text{lead}}(x, a) + \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{B}_k^{\text{fut}}(x_t, a_t) \mid (x_h, a_h) = (x, a) \right]$ (surplus bound, Proposition B.4)
$\mathbf{B}_{k,h}^{\text{lead}}(x, a) := H \wedge \sqrt{\frac{\text{Var}_{h,x,a}^{(k)} \log\left(\frac{Mn_k(x,a)}{\delta}\right)}{n_k(x,a)}}$ (lead bound on surplus)
$\mathbf{B}_k^{\text{fut}}(x, a) = H^3 \wedge H^3 \left(\sqrt{\frac{S \log\left(\frac{Mn_k(x,a)}{\delta}\right)}{n_k(x,a)}} + \frac{S \log\left(\frac{Mn_k(x,a)}{\delta}\right)}{n_k(x,a)} \right)^2$ (bound on future surpluses)

Part I

Precise Results and Analysis

B Precise Statement and Rigorous Proof Sketch of Main Regret Bounds

In this section, we present a precise statements and formal proofs of the upper bounds, Corollary 2.1 and Theorem 2.4, from the main text. These bounds both makes the improvements in the benign instances of Definition 2.2, and takes advantage of other possibly-favorable instance-specific quantities. The remainder of the section is organized as follows. In Section B.1, we introduce the relevant problem-dependent quantities in terms of which we state our more refined bounds. We then state Corollary B.1, a more precise analogue of the log-regret bounds in Corollary 2.1, followed by Theorem B.2, which refines the regret bound Theorem 2.4 in interpolating between the $\log T$ and \sqrt{T} regimes.

Next, in Section B.2, we set up the preliminaries for the proof of our upper bound, including (a) Theorem B.3, the granular clipping bound strengthening Proposition 3.1, (b) Proposition B.4, which upper bounds the surpluses $\mathbf{E}_{k,h}(x, a)$ for StrongEuler, and (c) Lemma B.6 which combines the two into a useful form.

Then, in Section B.2, we present a rigorous proof of Corollary B.1 based on integration tools developed in Section C. Finally, we modify the arguments slightly to obtain the interpolation in Theorem B.2. The proof of Proposition 3.1 is given in Section F, Theorem B.3 is given in Section D, and the remainder of technical results in the present section are established in Section C.

We emphasize that the tools in this section provide a general recipe for establishing similar regret bounds for the existing model-based optimistic algorithms in the literature. We have attempted to present our tools in a modular fashion in hope that they can be borrowed to automate the proofs of similar guarantees in related settings.

B.1 More Precise Statement of Regret Bound Theorem 2.4

We shall begin by stating a more precise version of Theorem 2.4. Following [16], we begin by defining the variances of the value optimal functions::

Definition B.1 (Variance Terms). We define the variance of a triple (x, a, h) as

$$\text{Var}_{h,x,a}^* := \text{Var}[R(x, a)] + \text{Var}_{x' \sim p(x,a)}[\mathbf{V}_{h+1}^*(x')],$$

and the statewise maximal variances as $\text{Var}_{x,a}^* := \max_h \text{Var}_{h,x,a}^*$, and the maximal variance as $\overline{\text{Var}} := \max_{x,a,h} \text{Var}_{h,x,a}^*$.

Remark B.1 (Typical Bounds on the variance). While $\overline{\text{Var}} \leq H^2$ for general MDPs (see e.g. [1]), we have $\overline{\text{Var}}$ is smaller for the benign instances in Definition 2.2. We briefly summarize this discussion from [16]: If \mathcal{M} has \mathcal{G} bounded rewards, then $\mathbf{V}_{h+1}^*(x) \leq \mathcal{G}$ for any x , and thus $\text{Var}_{h,x,a}^* \leq 1 + \mathcal{G}^2$, which is $\lesssim 1$ if $\mathcal{G} \lesssim 1$. For contextual bandits, $p = p(x, a)$ does not depend on x, a , and $\mathbf{V}_{h+1}^*(x) = (\max_a R(x, a)) + (\mathbb{E}_{x' \sim p} \mathbf{V}_{h+2}^*(x'))$, where the second term does not dependent Hence, $\text{Var}_{x' \sim p}[\mathbf{V}_{h+1}^*(x')] \leq \text{Var}[(\max_a R(x, a))] \leq 1$, and thus $\text{Var}_{h,x,a}^* \leq 2$.

We can then define an associated “effective horizon”, which replaces H with a possibly smaller problem dependent quantity:

Definition B.2 (Effective Horizon). Suppose that \mathcal{M} has \mathcal{G} -bounded rewards, as in Definition 2.2) We define the *effective horizon* as

$$\overline{H}_T := \min \left\{ \overline{\text{Var}}, \frac{\mathcal{G}^2}{\overline{\text{Var}}} \log T \right\}.$$

Since any horizon- H MDP has H -bounded rewards, \overline{H}_T always satisfies $\overline{H}_T \leq \min \{H^2, H \log T\}$.

We note that the bound $\overline{\text{Var}} \lesssim 1$ for contextual bandits implies $\overline{H}_T \lesssim 1$, whereas if \mathcal{M} has \mathcal{G} -bounded rewards with $\mathcal{G} \lesssim 1$, $\overline{H}_T \lesssim 1 \wedge \frac{1}{H} \log T$.

Lastly, we shall introduce one more condition we call *transition suboptimality*, which is a notion of distributional closeness that enables the improved clipping and sharper regret bounds for the special case of contextual bandits (Definition 2.2):

Definition B.3 (Transition Sub-optimality). Given $\alpha \in [0, 1]$, we say that a tuple (x, a, h) is α -transition suboptimal if there exists an $a^* \in \pi_h^*(x)$ such that

$$p(x'|x, a) - p(x'|x, a^*) \leq \alpha p(x'|x, a) \quad \forall x' \in \mathcal{S}.$$

Intuitively, the condition states that the transition distributions $p(x, a)$ and $p(x, a^*)$ are close in a pointwise, multiplicative sense. This is motivated by the contextual bandit setting of Definition 2.2, where each (x, a, h) is exactly 0-transition suboptimal. For arbitrary MDPs, the bound $p(x'|x, a^*) \geq 0$ implies that every triple (x, a, h) is at most 1-transition suboptimal.⁵

With these definitions in place, we can state the more precise analogue of Corollary 2.1 as follows:

Corollary B.1 (Logarithmic Regret Bound for StrongEuler). Fix $\delta \in (0, 1/2)$, and let $A = |\mathcal{A}|$, $S = |\mathcal{S}|$, $M = (SAH)^2$. Then with probability at least $1 - \delta$, StrongEuler run with confidence parameter δ enjoys the following regret bound for all $K \geq 2$:

$$\begin{aligned} \text{Regret}_K \lesssim & \sum_{(x,a) \in \mathcal{Z}_{\text{sub}}} \frac{\text{Var}_{x,a}^*(1 \vee \alpha H)}{\text{gap}(x,a)} \log\left(\frac{M}{\delta} \left(\frac{H}{\text{gap}(x,a)} \wedge T\right)\right) + |\mathcal{Z}_{\text{opt}}| \frac{\overline{\text{Var}} H}{\text{gap}_{\min}(x,a)} \log\left(\frac{M}{\delta} \left(\frac{H}{\text{gap}(x,a)} \wedge T\right)\right) \\ & + H^4 SA(S \vee H) \log \frac{TM}{\delta} \min \left\{ \log \frac{HM}{\text{gap}_{\min}}, \log \frac{TM}{\delta} \right\} \end{aligned}$$

In particular, if \mathcal{M} is an instance of contextual bandits, then $\overline{\text{Var}}$ can be replaced by 1, \overline{H}_T can be replaced by 1 and $\max\{\alpha H, 1\} = 1$. If \mathcal{M} has $\mathcal{G} \lesssim 1$ bounded rewards, then $\overline{\text{Var}}$ can be replaced by 1 in the above bound,

Moreover, our more precise analogue of Theorem 2.4, which interpolates between the $\log T$ and \sqrt{T} regimes, is as follows:

Theorem B.2 (Main Regret Bound for StrongEuler). Fix $\delta \in (0, 1/2)$, and let $A = |\mathcal{A}|$, $S = |\mathcal{S}|$, $M = (SAH)^2$. Let \overline{H}_T be as in Definition B.2, and suppose that each tuple (x, a, h) is α -transition suboptimal. Further, define $\mathcal{Z}_{\text{sub}}(\epsilon) := \{(x, a) \in \mathcal{Z}_{\text{sub}} : \text{gap}(x, a) < \epsilon\}$. Then with probability at least $1 - \delta$, StrongEuler run with confidence parameter δ enjoys the following regret bound for all $K \geq 2$:

$$\begin{aligned} \text{Regret}_K \lesssim & \min_{\epsilon > 0} \left\{ \sqrt{\overline{H}_T |\mathcal{Z}_{\text{sub}}(\epsilon)| T \log \frac{MT}{\delta}} + \sum_{(x,a) \in \mathcal{Z}_{\text{sub}} \setminus \mathcal{Z}_{\text{sub}}(\epsilon)} \frac{\max\{\alpha H, 1\} \text{Var}_{x,a}^* \left(\frac{M}{\delta} \left(\frac{H}{\text{gap}(x,a)} \wedge T\right)\right)}{\text{gap}(x,a)} \right\} \\ & + \min \left\{ \sqrt{\overline{H}_T |\mathcal{Z}_{\text{opt}}| T \log \frac{MT}{\delta}}, |\mathcal{Z}_{\text{opt}}| \frac{H \overline{\text{Var}}}{\text{gap}_{\min}} \log\left(\frac{M}{\delta} \left(\frac{H}{\text{gap}(x,a)} \wedge T\right)\right) \right\} \\ & + H^4 SA(S \vee H) \log \frac{MT}{\delta} \min \left\{ \log \frac{MH}{\text{gap}_{\min}}, \log \frac{MT}{\delta} \right\} \\ \lesssim & \sqrt{HSAT \log(T) \log\left(\frac{MT}{\delta}\right)} + H^4 SA(S \vee H) \log^2 \frac{TM}{\delta}, \end{aligned}$$

where the second inequality follows from the first with $\max\{|\mathcal{Z}_{\text{sub}}(\epsilon)|, |\mathcal{Z}_{\text{opt}}|\} \leq SA$. In particular, if \mathcal{M} is an instance of contextual bandits, then $\overline{\text{Var}}$ can be replaced by 1, \overline{H}_T can be replaced by 1 and $\max\{\alpha H, 1\} = 1$. If \mathcal{M} has $\mathcal{G} \lesssim 1$ bounded rewards, then $\overline{\text{Var}}$ can be replaced by 1 in the above bound, and \overline{H}_T replaced by $\min\{1, \frac{\log T}{H}\}$.

We observe that Theorem 2.4, and Corollary 2.1 are direct consequences of the above theorem.

⁵The condition can be relaxed somewhat to only needing to hold for a set \mathcal{S} for which $p(x' \in \mathcal{S} \mid x, a)$ is close to 1; for simplicity, consider the unrelaxed notion as defined as above.

Remark B.2 (Bounds on $|\mathcal{Z}_{\text{opt}}|$). Note that $|\mathcal{Z}_{\text{opt}}| \leq \sum_{x,h} |\pi_h^*(x)|$; in particular if for each (x, h) there is exactly one optimal action, then $|\mathcal{Z}_{\text{opt}}| \leq H|\mathcal{S}|$. If in addition the same action is optimal at x for each $h \in [H]$, then $|\mathcal{Z}_{\text{opt}}| = |\mathcal{S}|$. For many environments $|\mathcal{Z}_{\text{opt}}| \lesssim |\mathcal{S}|$; for instance, a race car doing many laps around a track may have h -dependent optimal actions in the first and last laps, but for the steady-state laps the optimal action will depend just on the current state.

Remark B.3 (Coupling Variances and Gaps). For state action pairs $(x, a) \in \mathcal{Z}_{\text{sub}}$, Corollary B.1 Theorem B.2 suffer for the term $\frac{(1 \vee \alpha H) \text{Var}_{x,a}^*}{\text{gap}(x,a)}$, where $\text{Var}_{x,a}^* := \max_h \text{Var}_{h,x,a}^*$ is the maximal variance over stages, and $\text{gap}(x, a) = \min_h \text{gap}_h(x, a)$ is the minimal gap. This quantity can be refined to depend on roughly $\max_h \frac{\text{Var}_{h,x,a}^*}{\text{gap}_h(x,a)}$, coupling the variance and gap terms. To do so, one needs to bin the gaps into intervals of $[2^{j-1}H, 2^j]$ or integers $j \in \mathbb{N}$, and apply numerous careful manipulations. In the interest of brevity, we defer the details to a later work.

B.2 Rigorous proof of upper bounds: Preliminaries

We now turn to a rigorous proof of the regret bounds for StrongEuler: Corollary B.1 and Theorem B.2 (and consequently Theorem 2.4 and Corollary 2.1).

We first state our generalized surplus clipping bound in terms of the transition-suboptimality condition, which generalizes Proposition 3.1:

Theorem B.3. *Suppose that each tuple (x, a, h) is $\alpha_{x,a,h}$ transition-suboptimal, and set $\text{g}\ddot{\text{a}}\text{p}_h(x, a) := \frac{\text{g}\ddot{\text{a}}\text{p}_{\text{min}}}{2H} \vee \frac{\text{g}\ddot{\text{a}}\text{p}_h(x,a)}{4(H\alpha_{x,a,h} \vee 1)}$. Then, if π_k is induced by a strongly optimistic algorithm with surpluses $\mathbf{E}_{k,h}(x, a)$,*

$$\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq 2e \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \text{clip} [\mathbf{E}_{k,h}(x, a) \mid \text{g}\ddot{\text{a}}\text{p}_h(x, a)].$$

If the algorithm is optimistic but not strongly optimistic, then the above holds by replacing $\alpha_{x,a,h}$ with 1 in the definition of $\text{g}\ddot{\text{a}}\text{p}_h(x, a)$.

The proof of the above theorem is given in Section D. We remark that the above theorem specializes to Proposition 3.1 by noting that each tuple (x, a, h) is 0-transition suboptimal for contextual bandits. For simplicity, we shall assume in the proof of Theorem B.2 that each state is α -suboptimal for a common α ; the bound can be straightforwardly refined to allow α to vary across (x, a, h) .

Next, in order to ensure optimal H -dependence when interpolating with the $\mathcal{O}(\sqrt{T})$ regret bounds, we introduce policy-dependent variance quantities:

Definition B.4. Define the variances

$$\text{Var}_{h,x,a}^{\pi} := \text{Var}[R(x, a)] + \text{Var}_{x' \sim p(x,a)}[\mathbf{V}_{h+1}^{\pi}(x')].$$

where we recall that $\text{Var}_{h,x,a}^* := \text{Var}_{h,x,a}^{\pi^*}$. Further, define $\text{Var}_{h,x,a}^{(k)} = \min \left\{ \text{Var}_{h,x,a}^*, \text{Var}_{h,x,a}^{\pi_k} \right\}$.

We are now ready to state the formal version of Proposition 3.2, which upper bounds the surpluses of StrongEuler, and verifies that the algorithm satisfies strong optimism:

Proposition B.4 (Surplus Bound for StrongEuler). *There exists a universal constant $\mathbf{c} \geq 1$ and event $\mathcal{A}^{\text{conc}}$, with $\mathbb{P}[\mathcal{A}^{\text{conc}}] \geq 1 - \delta/2$, such that on $\mathcal{A}^{\text{conc}}$, for all $x \in \mathcal{S}$, $a \in \mathcal{A}$, $h \in [H]$ and $k \geq 1$,*

$$0 \leq \frac{1}{\mathbf{c}} \mathbf{E}_{k,h}(x, a) \leq \mathbf{B}_{k,h}^{\text{lead}}(x, a) + \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{B}_k^{\text{fut}}(x_t, a_t) \mid (x_h, a_h) = (x, a) \right].$$

where have defined the terms

$$\mathbf{B}_{k,h}^{\text{lead}}(x,a) := H \wedge \sqrt{\frac{\text{Var}_{h,x,a}^{(k)} \log\left(\frac{Mn_k(x,a)}{\delta}\right)}{n_k(x,a)}}, \quad \text{and}$$

$$\mathbf{B}_k^{\text{fut}}(x,a) = H^3 \wedge H^3 \left(\sqrt{\frac{S \log\left(\frac{Mn_k(x,a)}{\delta}\right)}{n_k(x,a)}} + \frac{S \log\left(\frac{Mn_k(x,a)}{\delta}\right)}{n_k(x,a)} \right)^2.$$

The above proposition is proven in Appendix F. Here \mathbf{B}^{lead} denotes a ‘‘lead term’’ in the analysis, which contributes to the dominate factors in our regret bounds. \mathbf{B}^{fut} notates ‘‘future’’ bound terms under a rollout of π_k starting at a given triple (x, a, h) ; these terms are responsible for the lower $\tilde{\mathcal{O}}(SAH^4(S \vee H))$ -term in the regret.

Remark B.4 (Remarks on Proposition B.4). First, the dominant term in the upper bound on $\mathbf{E}_{k,h}$ is $\mathbf{B}_{k,h}^{\text{lead}}(x,a)$, which decays as $\tilde{\mathcal{O}}(n_k(x,a)^{-1/2})$. The terms $\mathbf{B}_k^{\text{fut}}(x,a)$ decay more rapidly $\tilde{\mathcal{O}}(n_k(x,a)^{-1})$, and will thus be responsible for the (nearly gap-free) portion of the regret. Second, in order to analyze similar optimistic algorithms in the same vein (e.g. [1, 4, 5]), one would instead prove the appropriate analogue to Proposition B.4 and follow the remaining steps of the present proof. Little would change, except one would be forced to replace $\text{Var}_{h,x,a}^*$ with a more pessimistic, less problem-dependent quantity. Lastly, note that the lead term $\mathbf{B}_{k,h}^{\text{lead}}(x,a)$ depends on the *minimum* of the variance of the optimal value function, $\text{Var}_{h,x,a}^*$ and of the variance of the value function for π_k , $\text{Var}_{h,x,a}^{\pi_k}$. As in the aforementioned works, this dependence on $\text{Var}_{h,x,a}^{\pi_k}$ is crucial for obtaining the correct minimax $\tilde{\mathcal{O}}(\sqrt{HSAT})$ regret.

Next, let us combine Proposition B.4 with our main clipping theorem, Theorem B.3. Since $\mathbf{E}_{k,h}(x,a) \lesssim \mathbf{B}_{k,h}^{\text{lead}}(x,a) + \mathbb{E}^{\pi_k}[\sum(\dots) | (\dots)]$, combining the two results into a convenient form requires that we reason about how to distribute clipping operations across sums of terms. To this end, we invoke the following technical lemma:

Lemma B.5 (Distributing the clipping operator). *Let $m \geq 2$, $a_1, \dots, a_m \geq 0$, and $\epsilon \geq 0$. $\text{clip}[\sum_{i=1}^m a_i | \epsilon] \leq 2 \sum_{i=1}^m \text{clip}[a_i | \frac{\epsilon}{2m}]$.*

Proof. Let us assume without loss of generality $0 \leq a_1 \leq \dots \leq a_m$, and that $\sum_{i=1}^m a_i \geq \epsilon$. Defining the index $i^* := \min\{i : a_i \geq \frac{\epsilon}{2m}\}$, we observe that $a_{i^*} \geq \frac{\epsilon}{2m}$, and since (a_i) are non-decreasing by assumption, $\sum_{i=i^*}^m a_i = \sum_{i=i^*}^m \text{clip}[a_i | \frac{\epsilon}{2m}] \leq \sum_{i=1}^m \text{clip}[a_i | \frac{\epsilon}{2m}]$. It therefore suffices to show that $\sum_{i=1}^m a_i \leq 2 \sum_{i=i^*}^m a_i$. To this end, we see that, since $a_i \leq \frac{\epsilon}{2m}$ for $i < i^*$, $\sum_{i=1}^{i^*-1} a_i \leq \sum_{i=1}^{i^*-1} \frac{\epsilon}{2m} \leq \frac{(i^*-1)\epsilon}{2m} \leq \epsilon/2$. On the other hand, since $\sum_{i=1}^m a_i \geq \epsilon$, we must have that $\sum_{i=i^*}^m a_i \geq \frac{\epsilon}{2}$, and thus $\sum_{i=1}^m a_i \leq 2 \sum_{i=i^*}^m a_i$, as needed. \square

Applying the above lemma careful, we arrive at the following useful regret decomposition:

Lemma B.6 (Clipped Regret Decomposition Lead and Future Bounds). *Let $\mathfrak{g}\check{\mathbf{a}}\mathfrak{p}_{\min} = \min_{x,a,h} \mathfrak{g}\check{\mathbf{a}}\mathfrak{p}_h(x,a)$. Then on the event $\mathcal{A}^{\text{conc}}$ the regret of StrongEuler is bounded by*

$$\text{Regret}_K \leq 4e \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x,a) \text{clip} \left[\mathbf{c}\mathbf{B}_{k,h}^{\text{lead}}(x,a) \mid \frac{\mathfrak{g}\check{\mathbf{a}}\mathfrak{p}_h(x,a)}{4} \right].$$

$$+ 8He \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x,a) \text{clip} \left[\mathbf{c}\mathbf{B}_k^{\text{fut}}(x',a') \mid \frac{\mathfrak{g}\check{\mathbf{a}}\mathfrak{p}_{\min}}{8SAH} \right],$$

where \mathbf{c} is a universal constant.

B.3 Proof of Corollary B.1: A proof via integration

Note that Lemma B.6 bounds Regret_K by a sum of *local* bounds terms $\mathbf{B}_{k,h}^{\text{lead}}(x,a)$ and $\mathbf{B}_k^{\text{fut}}$, which depend only on the number of samples $n_k(x,a)$ obtained from state action pair (x,a) . More precisely, we can represent the bound terms by defining the functions

$$g_{x,a}^{\text{lead}}(u) := \sqrt{\frac{\text{Var}_{x,a}^* \log\left(\frac{Mu}{\delta}\right)}{u}}, \quad \text{and} \quad g^{\text{fut}}(u) = H^3 \left(\sqrt{\frac{S \log\left(\frac{Mu}{\delta}\right)}{u}} + \frac{S \log\left(\frac{Mu}{\delta}\right)}{u} \right)^2.$$

Further, define $\epsilon_{x,a}^{\text{lead}} := \min_h \frac{\bar{g}_{\text{ap}_h(x,a)}}{4}$, $\epsilon^{\text{fut}} := \frac{\bar{g}_{\text{ap}_{\min}}}{8SAH}$, and lastly set

$$f_{x,a}^{\text{lead}}(u) := H \wedge \text{clip} \left[\mathbf{c}g_{x,a}^{\text{lead}}(u) | \epsilon_{x,a}^{\text{lead}} \right], \quad f^{\text{fut}}(u) := H^3 \wedge \text{clip} \left[\mathbf{c}g^{\text{fut}}(u) | \epsilon^{\text{fut}} \right].$$

Then, recalling the definitions of \mathbf{B}^{lead} , \mathbf{B}^{fut} , and the fact that $\mathbf{B}_{k,h}^{\text{lead}}(x,a) \leq H$ and $\mathbf{B}_k^{\text{fut}}(x,a) \leq H^3$, we can write

$$\text{Regret}_K \lesssim \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x,a) f_{x,a}^{\text{lead}}(n_k(x,a)) + H \sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x,a) f^{\text{fut}}(n_k(x,a)). \quad (4)$$

As described in Section 3, the crucial step now is to relate the empirical counts $n_k(x,a)$ to the visitation probabilities. Precisely, let us aggregate

$$\omega_k(x,a) := \sum_{h=1}^H \omega_{k,h}(x,a), \quad \bar{n}_k(x,a) := \sum_{j=1}^k \omega_j(x,a).$$

Note that if $\{\mathcal{F}_k\}$ denotes the filtration corresponding to the episodes k , then, $\mathbb{E}[n_k(x,a) \mid \mathcal{F}_{k-1}] = n_{k-1}(x,a) + \omega_{k-1}(x,a)$. In other words, $\bar{n}_{k-1}(x,a)$ is precise the sum of the increments $\mathbb{E}[n_j(x,a) - n_{j-1}(x,a) \mid \mathcal{F}_{j-1}]$ for $j = 1, \dots, k-1$ ⁶. Hence, by a now-standard martingale concentration argument, we find that $n_k(x,a)$ will be lower bounded by $\bar{n}_k(x,a)$, provided that the latter quantity is sufficiently large. More precisely:

Lemma B.7 (Sampling Event). *Define the event*

$$\mathcal{E}^{\text{samp}}(H_{\text{sample}}) := \left\{ \forall(x,a), \forall k \geq \tau_{H_{\text{sample}}}(x,a), \quad n_k(x,a) \geq \frac{1}{4} \bar{n}_k(x,a) \right\},$$

where $\tau_n(x,a) := \inf \{k : \bar{n}_k(x,a) \geq n\}$

Then, for some $H_{\text{sample}} \lesssim H \log \frac{M}{\delta}$, $\mathcal{E}^{\text{samp}}(H_{\text{sample}})$ holds with probability at least $1 - \delta/2$.

Lemma B.7 is proved in Appendix C.2 as a consequence of Lemma 6 in [6]. Together, the events $\mathcal{A}^{\text{conc}}$ and $\mathcal{E}^{\text{samp}}$ account for $1 - \delta$ probability with which our regret bounds hold. For short, we will let $\mathcal{E}^{\text{samp}}$ denote $\mathcal{E}^{\text{samp}}(H_{\text{sample}})$ when clear from context, and $\tau := \tau_{H_{\text{sample}}}$.

After neglecting the first $\tau(x,a)$ samples in the sum (5), we can approximately bound

$$\begin{aligned} \text{Regret}_K &\lesssim \sum_{x,a} \sum_{k \geq \tau(x,a)}^K \sum_{h=1}^H \omega_{k,h}(x,a) f_{x,a}^{\text{lead}}(\bar{n}_k(x,a)/4) \\ &\quad + \sum_{x,a} H \sum_{k=\tau(x,a)}^K \sum_{h=1}^H \omega_{k,h}(x,a) f^{\text{fut}}(\bar{n}_k(x,a)/4), \end{aligned}$$

where \lesssim denotes an informal inequality. Now, $\omega_{k,h}(x,a)$ and $\bar{n}_k(x,a)/4$ are directly related via $\bar{n}_k(x,a)/4 := \sum_{j=1}^k \sum_{h=1}^H \omega_{j,h}(x,a)$. Hence, we can view the above regret bounds as discrete integrals of the functions $f_{x,a}^{\text{lead}}$ and $f^{\text{fut}}(\bar{n}_k(x,a)/4)$. This argument is made precise by the following lemma, which comprises the workhorse of our argument:

Lemma B.8 (Integral Conversion). *Suppose that the event $\mathcal{E}^{\text{samp}}(H_{\text{sample}})$ holds. Then, for any collection of functions $f_{x,a}(\cdot)$ non-increasing functions from $\mathbb{N} \rightarrow \mathbb{R}$ bounded above by f_{\max} and any $\epsilon_{x,a,h} \geq 0$, we have that*

$$\sum_{k=1}^K \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x,a) f_{x,a}(n_k(x,a)) \leq 2ASH_{\text{sample}} f_{\max} + \sum_{x,a} \mathbb{I}(\bar{n}_K(x,a) \geq H) \int_H^{\bar{n}_K(x,a)} f(u/4) du$$

⁶Note that we induce $\bar{n}_k(x,a)$ to include a sum up to index k ; this makes the following arguments more convenient, and will only accrue constant factors in the analysis

Since $H_{\text{sample}} \lesssim H \log(M/\delta)$, and the functions $f_{x,a}^{\text{fut}}$, $f_{x,a}^{\text{lead}}$ are bounded by H^3 , we see that, on $\mathcal{E}^{\text{samp}} \cap \mathcal{A}^{\text{conc}}$, it holds that

$$\begin{aligned} \text{Regret}_K &\lesssim SAH^5 \log(M/\delta) + \sum_{x,a} \mathbb{I}(\bar{n}_K(x,a) \geq H) \int_H^{\bar{n}_K(x,a)} f_{x,a}^{\text{lead}}(u/4) du \\ &\quad + \sum_{x,a} \mathbb{I}(\bar{n}_K(x,a) \geq H) \int_H^T f_{x,a}^{\text{fut}}(u/4) du, \end{aligned} \quad (5)$$

where for the term on the second line, we have bounded $\bar{n}_K(x,a) \leq T$ and used that $f^{\text{fut}}(\cdot) \geq 0$.

All that remains is to evaluate the above integrals. This is directly addressed by the following technical lemma, proved in Section C.5:

Lemma B.9 (General Integration Computations). *Let $f(u) \leq \min\{f_{\max}, \text{clip}[g(u)|\epsilon]\}$ where $\epsilon \in [0, H]$ and $g(u)$ is a non-increasing function is specified in each of two cases that follow. Further, let $M \geq 1$, and $\delta \in (0, 1/2)$ be problem dependent constants. Finally, let \lesssim denote inequality up to a problem independent constant. Then, the following integral computations hold:*

- (a) *Suppose that $C > 0$ is a problem dependent constant satisfying $\log C \lesssim \log(2M)$, and that $g(u) \lesssim \sqrt{\frac{C \log(Mu/\delta)}{u}}$. Then,*

$$\int_H^N f(u/4) du \lesssim \min \left\{ \sqrt{CN \log \frac{MN}{\delta}}, \frac{C}{\epsilon} \log \left(\frac{M}{\delta} \cdot \min\{T, \frac{H}{\epsilon}\} \right) \right\}.$$

- (b) *Suppose that $C, C' > 0$ are a problem dependent constant satisfying $\log(CC') \lesssim \log 2M$, and that $g(u) \lesssim C \left(\sqrt{\frac{C' \log(Mu/\delta)}{u}} + \frac{C' \log(Mu/\delta)}{u} \right)^2$. Then,*

$$\begin{aligned} \int_H^N f(u/4) du &\lesssim (1 + C') f_{\max} \log\left(\frac{M}{\delta}\right) \\ &\quad + CC' \log\left(\frac{MN}{\delta}\right) \min \left\{ \log \frac{MN}{\delta}, \log \left(\frac{MH}{\epsilon}\right) \right\} \end{aligned}$$

Note that the special case $g(u) \lesssim \frac{C}{\log(Mu/\delta)} u$ can be obtained by setting $C' = 1$ in the above inequality.

Lastly, the above computations hold if $f(u/4)$ is replaced by $f(u/c)$ for any universal constant $c > 0$. Moreover, the above computations hold if $f(u) \lesssim \min\{f_{\max}, g(u)\}$ by taking $\epsilon = 0$ and setting $\frac{1}{\epsilon} = \infty$.

Remark B.5 (Integration without anytime bounds). *If instead we consider functions $g(u)$ satisfying the looser bounds (a) $g(u) \lesssim \sqrt{\frac{C \log(MT/\delta)}{u}}$ and (b) $g(u) \lesssim C \left(\sqrt{\frac{C' \log(MT/\delta)}{u}} + \frac{C' \log(MT/\delta)}{u} \right)^2$ for $T \geq N$, then we can recover the bounds*

$$\int_H^N f(u/4) du \lesssim \begin{cases} \min \left\{ \sqrt{CN \log \frac{MT}{\delta}}, \frac{C}{\epsilon} \log \frac{MT}{\delta} \right\} & \text{case (a)} \\ (1 + C') f_{\max} \log \frac{MT}{\delta} + CC' \log \left(\frac{MT}{\delta}\right) \min \left\{ \log \frac{MT}{\delta}, \log \frac{\log(MT/\delta)}{\epsilon} \right\} & \text{case (b)} \end{cases}$$

These sorts of bounds arise when the confidence intervals are derived via union bounds over all time T , rather than via anytime estimates. In particular, we see that using a naive union bounded over all time T incurs a dependence on $\log T \cdot (\log \log T)$, and thus does not imply a strictly $\mathcal{O}(\log T)$ regret.

Let us conclude by applying the above lemma to the terms at hand. First, applying the Part (a) with $f = f_{x,a}^{\text{lead}}$, $g = g_{x,a}^{\text{lead}}$, $C = \text{Var}_{x,a}^*$, and $H \geq \epsilon = \epsilon_{x,a}^{\text{lead}} := \min_h \frac{\text{gap}_h(x,a)}{4} \gtrsim \frac{\text{gap}(x,a)}{(1 \vee \alpha)H}$ for

$(x, a) \in \mathcal{Z}_{\text{sub}}$, and $H \geq \epsilon x$, $a \gtrsim \frac{\text{gap}_{\min}}{H}$ for $(x, a) \in \mathcal{Z}_{\text{opt}}$, we have that

$$\begin{aligned}
& \sum_{x,a} \mathbb{I}(\bar{n}_K(x, a) \geq H) \int_H^{\bar{n}_K(x, a)} f_{x,a}^{\text{lead}}(u/4) du \\
& \lesssim \sum_{x,a} \frac{\text{Var}_{x,a}^*}{\min_h \check{\text{gap}}_h(x, a)} \log\left(\frac{MH}{\delta \min_h \check{\text{gap}}_h(x, a)}\right) \\
& \lesssim \sum_{(x,a) \in \mathcal{Z}_{\text{sub}}} \frac{\text{Var}_{x,a}^*(1 \vee \alpha H)}{\text{gap}(x, a)} \log \frac{MH}{\delta \text{gap}(x, a)} + \sum_{(x,a) \in \mathcal{Z}_{\text{opt}}} \frac{\text{Var}_{x,a}^* H}{\text{gap}_{\min}(x, a)} \log \frac{MH}{\delta \text{gap}_{\min}} \\
& \leq \sum_{(x,a) \in \mathcal{Z}_{\text{sub}}} \frac{\text{Var}_{x,a}^*(1 \vee \alpha H)}{\text{gap}(x, a)} \log \frac{MH}{\delta \text{gap}(x, a)} + |\mathcal{Z}_{\text{opt}}| \frac{\bar{\text{Var}} H}{\text{gap}_{\min}(x, a)} \log \frac{MH}{\delta \text{gap}_{\min}}.
\end{aligned}$$

Similarly, applying the Part (b) with $f = f^{\text{fut}}$, $g = g^{\text{fut}}$, $C' = S$ and $C = H^3$, and $\epsilon = \epsilon^{\text{fut}} := \frac{\check{\text{gap}}_{\min}}{8SAH} \gtrsim \frac{\text{gap}_{\min}}{H^2}$ (and also satisfying $\epsilon^{\text{fut}} \leq H$), we can bound

$$\lesssim S^2 AH^3 \log\left(\frac{MT}{\delta}\right) \min\left\{\log \frac{MT}{\delta}, \log\left(\frac{MH}{\text{gap}_{\min}}\right)\right\},$$

Plugging the above two displays into (5) concludes the proof of Corollary B.1.

B.4 Proof of Theorem B.2

We conclude the section by proving the regret bound of Theorem B.2, which interpolates between the \sqrt{T} and $\log T$ regimes. Let us recall the subset $\mathcal{Z}_{\text{sub}}(\epsilon) := \{(x, a) : \text{gap}(x, a) < \epsilon\}$, as well as $\bar{H}_T := \min\left\{\bar{\text{Var}}, \frac{\mathcal{G}^2}{H} \log T\right\}$. Retracing the proof of Corollary B.1, it suffices to establish only two points:

$$\begin{aligned}
& \sum_{(x,a) \in \mathcal{Z}_{\text{sub}}(\epsilon)} \sum_{k=1}^K \sum_{h=1}^H \omega_{k,h}(x, a) \text{clip}\left[\mathbf{cB}_{k,h}^{\text{lead}}(x, a) \mid \frac{\check{\text{gap}}_h(x, a)}{4}\right] \\
& \lesssim \sqrt{\bar{H}_T |\mathcal{Z}_{\text{sub}}(\epsilon)| T \log \frac{MT}{\delta}} + SAH^2 \log(M\delta/T) \\
& \sum_{(x,a) \in \mathcal{Z}_{\text{opt}}} \sum_{k=1}^K \sum_{h=1}^H \omega_{k,h}(x, a) \text{clip}\left[\mathbf{cB}_{k,h}^{\text{lead}}(x, a) \mid \frac{\check{\text{gap}}_h(x, a)}{4}\right] \\
& \lesssim \sqrt{\bar{H}_T |\mathcal{Z}_{\text{opt}}| T \log \frac{MT}{\delta}} + SAH^2 \log(M\delta/T)
\end{aligned}$$

For both of these inequalities, we will discard the clipping, and thus the two bounds will be syntactically the same. Hence, let us simply prove the following bound:

$$\sum_{k=1}^K \sum_{h=1}^H \sum_{(x,a) \in \mathcal{Z}_{\text{opt}}} \mathbf{B}_{k,h}^{\text{lead}}(x, a) \lesssim \sqrt{\bar{H}_T |\mathcal{Z}_{\text{opt}}| T \log \frac{MT}{\delta}}.$$

Since $\bar{H}_T := \min\left\{\bar{\text{Var}}, \frac{\mathcal{G}^2}{H} \log T\right\}$, it suffices to prove the above bound first with \bar{H}_T replaced by $\bar{\text{Var}}$, and then replaced by $\frac{\mathcal{G}^2}{H} \log T$.

Bound with $\bar{\text{Var}}$: To obtain a bound involving $\bar{\text{Var}}$, we use the fact that $\mathbf{B}_{k,h}^{\text{lead}}(x, a) \lesssim g^{\text{lead}}(n_k(x, a))$, for the function $g^{\text{lead}}(u) = \sqrt{\bar{\text{Var}}} \log(Mu/\delta)/u$. Hence, following the integration arguments in the proof of Corollary B.1, clipped at $\epsilon = 0$, we can bound

$$\begin{aligned}
\sum_{k=1}^K \sum_{h=1}^H \omega_{k,h}(x, a) \mathbf{B}_{k,h}^{\text{lead}}(x, a) & \leq H \log(M/\delta) + \sqrt{\bar{\text{Var}} \bar{n}_K(x, a) \log(\bar{n}_K(x, a) M/\delta)} \\
& \leq H^2 \log(M/\delta) + \sqrt{\bar{\text{Var}} \bar{n}_K(x, a) \log(TM/\delta)}.
\end{aligned}$$

Hence, by Cauchy Schwartz, and the bound $\sum_{(x,a) \in \mathcal{Z}_{\text{opt}}} \bar{n}_K(x,a) \leq \sum_{(x,a)} \bar{n}_K(x,a) = T$,

$$\begin{aligned} \sum_{(x,a) \in \mathcal{Z}_{\text{opt}}} \sum_{k=1}^K \sum_{h=1}^H \omega_{k,h}(x,a) \mathbf{B}_{k,h}^{\text{lead}}(x,a) &\lesssim SAH \log(M/\delta) + \sum_{x,a \in \mathcal{Z}_{\text{opt}}} \sqrt{\text{Var} \bar{n}_K(x,a) \log(TM/\delta)} \\ &\leq SAH \log(M/\delta) + \sqrt{\text{Var} |\mathcal{Z}_{\text{opt}}| \sum_{x,a \in \mathcal{Z}_{\text{opt}}} \bar{n}_K(x,a) \log(TM/\delta)} \\ &\leq SAH \log(M/\delta) + \sqrt{\text{Var} |\mathcal{Z}_{\text{opt}}| T \log(TM/\delta)}, \end{aligned}$$

as needed.

Bound with \bar{H}_T : This bound requires a little more subtlety. Define the function $f(u) = (1/\max\{u, 1\})$. Then, using the definition of $\mathbf{B}_{k,h}^{\text{lead}}(x,a)$ from Proposition B.4, we have

$$\begin{aligned} \sum_{(x,a) \in \mathcal{Z}_{\text{opt}}} \sum_{h=1}^H \mathbf{B}_{k,h}^{\text{lead}}(x,a) &\lesssim \sum_{k=1}^K \sum_{h=1}^H \sum_{(x,a) \in \mathcal{Z}_{\text{opt}}} \omega_{k,h}(x,a) (H \wedge \sqrt{\log(Mn_k(x,a)/\delta) f(n_k(x,a)) \text{Var}_{h,x,a}^*}) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \sum_{(x,a) \in \mathcal{Z}_{\text{opt}}} \omega_{k,h}(x,a) (H \wedge \sqrt{\log(MT/\delta) f(n_k(x,a)) \text{Var}_{h,x,a}^*}). \end{aligned}$$

Applying the recipe we used for Corollary B.1 will not quite carry over in this setting. Instead, we apply an argument based on Cauchy-Schwartz, deferred to Section C.4:

Lemma B.10 (Cauchy-Schwartz Integration Lemma for \mathcal{G} -bounds). *Let $\{V_{x,a,k,h}\}$ be a sequence of numbers, and let $f(u)$ be a nonnegative, non-decreasing function, $f_{\max} > 0$, \mathbf{L} a problem dependent parameter, and let $\mathcal{Z}_0 \subset \mathcal{S} \times \mathcal{A}$. Then, on $\mathcal{E}^{\text{samp}}$,*

$$\begin{aligned} \sum_{(x,a) \in \mathcal{Z}_0} \sum_{k=1}^K \sum_{h=1}^H \omega_{k,h}(x,a) f_{\max} \wedge \sqrt{\mathbf{L} f(n_k(x,a)) V_{x,a,k,h}} \\ \leq |\mathcal{Z}_0| H_{\text{sample}} f_{\max} + \sqrt{\mathbf{L} \sum_{k=1}^K \mathbb{E}^{\pi_k} \left[\sum_{h=1}^H V_{x,a,k,h} \right] \cdot |\mathcal{Z}| (Hf(H) + \int_1^T f(u) du)}. \end{aligned}$$

We apply the above lemma with $V_{x,a,k,h} = \text{Var}_{x,a,h}^{\pi_k}$, $f_{\max} = H$ and $f(u) = 1/\max\{u, 1\}$, and $\mathbf{L} = \log(MT/\delta)$. It is easy to see that the term $|\mathcal{Z}_0| H_{\text{sample}} f_{\max} \lesssim SAH^2 \log(M/\delta)$ will already be absorbed into terms already present in the final bound. On the other hand, by a now-standard law of total variance argument,

$$\sum_{k=1}^K \mathbb{E}^{\pi_k} \left[\sum_{h=1}^H \text{Var}_{x_h, a_h, h}^{\pi_k} \right] \leq K \max_{\pi} \mathbb{E}^{\pi} \left[\sum_{h=1}^H \text{Var}_{x_h, a_h, h}^{\pi} \right] \leq T\mathcal{G}^2/H, \quad (6)$$

where the last inequality is from the proof of [16], Proposition 6. On the other hand, we can bound $(Hf(H) + \int_1^T f(u) du) \leq 1 + \log T$. This finally yields

$$\sum_{k=1}^K \mathbb{E}^{\pi_k} \left[\sum_{h=1}^H \text{Var}_{x_h, a_h, h}^{\pi_k} \right] \lesssim SAH^2 \log(M/\delta) + \sqrt{\log(MT/\delta) T \cdot \mathcal{G}^2/H \log(T)},$$

as needed.

C Proof of Technical Lemmas

C.1 Proof of clipping with future bounds, Lemma B.6

Since strong optimism holds on $\mathcal{A}^{\text{conc}}$, Theorem B.3 yields

$$\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq 2e \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x,a) \text{clip} [\mathbf{E}_{k,h}(x,a) \mid \mathbf{g}\check{\text{a}}\text{p}_h(x,a)].$$

Applying Lemma B.5 with $m = 2$, $a_1 = \mathbf{cB}_{k,h}^{\text{lead}}(x, a)$, and

$$a_2 = \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{cB}_k^{\text{fut}}(x, a) \mid (x_h, a_h) = (x, a) \right],$$

$$\begin{aligned} \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} &\leq 4e \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \text{clip} \left[\mathbf{cB}_{k,h}^{\text{lead}}(x, a) \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_h(x, a)}{4} \right]. \\ &\quad 4e \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \text{clip} \left[\mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{cB}_k^{\text{fut}}(x_t, a_t) \mid (x_h, a_h) = (x, a) \right] \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_h(x, a)}{4} \right]. \end{aligned}$$

The term on the right hand side of the first line of the above display is exactly as needed. Let us turn our attention to the term on the second line. We have that

$$\mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{cB}_k^{\text{fut}}(x_t, a_t) \mid (x_h, a_h) = (x, a) \right] = \sum_{x', a'} \sum_{t=h}^H \mathbf{cB}_k^{\text{fut}}(x', a') \mathbb{P}[(x_t, a_t) = (x', a') \mid (x_h, a_h) = (x, a)].$$

Hence, applying Lemma B.5 with the terms a_i -terms corresponding to $\mathbf{B}_k^{\text{fut}}(x', a') \mathbb{P}[(x_t, a_t) = (x', a') \mid (x_h, a_h) = (x, a)]$ and the number of such terms m bounded by $8SAH$, we have

$$\begin{aligned} &\text{clip} \left[\mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{B}_k^{\text{fut}}(x_t, a_t) \mid (x_h, a_h) = (x, a) \right] \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_h(x, a)}{4} \right] \\ &\leq 2 \sum_{x', a'} \sum_{t=h}^H \text{clip} \left[\mathbf{cB}_k^{\text{fut}}(x', a') \mathbb{P}[(x_t, a_t) = (x', a') \mid (x_h, a_h) = (x, a)] \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_h(x, a)}{8SAH} \right]. \end{aligned}$$

Since $\text{clip}[\alpha x] \leq \alpha \text{clip}[x]$ for $\alpha \leq 1$, and since the probabilities $\mathbb{P}[(x_t, a_t) = (x', a') \mid (x_h, a_h) = (x, a)]$ are bounded by 1, we can bound the above by

$$2 \sum_{x', a'} \sum_{t=h}^H \mathbb{P}[(x_t, a_t) = (x', a') \mid (x_h, a_h) = (x, a)] \text{clip} \left[\mathbf{cB}_k^{\text{fut}}(x', a') \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_h(x, a)}{8SAH} \right].$$

Hence,

$$\begin{aligned} &4e \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \text{clip} \left[\mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{cB}_k^{\text{fut}}(x_t, a_t) \mid (x_h, a_h) = (x, a) \right] \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_h(x, a)}{4} \right] \\ &\leq 8e \sum_{h=1}^H \sum_{x,a} \sum_{x', a'} \omega_{k,h}(x, a) \sum_{t=h}^H \mathbb{P}[(x_t, a_t) = (x', a') \mid (x_h, a_h) = (x, a)] \text{clip} \left[\mathbf{cB}_k^{\text{fut}}(x', a') \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_h(x, a)}{8SAH} \right] \\ &\leq 8e \sum_{h=1}^H \sum_{x,a} \sum_{x', a'} \omega_{k,h}(x, a) \sum_{t=h}^H \mathbb{P}[(x_t, a_t) = (x', a') \mid (x_h, a_h) = (x, a)] \text{clip} \left[\mathbf{cB}_k^{\text{fut}}(x', a') \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_{\min}}{8SAH} \right] \\ &= 8e \sum_{h=1}^H \sum_{t=h}^H \sum_{x', a'} \omega_{k,h}(x', a') \text{clip} \left[\mathbf{cB}_k^{\text{fut}}(x', a') \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_{\min}}{8SAH} \right] \\ &\leq 8He \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \text{clip} \left[\mathbf{cB}_k^{\text{fut}}(x', a') \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_{\min}}{8SAH} \right]. \end{aligned}$$

Altogether,

$$\begin{aligned} \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} &\leq 4e \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \text{clip} \left[\mathbf{cB}_{k,h}^{\text{lead}}(x, a) \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_h(x, a)}{4} \right]. \\ &\quad + 8He \sum_{h=1}^H \sum_{x,a} \omega_{k,h}(x, a) \text{clip} \left[\mathbf{cB}_k^{\text{fut}}(x', a') \mid \frac{\mathbf{g}\check{\mathbf{a}}\mathbf{p}_{\min}}{8SAH} \right]. \end{aligned}$$

Summing over $k = 1, \dots, K$ proves the inequality.

C.2 Proof of sampling lemma (Lemma B.7)

Recall $(\mathcal{E}^{\text{samp}})' := \{\forall k, s, a : n_k(x, a) \geq \frac{1}{2}\bar{n}_{k-1}(x, a) - H \log \frac{2HSA}{\delta}\}$; Lemma 6 in [6] shows that this event occurs with probability at least $1 - \delta/2$. We show that $(\mathcal{E}^{\text{samp}})' \subseteq \mathcal{E}^{\text{samp}}(H_{\text{sample}})$, for $H_{\text{sample}} = 4H \log \frac{2HSA}{\delta} \lesssim H \log \frac{M}{\delta}$.

Noting that $\bar{n}_k \leq \bar{n}_{k-1} + H$, $(\mathcal{E}^{\text{samp}})'$ implies that $n_k \geq \frac{1}{2}\bar{n}_k(x, a) - H \log \frac{2HSA}{\delta} - H = \frac{1}{2}\bar{n}_k(x, a) - H \log \frac{2eHSA}{\delta}$. Hence, for any $k \geq \tau(x, a)$, we have $\bar{n}_k \geq 4H \log \frac{2HSA}{e\delta}$ and thus $n_k(x, a) \geq \frac{\bar{n}_k}{4} + \frac{\bar{n}_k}{4} - H \log \frac{2eHSA}{\delta} \geq \frac{\bar{n}_k}{4}$. Bounding $\log \frac{2HSA}{e\delta} \leq \tilde{L}(1)$ concludes the proof.

C.3 Proof of integral conversion, Lemma B.8

Recall that $\tau(x, a)$ denote $\inf\{k : \bar{n}_k(x, a) \geq H_{\text{sample}}\}$. Then,

$$\begin{aligned} \sum_{k=1}^K \sum_{x,a} \omega_k(x, a) f_{x,a}(n_k(x, a)) &= \sum_{x,a} \sum_{k=1}^{\tau(x,a)-1} \omega_k(x, a) f_{x,a}(n_k(x, a)) + \sum_{x,a} \sum_{k=\tau(x,a)} f_{x,a}(n_k) \\ &\leq \sum_{x,a} \sum_{k=1}^{\tau(x,a)-1} \omega_k(x, a) f_{\max} + \sum_{x,a} \sum_{k=\tau(x,a)} f_{x,a}(n_k) \\ &\leq SAH_{\text{sample}} f_{\max} + \sum_{x,a} \sum_{k=\tau(x,a)} f_{x,a}(n_k(x, a)) \\ &\leq SAH_{\text{sample}} f_{\max} + \sum_{x,a} \sum_{k=\tau(x,a)} f_{x,a}(\bar{n}_k(x, a)/4), \end{aligned}$$

since $\sum_{k=1}^{\tau(x,a)-1} \omega_k(x, a) = \bar{n}_{\tau(x,a)-1}(x, a) \leq H_{\text{sample}} \text{madn } f(\cdot) \leq f_{\max}$. We now appeal to the following integration lemma, which we prove momentarily.

Lemma C.1 (Integration over $\omega_k(x, a)$). *Let $f : [H, \infty) \rightarrow \mathbb{R}_{>0}$ be a non-increasing function. Then,*

$$\sum_{k=\tau(x,a)}^K \omega_k(x, a) f(\bar{n}_k(x, a)) \leq H f(H) + \int_H^{\bar{n}_K(x,a)} f(u) du. \quad (7)$$

To conclude the proof of Lemma B.8, we apply the above for each (x, a) with the functions $f(u) \leftarrow f_{x,a}(u/4)$, and note $H f_{x,a}(H/4) \leq H f_{\max} \leq H_{\text{sample}} f_{\max}$. \square

Proof of Lemma C.1. The proof generalizes Lemma E.5 in [5]. For ease of notation, define $k_0 = \tau(x, a)$. We can define the step function $g : [k_0, K] \rightarrow \mathbb{R}$ via $g(t) = \sum_{k=k_0}^{K-1} \omega_{k+1}(x, a) \mathbb{I}(t \in [k, k+1])$. Then, letting $G(t) := \bar{n}_{k_0}(x, a) + \int_0^t g(u) du$, we see that $G'(t) = g(t)$ almost everywhere, G is non-decreasing, and $G(k) = \bar{n}_k(x, a)$ for all $k \in [k_0, K]$. We can therefore express

$$\begin{aligned} \sum_{k>\tau(x,a)}^K \omega_k(x, a) f(\bar{n}_k(x, a)) &= \sum_{k=k_0+1}^K \omega_k(x, a) f(\bar{n}_k(x, a)) = \sum_{k=k_0+1}^K \left(\int_{k-1}^k g(t) dt \right) f(G(k)). \\ &\stackrel{(i)}{\leq} \sum_{k=k_0+1}^K \left(\int_{k-1}^k g(t) f(G(t)) dt \right) = \int_{k=k_0}^K g(t) f(G(t)) dt \\ &\stackrel{(ii)}{=} \int_{G(k_0)}^{G(K)} f(u) du \stackrel{(iii)}{=} \int_{\bar{n}_{k_0}(x,a)}^{\bar{n}_K(x,a)} f(u) du, \end{aligned}$$

where (i) uses the fact that $f \circ G$ is non-increasing, (ii) is the Fundamental Theorem of Calculus, with $G'(t) = g(t)$, and (iii) is $G(k) = \bar{n}_k(x, a)$ for $k \in [k_0, K]$. Hence, we have the bound

$$\begin{aligned} \sum_{k \geq k_0}^K \omega_k(x, a) f(\bar{n}_k(x, a)) &\leq \omega_{k_0}(x, a) f(\bar{n}_{k_0}(x, a)) + \int_{\bar{n}_{k_0}(x, a)}^{\bar{n}_K(x, a)} f(u) du \\ &\stackrel{(i)}{\leq} H f(\bar{n}_{k_0}(x, a)) + \int_{\bar{n}_{k_0}(x, a)}^{\bar{n}_K(x, a)} f(u) du \stackrel{(ii)}{\leq} H f(H) + \int_H^{\bar{n}_K(x, a)} f(u) du, \end{aligned}$$

where (i) uses $\omega_{k_0} \leq H$, and that $f(u) \geq 0$, and (ii) uses the fact that f is nonincreasing, and $\bar{n}_{k_0}(x, a) \geq H_{\text{sample}} \geq H$. \square

C.4 Proof of interal conversion for \mathcal{G} -bounds, Lemma B.10

Let $\tau(x, a) = \tau_{H_{\text{sample}}}(x, a)$. Then, as in the proof of Lemma B.8,

$$\begin{aligned} &\sum_{(x, a) \in \mathcal{Z}} \sum_{k=1}^K \sum_{h=1}^H \omega_{k, h}(x, a) (f_{\max} \wedge \sqrt{\tilde{\mathbf{L}} f(n_k(x, a)) V_{x, a, k, h}}) \\ &\lesssim \sum_{(x, a) \in \mathcal{Z}} \sum_{k=\tau(x, a)}^K \sum_{h=1}^H \omega_{k, h}(x, a) \sqrt{\tilde{\mathbf{L}} f(n_k(x, a)) V_{x, a, k, h}} + |\mathcal{Z}_{\text{opt}}| H_{\text{sample}} f_{\max}. \end{aligned}$$

By Cauchy-Schwartz

$$\begin{aligned} &\sum_{(x, a) \in \mathcal{Z}} \sum_{k=\tau(x, a)}^K \sum_{h=1}^H \omega_{k, h}(x, a) \tag{1} \\ &\sqrt{\sum_{(x, a) \in \mathcal{Z}} \sum_{k=\tau(x, a)}^K \sum_{h=1}^H \omega_{k, h}(x, a) V_{x, a, k, h}} \times \mathbf{L}^{1/2} \sqrt{\sum_{(x, a) \in \mathcal{Z}} \sum_{k=\tau(x, a)}^K \sum_{h=1}^H \omega_{k, h}(x, a) f(n_k(x, a))}. \tag{2} \end{aligned}$$

The first term in the above product can be bounded as

$$\sum_{(x, a) \in \mathcal{Z}} \sum_{k=\tau(x, a)}^K \sum_{h=1}^H \omega_{k, h}(x, a) V_{x, a, k, h} \leq \sum_{k=1}^K \sum_{h=1}^H \sum_{x, a} \omega_{k, h}(x, a) V_{x, a, k, h} = \sum_{k=1}^K \mathbb{E}^{\pi_k} \left[\sum_{h=1}^H V_{x, a, k, h} \right].$$

Using Lemma C.1, the second term can be bounded as

$$\sqrt{\sum_{(x, a) \in \mathcal{Z}} \sum_{k=\tau(x, a)}^K \sum_{h=1}^H \omega_{k, h}(x, a) f(n_k(x, a))} \leq |\mathcal{Z}| (H f(H) + \int_1^T f(u) du).$$

C.5 General Integral computations (Lemma B.9)

For convenience, let us restate the lemma we are about to prove.

Lemma B.9 (General Integration Computations). *Let $f(u) \leq \min\{f_{\max}, \text{clip}[g(u)|\epsilon]\}$ where $\epsilon \in [0, H]$ and $g(u)$ is a non-increasing function is specified in each of two cases that follow. Further, let $M \geq 1$, and $\delta \in (0, 1/2)$ be problem dependent constants. Finally, let \lesssim denote inequality up to a problem independent constant. Then, the following integral computations hold:*

- (a) *Suppose that $C > 0$ is a problem dependent constant satisfying $\log C \lesssim \log(2M)$, and that $g(u) \lesssim \sqrt{\frac{C \log(Mu/\delta)}{u}}$. Then,*

$$\int_H^N f(u/4) du \lesssim \min \left\{ \sqrt{CN \log \frac{MN}{\delta}}, \frac{C}{\epsilon} \log \left(\frac{M}{\delta} \cdot \min\{T, \frac{H}{\epsilon}\} \right) \right\}.$$

(b) Suppose that $C, C' > 0$ are a problem dependent constant satisfying $\log(CC') \lesssim \log 2M$, and that $g(u) \lesssim C \left(\sqrt{\frac{C' \log(Mu/\delta)}{u}} + \frac{C' \log(Mu/\delta)}{u} \right)^2$. Then,

$$\int_H^N f(u/4) du \lesssim (1 + C') f_{\max} \log\left(\frac{M}{\delta}\right) + CC' \log\left(\frac{MN}{\delta}\right) \min\left\{\log\frac{MN}{\delta}, \log\left(\frac{MH}{\epsilon}\right)\right\}$$

Note that the special case $g(u) \lesssim \frac{C}{\log(Mu/\delta)} u$ can be obtained by setting $C' = 1$ in the above inequality.

Lastly, the above computations hold if $f(u/4)$ is replaced by $f(u/c)$ for any universal constant $c > 0$. Moreover, the above computations hold if $f(u) \lesssim \min\{f_{\max}, g(u)\}$ by taking $\epsilon = 0$ and setting $\frac{1}{\epsilon} = \infty$.

Proof. By inflating C by a problem-independent constant if necessary, we may assume without loss of generality that $g(u) = \sqrt{C \log(Mu/\delta)}/u$ in part (a) and $g(u) = C(\sqrt{C' \log(Mu/\delta)}/u + \sqrt{C' \log(Mu/\delta)}/u)^2$, with equality rather than approximate inequality \lesssim .

Next, define

$$n_{\text{end}} := \begin{cases} \max\{u : g(u/4) \geq \epsilon\} \epsilon > 0 \\ N & \epsilon = 0. \end{cases}$$

Throughout, we shall assume the case $\epsilon > 0$, as the $\epsilon = 0$ can be derived by just taking $n_{\text{end}} = N$. Note then that $f(u/4) = \text{clip}[g(u/4)|\epsilon] = 0$ for all $u > n_{\text{end}}$. Hence, it suffices to upper bound

$$\mathbb{I}(n_{\text{end}} \geq H) \int_H^{N \wedge n_{\text{end}}} f_{\max} \wedge g(u/4) du.$$

Lastly, let us define $\tilde{\mathbf{L}}(u) := \log(Mu/\delta)$ for $u \geq H$. We shall require the following inversion lemma, which is standard in the multi-arm bandits literature.

Lemma C.2 (Inversion Lemma). *There exists a universal constant $c > 0$ such that for all $b \geq 0$, $\tilde{\mathbf{L}}(u)/u \leq b$ as long as $u \geq \tilde{\mathbf{L}}(1 + b^{-1})/cb$. Moreover, for $u \lesssim \tilde{\mathbf{L}}(b^{-1})/cb$, it holds that $\tilde{\mathbf{L}}(u) \lesssim \tilde{\mathbf{L}}(1 + b^{-1})$.*

Proof. Let $u = \tilde{\mathbf{L}}(1/b)/cb$ for a constant c to be chosen shortly. Then,

$$\begin{aligned} \tilde{\mathbf{L}}(u)/u &= cb \frac{\tilde{\mathbf{L}}(\frac{1}{cb} \tilde{\mathbf{L}}(b^{-1}))}{\tilde{\mathbf{L}}(1 + b^{-1})} = cb \frac{\log \frac{1}{c} + \log(M/b\delta) + \log(\log \frac{M}{b\delta})}{\log(M/b\delta)} \\ &\leq \frac{cb \log \frac{1}{c}}{\log 2} + 2cb, \end{aligned}$$

where we use $\log \log(x) \leq x$ and $\tilde{\mathbf{L}}(1 + b^{-1}) \geq \tilde{\mathbf{L}}(1) \geq \log 2$. It is easy to see that this quantity is less than b for a constant c sufficiently small that does not depend on M, δ, b . The second statement follows from an analogous computation. \square

Proof of Part (a): Suppose $g(u) = \sqrt{\frac{C \tilde{\mathbf{L}}(u)}{u}}$. It is straightforward to bound

$$\begin{aligned} \mathbb{I}(n_{\text{end}} \geq H) \int_H^{N \wedge n_{\text{end}}} f_{\max} \wedge g(u/4) du &\lesssim \mathbb{I}(n_{\text{end}} \geq H) \sqrt{C \tilde{\mathbf{L}}(N \wedge n_{\text{end}})} \int_1^{N \wedge n_{\text{end}}} \frac{du}{\sqrt{u}} \\ &\lesssim \sqrt{C \tilde{\mathbf{L}}(N \wedge n_{\text{end}})} \cdot \sqrt{(N \wedge n_{\text{end}})} \\ &\lesssim \min \left\{ \sqrt{CN \tilde{\mathbf{L}}(N)}, \sqrt{C n_{\text{end}} \tilde{\mathbf{L}}(n_{\text{end}})} \right\} \\ &= \min \left\{ \sqrt{CN \log \frac{MN}{\delta}}, \sqrt{C n_{\text{end}} \tilde{\mathbf{L}}(n_{\text{end}})} \right\}. \end{aligned}$$

To conclude, let us find $n_{\text{end}}(x, a)$. By our inversion Lemma C.2, we can see that

$$\begin{aligned} n_{\text{end}} &\lesssim \frac{C}{\epsilon^2} \tilde{\mathbf{L}}\left(1 + \frac{C}{\epsilon^2}\right) \lesssim \frac{C}{\epsilon^2} \tilde{\mathbf{L}}\left(1 + \frac{C}{\epsilon}\right) \\ \tilde{\mathbf{L}}(n_{\text{end}}) &\lesssim \tilde{\mathbf{L}}\left(1 + \frac{C}{\epsilon}\right). \end{aligned}$$

Therefore,

$$\sqrt{C n_{\text{end}} \tilde{\mathbf{L}}(n_{\text{end}})} \lesssim \frac{C}{\epsilon} \tilde{\mathbf{L}}\left(1 + \frac{C}{\epsilon}\right).$$

Moreover, for if $\log C \lesssim \log M$ and $\epsilon \leq H$, we can bound $\tilde{\mathbf{L}}\left(1 + \frac{C}{\epsilon}\right) \lesssim \log \frac{MH}{\epsilon}$. Hence, we haveshow

$$\int_H^N f(u/4) du \lesssim \min \left\{ \sqrt{CN \log \frac{MN}{\delta}}, \frac{C}{\epsilon} \log \left(\frac{M}{\delta} \cdot \frac{H}{\epsilon} \right) \right\}.$$

To conclude, it remains to show that we can replace $\frac{H}{\epsilon}$ with N . For this, we use a simpler argument:

$$\begin{aligned} \int_H^N f(u/4) du &\lesssim \int_H^N \text{clip} \left[\frac{\sqrt{C \tilde{\mathbf{L}}(u)}}{u} \middle| \epsilon \right] \\ &\leq \int_H^N \text{clip} \left[\frac{\sqrt{C \tilde{\mathbf{L}}(N)}}{u} \middle| \epsilon \right] \\ &= \sqrt{\tilde{\mathbf{L}}(T)} \int_H^N \text{clip} \left[\frac{1}{\sqrt{u}} \middle| \epsilon' \right], \quad \text{where } \epsilon' = \sqrt{\tilde{\mathbf{L}}(T)}. \end{aligned}$$

Using similar arguments to above, we can bound $\int_H^N \text{clip} \left[\frac{1}{\sqrt{u}} \middle| \epsilon' \right] \lesssim \frac{1}{\epsilon'}$, yielding the bound $\int_H^N f(u/4) du \lesssim \frac{\tilde{\mathbf{L}}^{1/2}(T)}{\epsilon'} = \frac{\tilde{\mathbf{L}}(T)}{\epsilon}$.

Proof of Part (b): A first step This proof will require slightly more care than part (b). We shall first require the following lemma:

Claim C.3. *In the setting of Lemma B.9, if $g(u) = \frac{C \log \frac{Mu}{\delta}}{u} = \frac{C \tilde{\mathbf{L}}(u)}{u}$, then*

$$\int_H^N f(u/4) du \lesssim f_{\max} \log M + C \log(MT/\delta) \min \left\{ \log \left(\frac{MT}{\delta} \right), \log \left(\frac{MH}{\epsilon} \right) \right\}$$

Proof of Claim C.3. Define $n_0 = 2 + \log(M/\delta)$. Then, we have

$$\begin{aligned} \mathbb{I}(n_{\text{end}} \geq H) \int_H^{N \wedge n_{\text{end}}} f_{\max} \wedge g(u/4) du &\leq f_{\max} n_0 + \mathbb{I}(N \wedge n_{\text{end}} \geq n_0) \cdot \int_{n_0}^{N \wedge n_{\text{end}}} g(u/4) du. \\ &\leq f_{\max} n_0 + \int_{n_0}^{n_0 + N \wedge n_{\text{end}}} g(u/4) du \\ &\lesssim f_{\max} \log(M/\delta) + \int_{n_0}^{n_0 + N \wedge n_{\text{end}}} g(u/4) du. \end{aligned}$$

Now take $g(u) = C \tilde{\mathbf{L}}(u)/u$. Since $\tilde{\mathbf{L}}(u) \lesssim \log(M/\delta) + \log(u)$ for $u \geq n_0 \geq 2$, it is straightforward to bound

$$\begin{aligned} \int_{n_0}^{n_0 + N \wedge n_{\text{end}}} g(u/4) du &\lesssim C \log(M/\delta) \log \left(1 + \frac{N \wedge n_{\text{end}}}{n_0} \right) + C \log^2 \left(1 + \frac{N \wedge n_{\text{end}}}{n_0} \right) \\ &\lesssim C \log(MT/\delta) \log \left(1 + \frac{N \wedge n_{\text{end}}}{n_0} \right), \end{aligned}$$

where in the final inequality, we use $N \leq T$, $M/\delta \geq 2$, and $n_0 \geq 1$. By the same token, we can crudely bound the above by $C \lesssim \log^2(MT/\delta)$.

Let us now develop a more refined bound by taking advantage of n_{end} . By our inversion lemma, we have

$$n_{\text{end}} \lesssim \frac{C}{\epsilon} \tilde{\mathbf{L}} \left(1 + \frac{C}{\epsilon} \right) = \frac{C}{\epsilon} \left(\log 1 + \frac{C}{\epsilon} \right) + \log \frac{M}{\delta}.$$

Since $n_0 = 1 + \log \frac{M}{\delta}$,

$$\frac{n_{\text{end}}}{n_0} \lesssim \frac{C}{\epsilon} \log \left(1 + \frac{C}{\epsilon} \right) + \frac{C \log \frac{M}{\delta}}{\epsilon \log \frac{M}{\delta}} \lesssim \frac{C}{\epsilon} \log \left(1 + \frac{C}{\epsilon} \right).$$

Hence, with some algebra we can bound

$$\log \left(1 + \frac{n_{\text{end}}}{n_0} \right) \lesssim \log \left(1 + \frac{C}{\epsilon} \right)$$

This leads to the more refined bound $\int_{n_0}^{n_0 + N \wedge n_{\text{end}}} g(u/4) du \lesssim C \log \left(1 + \frac{C}{\epsilon} \right) \log(\widetilde{MT})$. Again, since $\log C \lesssim \log M$ and $\epsilon \leq H$, we bound again $\log \left(1 + \frac{C}{\epsilon} \right) \lesssim \log \frac{MH}{\epsilon}$. \square

Concluding the proof of Part (b) Define

$$n_0 := \left\{ \inf u : \sqrt{\frac{C' \tilde{\mathbf{L}}(u)}{u}} \leq 1 \right\}.$$

Then, we have

$$\int_H^N f(u/4) du \leq f_{\max} n_0 + \int_{n_0}^N f_{\max} \wedge \text{clip}[g(u/4)|\epsilon] du.$$

Note that for $u \geq n_0$, $g(u/4) \lesssim h(u/4)$, where $h(u) \leq \frac{CC'}{u} \tilde{\mathbf{L}}(u)$. Hence, applying the bound from Lemma C.3 with $C \leftarrow CC'$, we have

$$\int_H^N f(u/4) du \lesssim f_{\max} \log \widetilde{M} + CC' \log(\widetilde{MT}) \min \left\{ \log \left(\frac{MT}{\delta} \right), \log \left(\frac{MH}{\epsilon} \right) \right\}$$

On the otherhand, by our inversion lemma and using $C' \leq \widetilde{M}^{\mathcal{O}(1)}$, we can bound

$$n_0 \leq C' \tilde{\mathbf{L}}(C') \lesssim C' \log \left(\frac{M}{\delta} \right).$$

Combining these two pieces yields the bound. \square

D Proof of ‘clipping’ bound: Proposition 3.1 / Theorem B.3

In this section, we prove Theorem B.3 (of which Proposition 3.1 in the body is a direct consequence), which allows us to clip the surpluses when they are below a certain value. The center of our analysis is the following lemma, which tells us that if $\text{gap}_h(x, a) > 0$ for a pair (x, a, h) , then either the surplus $\mathbf{E}_{k,h}(x, a)$ is large, or expected difference in value functions at the next stage, $p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^{\pi_k})$, is large:

Lemma D.1 (Fundamental Gap Bound). *Then suppose that Alg is strongly optimistic, and consider a pair (x, a, h) with $\text{gap}_h(x, a) > 0$ which is α -transition optimal. Then*

$$\text{gap}_h(x, a) \leq \mathbf{E}_{k,h}(x, a) + \alpha \cdot p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^{\pi_k}).$$

If Alg is possibly not strongly optimistic, then the above holds still holds $\alpha = 1$.

Lemma D.1 is established in Section D.2. Notice that as α gets close to zero, the above bound implies that when $\mathbf{E}_{k,h}(x, a)$ is much smaller than the $\text{gap}_h(x, a)$, the difference in value functions at the next stage, $p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^{\pi_k})$, must become even larger to compensate. The extreme case is $\alpha = 0$, e.g. in contextual bandits, where the gap always lower bounds the surplus.

Continuing with the proof of Theorem B.3, we begin with the ‘‘half-clipping’’ which clips the surpluses at at most gap_{\min} :

Definition D.1 (Half Clipped Value Function). We define the half-clipped surplus $\ddot{\mathbf{E}}_{k,h}(x, a) := \text{clip}[\mathbf{E}_{k,h}(x, a) \mid \epsilon_{\text{clip}}]$, where $\epsilon_{\text{clip}} := \text{gap}_{\min}/(2H)$. We set $\ddot{\mathbf{V}}_{k,H+1}^{\pi_k}(x) = 0$ for all $x \in \mathcal{S}$, and recursively define

$$\ddot{\mathbf{Q}}_h^{\pi_k}(x, a) = r(x, a) + \ddot{\mathbf{E}}_{k,h}(x, a) + p(x, a)^\top \ddot{\mathbf{V}}_{k,h+1}^{\pi_k}, \quad \ddot{\mathbf{V}}_{k,h}^{\pi_k}(x) := \ddot{\mathbf{Q}}_h^{\pi_k}(x, \pi_{k,h}(x)),$$

denote the value and Q-functions of under π_k associated with MDP whose transitions are transitions $p(\cdot, \cdot)$ and non-stationary rewards $r(x, a) + \ddot{\mathbf{E}}_{k,h}(x, a)$ at stage h .

After the half-clipping has been introduced, it is no longer the case that π_k is optimal for this half clipped MDP. As a result, it is not certain that the half-clipped Q-function for π_k is *optimistic* in the sense that $\ddot{\mathbf{Q}}_h^{\pi_k}(x, a) \geq \mathbf{Q}_h^*(x, a)$. We shall instead show that if $\ddot{\mathbf{V}}_{k,h}^{\pi_k}$ is approximately optimistic, in the sense that its excess relative to \mathbf{V}^{π_k} , $\ddot{\mathbf{V}}_0^{\pi_k,h} - \mathbf{V}_0^{\pi_k}$ is at least a constant factor of the regret $\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}$:

Lemma D.2 (Lower Bound on Half-Clipped Surplus). *For $\epsilon_{\text{clip}} = \text{gap}_{\min}/2H$, it holds that*

$$\ddot{\mathbf{V}}_0^{\pi_k} - \mathbf{V}_0^{\pi_k} = \mathbb{E}^{\pi_k} \left[\sum_{h=1}^H \ddot{\mathbf{E}}_{k,h}(x_h, a_h) \right] \geq \frac{1}{2}(\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}),$$

The above bound is established in Section D.1. Hence, to establish the bound of Theorem B.3, it suffices to bound the gap $\ddot{\mathbf{V}}_0^{\pi_k,h} - \mathbf{V}_0^{\pi_k}$. For a given h , and an $x : \pi_{k,h}(x) \notin \pi_h^*(x)$, let us consider the difference

$$\ddot{\mathbf{V}}_h^{\pi_k}(x) - \mathbf{V}_h^{\pi_k}(x) = \ddot{\mathbf{E}}_{k,h}(x, \pi_{k,h}(x)) + p(x, \pi_{k,h}(x))^\top (\ddot{\mathbf{V}}_{h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}).$$

We now introduce the following lemma, proven Section D.3, which allows us to further clip the bonus for suboptimal actions $a \notin \pi_h^*(x)$, i.e. , actions with $\text{gap}_h(x, a) > 0$:

Lemma D.3 (Gap Clipping). *Suppose either Alg is strongly optimistic and each tuple is $\alpha_{x,a,h}$ -transition suboptimal. Then the fully-clipped surpluses*

$$\check{\mathbf{E}}_{k,h}(x, a) := \text{clip} \left[\mathbf{E}_{k,h}(x, a) \mid \epsilon_{\text{clip}} \vee \frac{\text{gap}_h(x, a)}{4(\alpha_{x,a,h}H \vee 1)} \right]$$

satisfy the bound

$$\ddot{\mathbf{V}}_h^{\pi_k}(x) - \mathbf{V}_h^{\pi_k}(x) \leq \check{\mathbf{E}}_{k,h}(x, \pi_{k,h}(x)) + \left(1 + \frac{1}{H}\right) p(x, \pi_{k,h}(x))^\top (\ddot{\mathbf{V}}_{h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k})$$

If Alg is just optimistic, then the above bound holds with $\alpha_{x,a,h} = 1$.

Unfolding the above lemma, and noting that even when Alg is not strongly optimistic, the clipping ensures that $\check{\mathbf{E}}_{k,h}(x, a) \geq 0$, so that we can bound

$$\begin{aligned}
\check{\mathbf{V}}_{k,0}^{\pi_k} - \mathbf{V}_0^{\pi_k} &= \mathbb{E}^{\pi_k} [\check{\mathbf{V}}_1^{\pi_k}(x_1) - \mathbf{V}_1^{\pi_k}(x_1)] \\
&\leq \mathbb{E}^{\pi_k} [\check{\mathbf{E}}_{k,h}(x_1, a_1) + \left(1 + \frac{1}{H}\right) p(x, \pi_{k,h}(x))^\top (\check{\mathbf{V}}_2^{\pi_k} - \mathbf{V}_2^{\pi_k})] \\
&= \mathbb{E}^{\pi_k} [\check{\mathbf{E}}_{k,h}(x_1, a_1) + \left(1 + \frac{1}{H}\right) (\check{\mathbf{V}}_2^{\pi_k, h}(x_2) - \mathbf{V}_2^{\pi_k}(x_2))] \\
&\leq \mathbb{E}^{\pi_k} \left[\sum_{h=1}^H \left(\prod_{h'=2}^h \left(1 + \frac{1}{H}\right) \right) \check{\mathbf{E}}_{k,h}(x_h, a_h) \right] \leq \left(1 + \frac{1}{H}\right)^H \mathbb{E}^{\pi_k} \left[\sum_{h=1}^H \check{\mathbf{E}}_{k,h}(x_h, a_h) \right] \\
&\leq e \mathbb{E}^{\pi_k} \left[\sum_{h=1}^H \check{\mathbf{E}}_{k,h}(x_h, a_h) \right] = e \sum_{x,a} \sum_{h=1}^H \omega_{k,h}(x, a) \check{\mathbf{E}}_{k,h}(x, a),
\end{aligned}$$

where we recall $\omega_{k,h}(x, a) = \mathbb{P}^{\pi_k}[(x_h, a_h) = (x, a)]$. Combining with our earlier bound $\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq 2(\check{\mathbf{V}}_{k,0}^{\pi_k}(x) - \mathbf{V}_0^{\pi_k}(x))$ from Lemma D.2, we find that $\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} \leq 2e \sum_{x,a} \sum_{h=1}^H \omega_{k,h}(x, a) \check{\mathbf{E}}_{k,h}(x, a)$, thereby demonstrating Theorem B.3.

D.1 Proof of Lemma D.2

We can with a crude comparison between the clipped and optimistic value functions.

Lemma D.4. *We have that $\check{\mathbf{E}}_{k,h}(x, \pi_{k,h}(x)) \geq \mathbf{E}_{k,h}(x, \pi_{k,h}(x)) - \epsilon_{\text{clip}}$, which implies*

$$\check{\mathbf{V}}_{k,h}^{\pi_k}(x) + (H - h + 1)\epsilon_{\text{clip}} \geq \bar{\mathbf{V}}_{k,h}(x) \geq \mathbf{V}_h^{\pi_k}(x). \quad (8)$$

Proof. The bound $\check{\mathbf{E}}_{k,h}(x, \pi_{k,h}(x)) \geq \mathbf{E}_{k,h}(x, \pi_{k,h}(x)) - \epsilon_{\text{clip}}$ follows directly from

$$\check{\mathbf{E}}_{k,h}(x, a) = \mathbf{E}_{k,h}(x, a) \mathbb{I}(\mathbf{E}_{k,h}(x, a) \geq \epsilon_{\text{clip}}) \geq \mathbf{E}_{k,h}(x, a) - \epsilon_{\text{clip}}.$$

Hence,

$$\begin{aligned}
\check{\mathbf{V}}_{k,h}^{\pi_k}(x) - \mathbf{V}_h^{\pi_k}(x) &\stackrel{(i.a)}{=} \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \check{\mathbf{E}}_{k,t}(x_t, \pi_{k,h}(x_t)) \mid x_h = x \right] \\
&\geq \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{E}_{k,t}(x_t, \pi_{k,h}(x_t)) - \epsilon_{\text{clip}} \mid x_h = x \right] \\
&= \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{E}_{k,t}(x_t, \pi_{k,h}(x_t)) \right] - (H - h + 1)\epsilon_{\text{clip}} \\
&\stackrel{(i.b)}{=} \bar{\mathbf{V}}_{k,h}(x) - \mathbf{V}_h^{\pi_k}(x) - (H - h + 1)\epsilon_{\text{clip}},
\end{aligned}$$

where (i.a) and (i.b) follow by recursively unfolding the identities $\check{\mathbf{V}}_{k,h}^{\pi_k}(x) - \mathbf{V}_h^{\pi_k}(x) = \check{\mathbf{E}}_{k,h}(x, a) + p(x, a)^\top (\check{\mathbf{V}}_{k,h+1}^{\pi_k}(x) - \mathbf{V}_h^{\pi_k}(x))$ and $\bar{\mathbf{V}}_{k,h}(x) - \mathbf{V}_h^{\pi_k}(x) = \mathbf{E}_{k,h}(x, a) + p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1}(x) - \mathbf{V}_h^{\pi_k}(x))$. □

We now turn to proving Lemma D.2.

Proof. The strategy is as follows. We shall introduce the events over \mathbb{P}^{π_k} , $\mathcal{E}_h := \{\pi_{k,h}(x_h) \notin \pi_h^*(x_h)\}$, which is the event that the policy $\pi_{k,h}$ does not prescribe an optimal action x_h . We further define the events

$$\mathcal{A}_h = \mathcal{E}_h \cap \bigcap_{h' < h} \mathcal{E}_{h'}^c,$$

which is the event that the policy π_k agrees with an optimal action on x_1, \dots, x_{h-1} , and disagrees on x_h . Below, our goal will be to establish the following two formulae for the suboptimality gap $\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}$ and $\ddot{\mathbf{V}}_0^{\pi_k} - \mathbf{V}_0^{\pi_k}$:

$$\ddot{\mathbf{V}}_0^{\pi_k} - \mathbf{V}_0^{\pi_k} \geq \sum_{h=1}^H \mathbb{E}^{\pi_k} [\mathbb{I}(\mathcal{A}_h) \{ \text{gap}(x_h, \pi_{k,h}(x_h)) - H\epsilon_{\text{clip}} + \mathbf{Q}_h^*(x_h, \pi_{k,h}(x_h)) - \mathbf{V}_h^{\pi_k}(x_h) \}] \quad (9)$$

and

$$\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k} = \sum_{h=1}^H \mathbb{E}^{\pi_k} [\mathbb{I}(\mathcal{A}_h) \{ \text{gap}(x_h, \pi_{k,h}(x_h)) + \mathbf{Q}_h^*(x_h, \pi_{k,h}(x_h)) - \mathbf{V}_h^{\pi_k}(x_h) \}] \quad (10)$$

Note that on \mathcal{A}_h , $\mathcal{E}_h = \{ \pi_{k,h}(x_h) \notin \pi_h^*(x_h) \}$ also occurs, and therefore $\text{gap}(x_h, \pi_{k,h}(x_h)) \geq \text{gap}_{\min}$. In particular, displays (9) and (10) both imply

$$\begin{aligned} \ddot{\mathbf{V}}_0^{\pi_k} - \mathbf{V}_0^{\pi_k} &\stackrel{(i)}{\geq} \sum_{h=1}^H \mathbb{E}^{\pi_k} [\mathbb{I}(\mathcal{A}_h) \left\{ \frac{1}{2} \text{gap}(x_h, \pi_{k,h}(x_h)) + \mathbf{Q}_h^*(x_h, \pi_{k,h}(x_h)) - \mathbf{V}_h^{\pi_k}(x) \right\}] \\ &\stackrel{(ii)}{\geq} \frac{1}{2} \sum_{h=1}^H \mathbb{E}^{\pi_k} [\mathbb{I}(\mathcal{A}_h) \{ \text{gap}(x_h, \pi_{k,h}(x_h)) + \mathbf{Q}_h^*(x_h, \pi_{k,h}(x_h)) - \mathbf{V}_h^{\pi_k}(x) \}] \\ &\stackrel{(iii)}{\geq} \frac{1}{2} (\mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}), \end{aligned}$$

where (i) uses $\epsilon_{\text{clip}} = \frac{\text{gap}_{\min}}{2H}$ and display (9), (ii) uses that $\mathbf{Q}_h^*(x_h, \pi_{k,h}(x_h)) - \mathbf{V}_h^{\pi_k}(x) \geq 0$, and (iii) uses display (10).

Let us start with proving (9). First, consider a stage h , state x , and suppose that $\pi_{k,h}(x) \notin \pi_h^*(x)$. Observe that by Lemma D.4, optimism, and the definition of $\text{gap}_h(x, a)$, we have that for any $a^* \in \pi_h^*(x)$,

$$\begin{aligned} H\epsilon_{\text{clip}} + \ddot{\mathbf{V}}_{k,h}^{\pi_k}(x) &\geq \overline{\mathbf{V}}_{k,h}(x) = \overline{\mathbf{Q}}_{k,h}(x, \pi_{k,h}(x)) \geq \overline{\mathbf{Q}}_{k,h}(x, a^*) \\ &\geq \mathbf{Q}_h^*(x, a^*) = \text{gap}(x, \pi_{k,h}(x)) + \mathbf{Q}_h^*(x, \pi_{k,h}(x)). \end{aligned}$$

Subtracting, we find that for $\pi_{k,h}(x) \notin \pi_h^*(x)$,

$$\ddot{\mathbf{V}}_{k,h}^{\pi_k}(x) - \mathbf{V}_h^{\pi_k}(x) \geq \text{gap}(x, \pi_{k,h}(x)) - H\epsilon_{\text{clip}} + \mathbf{Q}_h^*(x, \pi_{k,h}(x)) - \mathbf{V}_h^{\pi_k}(x). \quad (11)$$

Now, on the other hand, if $\pi_{k,h}(x) \in \pi_h^*(x)$, then,

$$\begin{aligned} \ddot{\mathbf{V}}_{k,h}^{\pi_k}(x) - \mathbf{V}_h^{\pi_k}(x) &= \ddot{\mathbf{E}}_{k,h}(x, \pi_{k,h}(x)) + r(x, \pi_{k,h}(x)) + p(x, \pi_{k,h}(x))^\top \ddot{\mathbf{V}}_{k,h+1}^{\pi_k} \\ &\quad - r(x, \pi_{k,h}(x)) - p(x, \pi_{k,h}(x))^\top \mathbf{V}_h^{\pi_k} \\ &= \ddot{\mathbf{E}}_{k,h}(x, \pi_{k,h}(x)) + p(x, \pi_{k,h}(x))^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}) \\ &\stackrel{(i)}{=} \ddot{\mathbf{E}}_{k,h}(x, \pi_{k,h}(x)) + p(x, \pi_{k,h}(x))^\top \partial \ddot{\mathbf{V}}_{h+1} \\ &\stackrel{(ii)}{\geq} p(x, \pi_{k,h}(x))^\top \partial \ddot{\mathbf{V}}_{h+1}, \end{aligned} \quad (12)$$

where in (i) we have defined the increment $\partial \ddot{\mathbf{V}}_h := \ddot{\mathbf{V}}_{k,h}^{\pi_k} - \mathbf{V}_h^{\pi_k}$ with $\partial \ddot{\mathbf{V}}_{H+1} = 0$, and (ii) holds since $\ddot{\mathbf{E}}_{k,h}(x, \pi_{k,h}(x)) = \mathbf{E}_{k,h}(x, \pi_{k,h}(x)) \mathbb{I}(\mathbf{E}_{k,h}(x, \pi_{k,h}(x)) \geq \epsilon_{\text{clip}}) \geq 0$.

Now, recalling that \mathcal{E}_h denotes the event that $\pi_{k,h}(x) \notin \pi_h^*(x)$, we have

$$\mathbb{E}^{\pi_k} [\partial \ddot{\mathbf{V}}_1] \geq \mathbb{E}^{\pi_k} [\mathbb{I}(\mathcal{E}_1) \{ \text{gap}(x_1, \pi_{k,1}(x_1)) - H\epsilon_{\text{clip}} + \mathbf{Q}_h^*(x_1, \pi_{k,1}(x_1)) - \mathbf{V}_h^{\pi_k}(x) \}] \quad (\text{by Eq. (11)})$$

$$+ \mathbb{E}^{\pi_k} [\mathbb{I}(\mathcal{E}_1^c) p(x_1, \pi_{k,1}(x_1))^\top \partial \ddot{\mathbf{V}}_2]. \quad (\text{by Eq. (13)})$$

We continue with

$$\begin{aligned} & \mathbb{E}^{\pi_k} \left[\mathbb{I}(\mathcal{E}_1^c) p(x_1, \pi_{k,1}(x_1))^\top \partial \ddot{\mathbf{V}}_2 \right] \\ & \geq \mathbb{E}^{\pi_k} \left[\mathbb{I}(\mathcal{E}_1^c) \mathbb{I}(\mathcal{E}_2) \{ \text{gap}(x_2, \pi_{k,2}(x_2)) - H \epsilon_{\text{clip}} + \mathbf{Q}_2^*(x_2, \pi_{k,2}(x_2)) - \mathbf{V}_2^{\pi_k}(x) \} \right] \\ & \quad + \mathbb{E}^{\pi_k} \left[\mathbb{I}(\mathcal{E}_1^c) \mathbb{I}(\mathcal{E}_2^c) p(x_2, \pi_{k,h}(x_2))^\top \partial \ddot{\mathbf{V}}_3 \right]. \end{aligned}$$

Recalling the event $\mathcal{A}_h = \mathcal{E}_h \cap \bigcap_{h' < h} \mathcal{E}_{h'}^c$, we can continue the above induction to find that,

$$\begin{aligned} \mathbb{E}^{\pi_k} [\partial \ddot{\mathbf{V}}_1] & \geq \sum_{h=1}^H \mathbb{E}^{\pi_k} \left[\mathbb{I}(\mathcal{A}_h) \{ \text{gap}(x_h, \pi_{k,h}(x_h)) - H \epsilon_{\text{clip}} + \mathbf{Q}_h^*(x_h, \pi_{k,h}(x_h)) - \mathbf{V}_h^{\pi_k}(x) \} \right] \\ & \quad + \underbrace{\mathbb{E}^{\pi_k} \left[\mathbb{I} \left(\bigcap_{h=1}^H \mathcal{E}_h^c \right) p(x_h, \pi_{k,h}(x_h))^\top \partial \ddot{\mathbf{V}}_{H+1} \right]}_{=0}, \end{aligned}$$

as needed. Now let's prove (10). We can always write

$$\mathbf{V}_h^*(x) - \mathbf{V}_h^{\pi_k}(x) = \text{gap}_h(x, a) + \mathbf{Q}_h^*(x, \pi_{k,h}(x)) - \mathbf{V}_h^{\pi_k}(x),$$

where $\text{gap}_h(x, a) = 0$ when $\pi_k^*(x) \in \pi_{k,h}(x)$, that is, on \mathcal{E}^c . Hence, the same line of reasoning used to prove Eq. (9) (omitting the subtracted $\epsilon_{\text{clip}}H$), verifies Eq. (10). \square

D.2 Proof of Lemma D.1

Proof. For simplicity, set $a = \pi_{k,h}(x)$, and let $a^* \in \pi_h^*(x)$ be an action which witnesses the α transition-suboptimality condition. We then have

$$\begin{aligned} \bar{\mathbf{V}}_{k,h}(x) & \stackrel{(i)}{=} \bar{\mathbf{Q}}_{k,h}(x, a) \stackrel{(ii)}{\geq} \bar{\mathbf{Q}}_{k,h}(x, a^*) \\ & = \mathbf{Q}_h^*(x, a^*) + (\bar{\mathbf{Q}}_{k,h}(x, a^*) - \mathbf{Q}_h^*(x, a^*)) \\ & \stackrel{(iii)}{=} \text{gap}_h(x, a) + \mathbf{Q}_h^*(x, a) + (\bar{\mathbf{Q}}_{k,h}(x, a^*) - \mathbf{Q}_h^*(x, a^*)), \end{aligned}$$

where (i) is by definition of $\bar{\mathbf{V}}_{k,h}(x)$, (ii) is since $a = \pi_{k,h}(x) = \arg \max_{a'} \bar{\mathbf{Q}}_{k,h}(x, a')$, and (iii) is the definition of $\text{gap}_h(x, a)$. Rearranging, we have

$$\text{gap}_h \leq \bar{\mathbf{V}}_{k,h}(x) - \mathbf{Q}_h^*(x, a) - (\bar{\mathbf{Q}}_{k,h}(x, a^*) - \mathbf{Q}_h^*(x, a^*)) \quad (14)$$

If Alg is not necessarily strongly optimistic then we bound $\bar{\mathbf{Q}}_{k,h}(x, a^*) - \mathbf{Q}_h^*(x, a^*) \geq 0$ and $\mathbf{Q}_h^*(x, a) \geq \mathbf{V}_h^{\pi_k}(x)$, yielding

$$\begin{aligned} \text{gap}_h(x, a) & \leq \bar{\mathbf{V}}_{k,h}(x) - \mathbf{V}_h^{\pi_k}(x) \\ & = \bar{\mathbf{Q}}_{k,h}(x, a) - \mathbf{V}_h^{\pi_k}(x) \\ & = \mathbf{E}_{k,h}(x, a) + r(x, a) + p(x, a)^\top \bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_h^{\pi_k}(x) \\ & = \mathbf{E}_{k,h}(x, a) + p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^{\pi_k}) \end{aligned}$$

which corresponds to the desired bound for $\alpha = 1$.

When Alg is strongly optimistic, we handle (14) more carefully. Specifically, we compute

$$\begin{aligned} \bar{\mathbf{V}}_{k,h}(x) - \mathbf{Q}_h^*(x, a) & = \mathbf{E}_{k,h}(x, a) + r(x, a) + p(x, a)^\top \bar{\mathbf{V}}_{k,h+1} - (r(x, a) + p(x, a)^\top \mathbf{V}_{h+1}^*) \\ & = \mathbf{E}_{k,h}(x, a) + p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^*). \end{aligned}$$

Moreover, recalling that $a^* \in \pi_h^*(x)$, we have

$$\begin{aligned} \bar{\mathbf{Q}}_{k,h}(x, a^*) - \mathbf{Q}_h^*(x, a^*) & = r(x, a^*) + \mathbf{E}_{k,h}(x, a^*) + p(x, a^*)^\top \bar{\mathbf{V}}_{k,h+1} - r(x, a^*) - p(x, a^*)^\top \mathbf{V}_{h+1}^* \\ & = \mathbf{E}_{k,h}(x, a^*) + p(x, a^*)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^*). \end{aligned}$$

where the last inequality uses strong optimism of Alg. Hence,

$$\begin{aligned}
\text{gap}_h(x, a) &\leq \bar{\mathbf{V}}_{k,h}(x) - \mathbf{Q}_h^*(x, a) - (\bar{\mathbf{Q}}_{k,h}(x, a^*) - \mathbf{Q}^*(x, a^*)) \\
&= \mathbf{E}_{k,h}(x, a) + p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^*) - (\mathbf{E}_{k,h}(x, a^*) + p(x, a^*)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^*)) \\
&= \mathbf{E}_{k,h}(x, a) - \mathbf{E}_{k,h}(x, a^*) + (p(x, a) - p(x, a^*))^\top (\bar{\mathbf{V}}_{k,h} - \mathbf{V}_{h+1}^*) \\
&\leq \mathbf{E}_{k,h}(x, a) + (p(x, a) - p(x, a^*))^\top (\bar{\mathbf{V}}_{k,h} - \mathbf{V}_{h+1}^*) \quad (\text{Strong Optimism}) \\
&\leq \mathbf{E}_{k,h}(x, a) + \alpha p(x, a)^\top (\bar{\mathbf{V}}_{k,h} - \mathbf{V}_{h+1}^*),
\end{aligned}$$

where the last line uses the component-wise inequalities $p(x, a) - p(x, a^*) \leq \alpha p(x, a)$ due to the fact that a^* witnesses the α transition-suboptimality, and $\bar{\mathbf{V}}_{k,h} - \mathbf{V}_{h+1}^* \geq 0$ due to optimism. \square

D.3 Proof of Lemma D.3

Proof. For ease, we suppress the dependence of α on (x, a, h) . By our fundamental gap bound (Lemma D.1) and then Lemma D.4, we have that

$$\begin{aligned}
\text{gap}_h(x, a) &\leq \mathbf{E}_{k,h}(x, a) + \alpha \cdot p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^{\pi_k}) \\
&\leq \ddot{\mathbf{E}}_{k,h}(x, a) + \alpha \cdot p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}) + (H - h + 1)\alpha\epsilon_{\text{clip}} \\
&\leq \ddot{\mathbf{E}}_{k,h}(x, a) + \alpha \cdot p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}) + \text{gap}_h(x, a)/2,
\end{aligned}$$

where the inequality bounds $\alpha(H - h + 1)\epsilon_{\text{clip}} \leq \alpha \text{gap}_{\min}/2 \leq \alpha \cdot \text{gap}_h(x, a)/2 \leq \text{gap}_h(x, a)/2$. This yields

$$\frac{1}{2} \text{gap}_h(x, a) \leq \ddot{\mathbf{E}}_{k,h}(x, a) + \alpha \cdot p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}).$$

Now, fix a constant $c \in (0, 1]$ to be chosen later. Either we have that $\ddot{\mathbf{E}}_{k,h}(x, a) \geq \frac{c}{2} \text{gap}_h(x, a)$, or otherwise,

$$\alpha \cdot p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}) \geq (1 - c) \frac{1}{2} \text{gap}_h(x, a) \geq \frac{1 - c}{c} \ddot{\mathbf{E}}_{k,h}(x, a),$$

which can be rearranged into

$$\ddot{\mathbf{E}}_{k,h}(x, a) \leq \frac{c\alpha}{1 - c} p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}).$$

Hence, we have

$$\begin{aligned}
\ddot{\mathbf{E}}_{k,h}(x, a) &\leq \ddot{\mathbf{E}}_{k,h}(x, a) \mathbb{I} \left\{ \ddot{\mathbf{E}}_{k,h}(x, a) \geq \frac{c}{2} \text{gap}_h(x, a) \right\} \\
&\quad + \frac{c\alpha}{1 - c} p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}) \mathbb{I} \left\{ \ddot{\mathbf{E}}_{k,h}(x, a) < \frac{c}{2} \text{gap}_h(x, a) \right\} \\
&\leq \ddot{\mathbf{E}}_{k,h}(x, a) \mathbb{I} \left\{ \ddot{\mathbf{E}}_{k,h}(x, a) \geq \frac{c}{2} \text{gap}_h(x, a) \right\} + \frac{c\alpha}{1 - c} p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}),
\end{aligned}$$

and thus,

$$\begin{aligned}
\ddot{\mathbf{E}}_{k,h}(x, a) + p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}) &\leq \ddot{\mathbf{E}}_{k,h}(x, a) \mathbb{I} \left\{ \ddot{\mathbf{E}}_{k,h}(x, a) \geq \frac{c}{2} \text{gap}_h(x, a) \right\} \\
&\quad + \left(1 + \frac{c\alpha}{1 - c}\right) p(x, a)^\top (\ddot{\mathbf{V}}_{k,h+1}^{\pi_k} - \mathbf{V}_{h+1}^{\pi_k}).
\end{aligned}$$

In particular, choosing $c = \frac{1}{2} \min\{1, (\alpha H)^{-1}\}$, we have $(1 + \frac{c\alpha}{1 - c}) \leq 1 + \frac{1}{H}$, and

$$\frac{1}{2} = \frac{(1 \wedge (\alpha H)^{-1})}{4} = \frac{1}{4(\alpha H \vee 1)},$$

so that $\ddot{\mathbf{E}}_{k,h}(x, a) \mathbb{I} \left\{ \ddot{\mathbf{E}}_{k,h}(x, a) \geq \frac{c}{2} \text{gap}_h(x, a) \right\} = \ddot{\mathbf{E}}_{k,h}(x, a)$. This concludes the proof. \square

Part II

StrongEuler and its surpluses

E The StrongEuler Algorithm

Before continuing, let us define a logarithmic factor we shall use throughout:

$$\mathbf{L}(u) := \sqrt{2 \log(10M^2 \max\{u, 1\}/\delta)}, \quad (15)$$

where we recall that $M = SAH \geq 2$. This section formally presents StrongEuler, which makes two subtle modification of the EULER algorithm of [16].

First, similar to [5, 6], StrongEuler refines the log factors in the bonuses to depend on the number of samples $n_k(x, a)$ via $\mathbf{L}(n_k(x, a)) \propto \log(Mn_k(x, a)/\delta)$, rather than the overall time $T = KH$ via $\mathbf{L}(n_k(x, a)) \propto \log(MT/\delta)$, which is necessary to ensure the optimal $\log T$ regret. Following [5, 6], our confidence bounds can be slightly refined using law-of-iterated logarithm bounds, but for simplicity we do not pursue this direction here.

Second, StrongEuler satisfies *strong optimism*. We remind the reader that strong optimism is not necessary to achieve gap dependent bounds, but can achieve sharper bounds for settings with simple transition dynamics like contextual bandits. The EULER algorithm, or its predecessors (e.g. [1]), would also achieve-gap dependent bounds due to our analysis. Moreover, running these algorithms with the refined $\log(Mn_k(x, a)/\delta)$ log factors would also yield $\log T$ - asymptotic regret, whereas implementing $\log(MT/\delta)$ confidence intervals may yield asymptotic regret that scales as $\log^2 T$ (see Remark B.5).

The EULER algorithm proceeds by standard optimistic value iteration, with carefully chosen exploration bonuses, and keeps track of various variance-related quantities:

Algorithm 1: StrongEuler

```

1 Input:
2 Initialized: For each  $a \in \mathcal{A}, x' \in \mathcal{S}, n_1(x, a) = 0, n_1(x' | x, a) = 0, \text{rsum}_1 = 0, \text{rsumsq}_1 = 0,$ 
    $\widehat{p}_1(x, a) = 0, \widehat{\text{Var}}_1[R(x, a)] = 0$ 
3 for  $k = 1, 2, \dots$  do
4    $\overline{\mathbf{V}}_{k, H+1} \leftarrow 0$ 
5   for  $h = H, H - 1, \dots, 1$  do
6     for  $x \in \mathcal{S}$  do
7       for  $a \in \mathcal{A}$  do
8         Call ConstructBonuses.
9          $\overline{\mathbf{Q}}_{k, h}(x, a) \leftarrow \min\{H - h + 1, \widehat{r}(x, a) + \widehat{p}_{k, h}(x, a)^\top \overline{\mathbf{V}}_{k, h+1} +$ 
10           $\mathbf{b}_{k, h}^{\text{prob}}(x, a) + \mathbf{b}_k^{\text{rw}}(x, a) + \mathbf{b}_{k, h}^{\text{str}}(x, a)\}$ 
11       end
12        $\pi_{k, h}(x) := \arg \max_a \overline{\mathbf{Q}}_{k, h}(x, a), \widehat{a} \leftarrow \pi_{k, h}(x)$ 
13        $\overline{\mathbf{V}}_{k, h}(x) := \overline{\mathbf{Q}}_{k, h}(x, \widehat{a})$ 
14        $\underline{\mathbf{V}}_{k, h}(x) =$ 
15          $\max\{0, \widehat{r}(x, \widehat{a}) - \mathbf{b}_{k, h}^{\text{rw}}(x, \widehat{a}) + \widehat{p}_{k, h}(x, \widehat{a})^\top \underline{\mathbf{V}}_{k, h+1} - \mathbf{b}_{k, h}^{\text{prob}}(x, \widehat{a}) - \mathbf{b}_{k, h}^{\text{str}}(x, \widehat{a})\}.$ 
16     end
17   end
18 Call RolloutAndUpdate( $k$ ).
19 end

```

The RolloutAndUpdate function (Algorithm 2 below) executes one trajectory according to the policy π_k , and records all count- and variance- data regarding the relevant rewards and transition probabilities. Finally, the bonuses are defined in Algorithm 3.

Algorithm 2: RolloutAndUpdate(k)

1 **Input:** Global current episode k , global counts and empirical probabilities. Initialize $k + 1$ -th episode counts: $n_{k+1}(\cdot, \cdot) \leftarrow n_k(\cdot, \cdot)$, $n_{k+1}(\cdot | \cdot, \cdot) \leftarrow n_k(\cdot | \cdot, \cdot)$, $\text{rsum}_{k+1}(\cdot, \cdot) \leftarrow \text{rsum}_k(\cdot, \cdot)$, $\text{rsumsq}_{k+1}(\cdot, \cdot) \leftarrow \text{rsumsq}_k(\cdot, \cdot)$.
2 **for** $h = 1, \dots, H$ **do**
3 Observe state x_h , play $a_h = \pi_{k,h}(x_h)$, receive reward R and view next state x_{h+1} .
4 $n_k(x_h, a_h) += 1$, $n_k(x_{h'} | x_h, a_h) += 1$, $\text{rsum}(x, a) += R$, $\text{rsum}(x, a) += R^2$
5 **end**
6 **for** $a \in \mathcal{A}$, $x \in \mathcal{S}$ **do**
7 **for** $x' \in \mathcal{S}$ **do**
8 $\hat{p}_{k+1}(x' | x, a) = \frac{n_k(x_{h'} | x_h, a_h)}{n_k(x_h, a_h)}$
9 **end**
10 $\bar{r}_{k+1}(x, a) = \frac{\text{rsum}_{k+1}}{n_k(x_h, a_h)}$, $\widehat{\text{Var}}_{k+1}[R(x, a)] = \frac{\text{rsumsq}_{k+1}}{n_k(x_h, a_h)} - \bar{r}_{k+1}(x, a)^2$.
11 **end**
12 ,

Algorithm 3: ConstructBonuses

1 **Bonuses:**

$$\mathbf{b}_k^{\text{rw}}(x, a) := 1 \wedge \left(\sqrt{\frac{2\widehat{\text{Var}}_k[R(x, a)]\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{8\mathbf{L}(n_k(x, a))}{3(n_k(x, a) - 1)} \right) \quad (16)$$

$$\begin{aligned} \mathbf{b}_{k,h}^{\text{prob}}(x, a) := & H \wedge \left(\sqrt{\frac{2\text{Var}_{\hat{p}_k(x,a)}[\bar{\mathbf{V}}_{k,h+1}]\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{8H\mathbf{L}(n_k(x, a))}{3(n_k(x, a) - 1)} \right. \\ & \left. + \sqrt{\frac{2\mathbf{L}(n_k(x, a))\|\mathbf{V}_{k,h+1} - \bar{\mathbf{V}}_{k,h+1}\|_{2,\hat{p}_k(x,a)}^2}{n_k(x, a)}} \right). \end{aligned} \quad (17)$$

$$\mathbf{b}_{k,h}^{\text{str}}(x, a) := \|\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}\|_{2,\hat{p}_k(x,a)} \sqrt{\frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{8}{3} \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)} \quad (18)$$

F Analysis of StrongEuler: Proof of Proposition B.4

Proposition B.4 requires demonstrating a lower bound on the surplus, $0 \leq \mathbf{E}_{k,h}(x, a)$, thereby establishing strong optimism, as well as an upper bound on the surplus, which we shall use to analyze the same complexity. We address strong optimism first in the next subsection, and then the upper bound in the following subsection. Throughout, we will assume that a good event $\mathcal{A}^{\text{conc}}$ holds. To keep the proofs modular, the event $\mathcal{A}^{\text{conc}}$ will only appear as an assumption in the supporting lemmas used in Sections F.1 and F.2. Then, in Section F.3, we formally define $\mathcal{A}^{\text{conc}}$ in terms of 6 constituent events, establish $\mathbb{P}[\mathcal{A}^{\text{conc}}] \geq 1 - \frac{\delta}{2}$, and conclude with proofs of the supporting lemmas which rely on $\mathcal{A}^{\text{conc}}$. We remark that many of the arguments in this section are similar to those from [16], with the main differences being strong optimism and the additional care paid to log-factors, necessary for log T regret. Again, recall the definition $\mathbf{L}(u) := \sqrt{2 \log(10M^2 \max\{u, 1\}/\delta)}$.

F.1 Proof of Optimism

Here we establish the optimism of StrongEuler, and in particular, the bound $\mathbf{E}_{k,h}(x, a) \geq 0$.

Proposition F.1. *Under the good event $\mathcal{A}^{\text{conc}}$,*

- (a) StrongEuler is optimistic: $\pi_{k,h}(x) = \arg \max_a \bar{\mathbf{Q}}_{k,h}(x, a)$, where $\bar{\mathbf{Q}}_{k,h}(x, a) \geq \mathbf{Q}_h^*(x, a)$ for all h, x, a . In particular, $\bar{\mathbf{V}}_{k,h}(x) \geq \mathbf{V}_h^*(x)$ for $h \in [0 : H]$.

(b) StrongEuler is strongly optimistic $\mathbf{E}_{k,h}(x,a) := \overline{\mathbf{Q}}_{k,h}(x,a) - r(x,a) - p(x,a)^\top \overline{\mathbf{V}}_{k,h+1}(x) \geq 0$.

(c) $\underline{\mathbf{V}}_{k,h} \leq \mathbf{V}_h^{\pi_k} \leq \mathbf{V}_h^* \leq \overline{\mathbf{V}}_{k,h}$

Proof. The policy choice $\pi_{k,h}(x) = \arg \max_a \overline{\mathbf{Q}}_{k,h}(x,a)$ holds by definition of the algorithm. We now give the remainder of the argument by inducting backwards on h . For $h = H + 1$, $\overline{\mathbf{V}}_{k,H+1} = \underline{\mathbf{V}}_{k,H+1} = \mathbf{V}_{k,H+1}^* = \mathbf{V}_{h+1}^{\pi_k} = 0$. Now, suppose as an inductive hypothesis that $\overline{\mathbf{V}}_{k,h+1} \geq \mathbf{V}_{h+1}^* \geq \mathbf{V}_{h+1}^{\pi_k} \geq \underline{\mathbf{V}}_{k,h+1}$, and $\mathbf{E}_{k,h+1}(x,a) \geq 0$ for all x,a .

First, we shall show that $\mathbf{E}_{k,h}(x,a) \geq 0$ for all x,a . This will establish the induction for point b. It also establishes (a), since then $\overline{\mathbf{Q}}_{k,h}(x,a) \geq r(x,a) + p(x,a)^\top \overline{\mathbf{V}}_{k,h+1}(x) \geq r(x,a) + p(x,a)^\top \mathbf{V}_{h+1}^* = \mathbf{Q}_h^*(x,a)$, proving optimism. To this end, note that

$$\begin{aligned} \mathbf{E}_{k,h}(x,a) &:= \overline{\mathbf{Q}}_{k,h}(x,a) - r(x,a) - p(x,a)^\top \overline{\mathbf{V}}_{k,h+1}(x) \\ &:= \min\{H - h + 1, \hat{r}(x,a) + \hat{p}_{k,h}(x,a)^\top \overline{\mathbf{V}}_{k,h+1} + \mathbf{b}_{k,h}^{\text{prob}}(x,a) + \mathbf{b}_k^{\text{rw}}(x,a) + \mathbf{b}_{k,h}^{\text{str}}(x,a)\} \\ &\quad - r(x,a) - p(x,a)^\top \overline{\mathbf{V}}_{k,h+1}(x). \end{aligned}$$

Since $r(x,a) + p(x,a)^\top \overline{\mathbf{V}}_{k,h+1}(x) \leq H - h + 1$, it suffices to show that

$$\hat{r}(x,a) + \hat{p}_{k,h}(x,a)^\top \overline{\mathbf{V}}_{k,h+1} + \mathbf{b}_{k,h}^{\text{prob}}(x,a) + \mathbf{b}_k^{\text{rw}}(x,a) + \mathbf{b}_{k,h}^{\text{str}}(x,a) - r(x,a) - p(x,a)^\top \overline{\mathbf{V}}_{k,h+1}(x) \geq 0.$$

Grouping the terms, it suffices to show that $\hat{r}(x,a) - r(x,a) + \mathbf{b}_k^{\text{rw}}(x,a) \geq 0$, and that

$$\begin{aligned} 0 &\leq (\hat{p}_{k,h}(x,a)^\top - p(x,a)^\top) \overline{\mathbf{V}}_{k,h+1}(x) + \mathbf{b}_{k,h}^{\text{prob}}(x,a) + \mathbf{b}_{k,h}^{\text{str}}(x,a) \\ &= \left\{ (\hat{p}_{k,h}(x,a)^\top - p(x,a)^\top) \mathbf{V}_{h+1}^* + \mathbf{b}_{k,h}^{\text{prob}}(x,a) \right\} \\ &\quad + \left\{ \hat{p}_{k,h}(x,a)^\top - p(x,a)^\top \right\} (\overline{\mathbf{V}}_{k,h+1}(x) - \mathbf{V}_{h+1}^*(x)) + \mathbf{b}_{k,h}^{\text{str}}(x,a). \end{aligned}$$

We lower bound $\hat{r}(x,a) - r(x,a) + \mathbf{b}_k^{\text{rw}}(x,a)$ and $(\hat{p}_{k,h}(x,a)^\top - p(x,a)^\top) \mathbf{V}_{h+1}^*(x) + \mathbf{b}_{k,h}^{\text{prob}}(x,a)$ by zero with the following lemma:

Lemma F.2. *On the good concentration event $\mathcal{A}^{\text{conc}}$, it holds that*

$$\begin{aligned} |\hat{r}(x,a) - r(x,a)| &\leq \mathbf{b}_k^{\text{rw}}(x,a), \\ |(\hat{p}_{k,h}(x,a)^\top - p(x,a)^\top) \mathbf{V}_{h+1}^*| &\leq \mathbf{b}_{k,h}^{\text{prob}}(x,a) \quad \text{if } \underline{\mathbf{V}}_{k,h+1} \leq \mathbf{V}_{h+1}^* \leq \overline{\mathbf{V}}_{k,h+1} \end{aligned}$$

We conclude the proof of (b) with the following lemma, which lets us bound

$$(\hat{p}_{k,h}(x,a)^\top - p(x,a)^\top) (\overline{\mathbf{V}}_{k,h+1}(x) - \mathbf{V}_{h+1}^*(x)) + \mathbf{b}_{k,h}^{\text{str}}(x,a) \geq 0$$

Precisely we apply the following lemma with $V_2 = \overline{\mathbf{V}}_{k,h+1}$ and $V_1 = \mathbf{V}_{h+1}^*$:

Lemma F.3. *Suppose that $\mathcal{A}^{\text{prob}} \supset \mathcal{A}^{\text{conc}}$ holds, and suppose that $V_1, V_2 : \mathcal{S} \rightarrow \mathbb{R}$ satisfies $\underline{\mathbf{V}}_{k,h+1} \leq V_1 \leq V_2 \leq \overline{\mathbf{V}}_{k,h+1}$. Then,*

$$|(\hat{p}_{k,h}(x,a)^\top - p(x,a)^\top) (V_1 - V_2)| \leq \mathbf{b}_{k,h}^{\text{str}}(x,a)$$

This finally establishes (b). We conclude by establishing (c). Here, we note that by definition $\mathbf{V}_h^{\pi_k} \leq \mathbf{V}_h^*$, and $\mathbf{V}_h^* \leq \overline{\mathbf{V}}_{k,h}$ as show above. Hence, it suffices to show $\underline{\mathbf{V}}_{k,h} \leq \mathbf{V}_h^{\pi_k}$. We begin with the inequality

$$\begin{aligned} \mathbf{V}_h^{\pi_k}(x) &= p(x, a^*)^\top \mathbf{V}_{h+1}^{\pi_k} + r(x, a^*) \\ &= \hat{p}(x, a^*)^\top \mathbf{V}_{h+1}^{\pi_k} + \hat{r}(x, a^*) + (r(x, a^*) - \hat{r}(x, a^*)) + (p(x, a^*)^\top - \hat{p}(x, a^*)^\top) \mathbf{V}_{h+1}^{\pi_k} \\ &= \hat{p}(x, a^*)^\top \mathbf{V}_{h+1}^{\pi_k} + \hat{r}(x, a^*) + (r(x, a^*) - \hat{r}(x, a^*)) + (p(x, a^*)^\top - \hat{p}(x, a^*)^\top) \mathbf{V}_{h+1}^* \\ &\quad + (p(x, a^*) - \hat{p}(x, a^*))^\top (\mathbf{V}_{h+1}^{\pi_k} - \mathbf{V}_{h+1}^*) \\ &\geq \hat{p}(x, a^*)^\top \mathbf{V}_{h+1}^{\pi_k} + \hat{r}(x, a^*) - \mathbf{b}_{k,h}^{\text{rw}}(x,a) - \mathbf{b}_{k,h}^{\text{prob}}(x,a) - \mathbf{b}_{k,h}^{\text{str}}(x,a), \end{aligned}$$

where the last inequality uses the bounds $(r(x, a^*) - \widehat{r}(x, a^*)) \geq -\mathbf{b}_{k,h}^{\text{rw}}(x, a)$ and $(p(x, a^*)^\top - \widehat{p}(x, a^*)^\top) \mathbf{V}_{h+1}^* \geq -\mathbf{b}_{k,h}^{\text{prob}}(x, a)$ on $\mathcal{A}^{\text{conc}}$ due to Lemma F.2, and bounds $(p(x, a^*) - \widehat{p}(x, a^*))^\top (\mathbf{V}_{h+1}^{\pi_k} - \mathbf{V}_{h+1}^*) \geq -\mathbf{b}_{k,h}^{\text{str}}(x, a)$ by applying Lemma F.3 with $V_1 = \mathbf{V}_{h+1}^{\pi_k}$ and $V_2 = \mathbf{V}_{h+1}^*$, which satisfy $\underline{\mathbf{V}}_{k,h+1} \leq V_1 \leq V_2 \leq \overline{\mathbf{V}}_{k,h+1}$ by our inductive hypothesis (namely, $\overline{\mathbf{V}}_{k,h+1} \geq \mathbf{V}_{h+1}^* \geq \mathbf{V}_{h+1}^{\pi_k} \geq \underline{\mathbf{V}}_{k,h+1}$). Since $\mathbf{V}_h^{\pi_k}(x) \geq 0$ as well, and since $\mathbf{V}_{h+1}^{\pi_k} \geq \underline{\mathbf{V}}_{k,h+1}$ by our inductive hypothesis, we therefore have

$$\mathbf{V}_h^{\pi_k}(x) \geq 0 \vee \widehat{p}(x, a^*)^\top \underline{\mathbf{V}}_{k,h+1} + \widehat{r}(x, a^*) - \mathbf{b}_{k,h}^{\text{rw}}(x, a) - \mathbf{b}_{k,h}^{\text{prob}}(x, a) - \mathbf{b}_{k,h}^{\text{str}}(x, a) = \underline{\mathbf{V}}_{k,h}(x),$$

This completes the induction. \square

F.2 Proof of Surplus Bound Upper Bound

Throughout, we assume the round k is fixed, and suppress the dependence of \widehat{p} , $\widehat{\text{Var}}$, and \widehat{r} on k . We use the shorthand $p = p(x, a)$ and $\widehat{p} = \widehat{p}(x, a)$, where the pair (x, a) are clear from context.

$$\begin{aligned} \mathbf{E}_{k,h}(x, a) &= \overline{\mathbf{V}}_{k,h}(x, a) - r(x, a) - p(x, a)^\top \overline{\mathbf{V}}_{k,h+1} \\ &\leq \mathbf{b}_k^{\text{rw}}(x, a) + \widehat{r}(x, a) + \widehat{p}(x, a)^\top \overline{\mathbf{V}}_{k,h+1} + \mathbf{b}_{k,h}^{\text{prob}}(x, a) - r(x, a) - p(x, a)^\top \overline{\mathbf{V}}_{k,h+1} \\ &= (\widehat{r}(x, a) - r(x, a) + (\widehat{p}(x, a) - p(x, a))^\top \mathbf{V}_{h+1}^*) \\ &\quad + \mathbf{b}_k^{\text{rw}}(x, a) + \mathbf{b}_{k,h}^{\text{prob}}(x, a) + (\widehat{p} - p)^\top (\overline{\mathbf{V}}_{k,t+1} - \mathbf{V}_{h+1}^*) \\ &\leq 2\mathbf{b}_k^{\text{rw}}(x, a) + 2\mathbf{b}_{k,h}^{\text{prob}}(x, a) + \mathbf{b}_{k,h}^{\text{str}}(x, a). \end{aligned}$$

where the last line is by Lemmas F.2 and F.3. Next, we state a standard lemma that lets us swap out the empirical variance for the true variance in upper bounding $\mathbf{b}_k^{\text{rw}}(x, a)$:

Lemma F.4. *Under the event $\mathcal{A}^{\text{conc}}$, $\mathbf{b}_k^{\text{rw}}(x, a) \lesssim \sqrt{\frac{\text{Var}[R(x, a)] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)}$.*

Next, we recall from the definition of \mathbf{b}^{prob} ,

$$\mathbf{b}_{k,h}^{\text{prob}}(x, a) \lesssim \sqrt{\frac{\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^*] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{H \mathbf{L}(n_k(x, a))}{n_k(x, a)} + \sqrt{\frac{\mathbf{L}(n_k(x, a)) \|\underline{\mathbf{V}}_{k,h+1} - \overline{\mathbf{V}}_{k,h+1}\|_{2, \widehat{p}}^2}{n_k(x, a)}}.$$

where we replaced $n_k(x, a) - 1$ by $n_k(x, a)$ in the denominator of one of the terms by taking advantage of the ‘ $H \wedge$ ’. Furthermore, we can bound

$$\begin{aligned} \sqrt{\frac{\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^*] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} &\leq \sqrt{\frac{\min\{\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^*], \text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^{\pi_k}]\} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} \\ &\quad + \left| \sqrt{\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^*]} - \sqrt{\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^{\pi_k}]} \right| \sqrt{\frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)}}. \end{aligned}$$

We can control the difference $|\sqrt{\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^*]} - \sqrt{\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^{\pi_k}]}|$ using the following lemma:

Lemma F.5. *Let X, Y be two real valued random variables, and let $\|\cdot\|_{p,2} := \sqrt{\mathbb{E}[(\cdot)^2]}$. Then $|\sqrt{\text{Var}[X]} - \sqrt{\text{Var}[Y]}| \leq \sqrt{\text{Var}[X - Y]} \leq \|X - Y\|_{2,p}$.*

Proof. The inequality $\text{Var}[X - Y] \leq \mathbb{E}[(X - Y)^2] = \|X - Y\|_{2,p}^2$ follows since $\text{Var}[Z] \leq \mathbb{E}[Z^2]$ for any random variable Z . For the first inequality, we can assume WLOG that X, Y are mean zero, in which case $\sqrt{\text{Var}[X]} = \|X\|_{2,p}$, and similarly for Y and $X - Y$. The result now follows from the fact that the norm $\|\cdot\|_{p,2}$ satisfies the triangle inequality. \square

We shall also need the following simple fact:

Fact F.6. *If $V_1(x) \leq V_2(x) \leq V_3(x) \leq V_4(x)$ for all $x \in \mathcal{S}$, then $\|V_2 - V_3\|_{2,p} \leq \|V_1 - V_4\|_{2,p}$.*

Since $\underline{\mathbf{V}}_{k,h+1} \leq \mathbf{V}_{h+1}^{\pi_k} \leq \mathbf{V}_{h+1}^* \leq \overline{\mathbf{V}}_{k,h+1}$ by Proposition F.1, Lemma F.5 and Fact F.6 above yield

$$|\sqrt{\text{Var}_{p(x,a)}[\mathbf{V}_{h+1}^*]} - \sqrt{\text{Var}_{p(x,a)}[\mathbf{V}_{h+1}^{\pi_k}]}| \leq \|\mathbf{V}_{h+1}^* - \mathbf{V}_{h+1}^{\pi_k}\|_{2,p} \leq \|\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1}\|_{2,p}.$$

Together with the elementary inequality, $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b} \lesssim \sqrt{a+b}$, this in turn yields

$$\begin{aligned} & \mathbf{b}_{k,h}^{\text{prob}}(x, a) + \mathbf{b}_{k,h}^{\text{rw}}(x, a) \\ & \lesssim \mathbf{b}_{k,h}^{\text{rw}}(x, a) + \sqrt{\frac{\min\{\text{Var}_{p(x,a)}[\mathbf{V}_{h+1}^*], \text{Var}_{p(x,a)}[\mathbf{V}_{h+1}^{\pi_k}]\} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{H\mathbf{L}(n_k(x, a))}{n_k(x, a)} \\ & \quad + \sqrt{\frac{\mathbf{L}(n_k(x, a)) \|\underline{\mathbf{V}}_{k,h+1} - \overline{\mathbf{V}}_{k,h+1}\|_{2,\widehat{p}+p}^2}{n_k(x, a)}} \\ & \lesssim \sqrt{\frac{\text{Var}[R(x, a)] + \min\{\text{Var}_{p(x,a)}[\mathbf{V}_{h+1}^*], \text{Var}_{p(x,a)}[\mathbf{V}_{h+1}^{\pi_k}]\} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{H\mathbf{L}(n_k(x, a))}{n_k(x, a)} \\ & \quad + \sqrt{\frac{\mathbf{L}(n_k(x, a)) \|\underline{\mathbf{V}}_{k,h+1} - \overline{\mathbf{V}}_{k,h+1}\|_{2,\widehat{p}+p}^2}{n_k(x, a)}}, \\ & \lesssim \sqrt{\frac{\text{Var}_{h,x,a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{H\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \sqrt{\frac{\mathbf{L}(n_k(x, a)) \|\underline{\mathbf{V}}_{k,h+1} - \overline{\mathbf{V}}_{k,h+1}\|_{2,\widehat{p}+p}^2}{n_k(x, a)}}, \end{aligned}$$

where we use the shorthand $\|V\|_{2,\widehat{p}+p} = \sqrt{\|V\|_{2,p}^2 + \|V\|_{2,\widehat{p}}^2}$, and where in the last, we recall that $\text{Var}_{h,x,a}^{(k)} = \min\{\text{Var}_{h,x,a}^{\pi_k}, \text{Var}_{h,x,a}^*\} = \text{Var}[R(x, a)] + \min\{\text{Var}_{p(x,a)}[\mathbf{V}_{h+1}^*], \text{Var}_{p(x,a)}[\mathbf{V}_{h+1}^{\pi_k}]\}$.

Next, substituting in $\mathbf{b}_{k,h}^{\text{str}}(x, a) := \|\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1}\|_{2,\widehat{p}(x,a)} \sqrt{\frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{8}{3} \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)}$, we obtain

$$\begin{aligned} & \mathbf{b}_k^{\text{rw}}(x, a) + \mathbf{b}_{k,h}^{\text{prob}}(x, a) + \mathbf{b}_{k,h}^{\text{str}}(x, a) \\ & \lesssim \sqrt{\frac{\text{Var}_{h,x,a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{H\mathbf{L}(n_k(x, a))}{n_k(x, a)} \\ & \quad + \sqrt{\frac{\mathbf{L}(n_k(x, a)) \|\underline{\mathbf{V}}_{k,h+1} - \overline{\mathbf{V}}_{k,h+1}\|_{2,\widehat{p}+p}^2}{n_k(x, a)}} + \sqrt{\frac{S\|\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1}\|_{2,\widehat{p}}^2 \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)} \\ & \lesssim \sqrt{\frac{\text{Var}_{h,x,a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \sqrt{\frac{S\|\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1}\|_{2,p+\widehat{p}}^2 \mathbf{L}(n_k(x, a))}{n_k(x, a)}} \end{aligned} \tag{19}$$

$$\begin{aligned} & \stackrel{(i)}{\leq} \sqrt{\frac{\text{Var}_{h,x,a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \|\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1}\|_{2,\widehat{p}+p}^2 \\ & \stackrel{(ii)}{=} \sqrt{\frac{\text{Var}_{h,x,a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)} + 2\|\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1}\|_{2,p}^2 + (\widehat{p} - p)(\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1})^2 \\ & \stackrel{(iii)}{\leq} \sqrt{\frac{\text{Var}_{h,x,a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)} + 2\|\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1}\|_{2,p}^2 + H(p - \widehat{p})^\top (\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1}), \end{aligned}$$

where (i) uses the inequality $a/b \leq a^2 + \frac{1}{b^2}$, and (ii) uses the facts that $\|V\|_{2,\widehat{p}+p}^2 = \|V\|_{2,p}^2 + \|V\|_{2,\widehat{p}}^2$ and $\|V\|_{2,\widehat{p}}^2 = \langle \widehat{p}, V^2 \rangle = \langle \widehat{p}, V^2 \rangle + \langle \widehat{p} - p, V^2 \rangle = \|V\|_{2,\widehat{p}}^2 + \langle \widehat{p} - p, V^2 \rangle$. Lastly, inequality (iii) uses $0 \leq \underline{\mathbf{V}}_{k,h+1} \leq \overline{\mathbf{V}}_{k,h+1} \leq H$.

We continue bounding $H(p - \widehat{p})^\top (\overline{\mathbf{V}}_{k,h+1} - \underline{\mathbf{V}}_{k,h+1})$ in much the same way that we bounded the term in Lemma F.3 in terms of \mathbf{b}^{str} , with the exception that we seek a term which depends on the true transition probability $p(x, a)$, and not the empirical $\widehat{p}(x, a)$:

Lemma F.7. Under $\mathcal{A}^{\text{conc}}$,

$$|(\hat{p}(x, a) - p(x, a))^\top (\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1})| \lesssim \|\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, p(x, a)} \sqrt{\frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)}.$$

The proof of the above lemma is omitted for the sake of brevity, and follows from a simplified version of the proof of Lemma F.3 where we need not pass through an empirical variance. Applying the bound in Lemma F.7, we have

$$\begin{aligned} H(p - \hat{p})^\top (\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}) &\lesssim H \|\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, p} \sqrt{\frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH^2\mathbf{L}(n_k(x, a))}{n_k(x, a)} \\ &\lesssim \|\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, p}^2 + \frac{H^2 S\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \frac{SH^2\mathbf{L}(n_k(x, a))}{n_k(x, a)}, \end{aligned}$$

where the last line uses the inequality $ab \leq (a^2 + b^2)/2$. Finally, combining the above with our previous bound, we arrive at

$$\begin{aligned} \mathbf{E}_{k, h}(x, a) &\lesssim \mathbf{b}_k^{\text{rw}}(x, a) + \mathbf{b}_{k, h}^{\text{prob}}(x, a) + \mathbf{b}_{k, h}^{\text{str}}(x, a) \\ &\lesssim \sqrt{\frac{\text{Var}_{h, x, a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH^2\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \|\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, p(x, a)}^2. \end{aligned}$$

From first principles, it is straightforward to show that $\mathbf{E}_{k, h}(x, a) \lesssim H$, which implies that

$$\mathbf{E}_{k, h}(x, a) \lesssim H \wedge \sqrt{\frac{\text{Var}_{h, x, a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}} \quad (20)$$

$$+ \left(H \wedge \frac{SH^2\mathbf{L}(n_k(x, a))}{n_k(x, a)} \right) + \|\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, p(x, a)}^2. \quad (21)$$

To conclude the proof, it remains to unravel the term $\|\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, p(x, a)}^2$.

Lemma F.8. Define the term

$$\mathbf{Z}_k(x, a) = H^2 \wedge H^2 \left(\sqrt{\frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)} \right)^2,$$

Then, we have the bound

$$\bar{\mathbf{V}}_{k, h}(x) - \underline{\mathbf{V}}_{k, h}(x) \lesssim \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \sqrt{\mathbf{Z}_k(x_t, a_t)} \mid x_h = x \right]. \quad (22)$$

As a consequence, we can compute

$$(\bar{\mathbf{V}}_{k, h}(x) - \underline{\mathbf{V}}_{k, h}(x))^2 \lesssim \mathbb{E}^{\pi_k} \left[\left(\sum_{t=h}^H \sqrt{\mathbf{Z}_k(x_t, a_t)} \right)^2 \mid x_h = x \right] \leq H \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \mathbf{Z}_k(x_t, a_t) \mid x_h = x \right].$$

Hence, we have

$$\begin{aligned} \|\bar{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, p}^2 &= \mathbb{E}_{x' \sim p(x, a)} (\bar{\mathbf{V}}_{k, h+1}(x) - \underline{\mathbf{V}}_{k, h+1}(x))^2 \\ &\lesssim H \mathbb{E}^{\pi_k} \left[\sum_{t=h+1}^H \mathbf{Z}_k(x_t, a_t) \mid (x_h, a_h) = (x, a) \right]. \end{aligned}$$

Since $H\mathbf{Z}_k(x, a) \geq H \wedge \frac{SH^2\mathbf{L}(n_k(x, a))}{n_k(x, a)}$. Hence, we can bound via (20)

$$\mathbf{E}_{k, h}(x, a) \lesssim \underbrace{H \wedge \sqrt{\frac{\text{Var}_{h, x, a}^{(k)} \mathbf{L}(n_k(x, a))}{n_k(x, a)}}}_{:= \mathbf{B}_{k, h}^{\text{lead}}(x, a)} + \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H \underbrace{H\mathbf{Z}_k(x_t, a_t)}_{:= \mathbf{B}_k^{\text{fut}}(x_t, a_t)} \mid (x_h, a_h) = (x, a) \right]$$

where we note that the summation in the expectation now begins at $t = h$ to account for the term $H \wedge \frac{SH^2 \mathbf{L}(n_k(x, a))}{n_k(x, a)}$ in (20), and recall that

$$\mathbf{B}_k^{\text{fut}}(x, a) := H^3 \wedge H^3 \left(\sqrt{\frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)} \right)^2 = H\mathbf{Z}_k(x, a)$$

To conclude, we recall the definitions,

$$\mathbf{L}(u) := \sqrt{2 \log(10M^2 \max\{u, 1\}/\delta)}.$$

so that, for $u \geq 1$, $\mathbf{L}(u) \lesssim \log \frac{Mu}{\delta}$.

F.3 Definition of $\mathcal{A}^{\text{conc}}$, and proofs of supporting lemmas

Before proving the lemmas above, we formally express the good event $\mathcal{A}^{\text{conc}}$ as a list of constituent concentration events, and verify that it occurs with probability at least $1 - \delta/2$:

Proposition F.9. *The event $\mathcal{A}^{\text{conc}} := \mathcal{A}^{\text{rw}} \cap \mathcal{A}^{\text{prob}} \cap \mathcal{A}^{\text{val}} \cap \mathcal{A}^{\text{var, val}} \cap \mathcal{A}^{\text{var, rw}}$ occurs with probability $1 - \delta/2$, where each of the constituent events occurs with probability at least $1 - \delta/12$:*

$$\begin{aligned} \mathcal{A}^{\text{rw}} &:= \left\{ \forall k, x, a, h : |\widehat{r}_k(x, a) - r(x, a)| \leq \sqrt{\text{Var}[R(x, a)] \frac{2\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \frac{2\mathbf{L}(n_k(x, a))}{3n_k(x, a)}} \right\} \\ \mathcal{A}^{\text{prob}} &:= \left\{ \forall k, x, x', a, h : |\widehat{p}(x' | x, a) - p(x' | x, a)| \leq \sqrt{p(x' | x, a)(1 - p(x' | x, a)) \frac{2\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \frac{2\mathbf{L}(n_k(x, a))}{3n_k(x, a)}} \right\} \\ \mathcal{A}^{\text{val}} &:= \left\{ \forall k, x, a, h : |(\widehat{p}(x, a) - p(x, a))^\top \mathbf{V}_{h+1}^*| \leq \sqrt{\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^*] \frac{2\mathbf{L}(n_k(x, a))}{n_k(x, a)} + \frac{2H\mathbf{L}(n_k(x, a))}{3n_k(x, a)}} \right\} \\ \mathcal{A}^{\text{var, prob}} &:= \left\{ \forall k, h, x, a : |\widehat{p}(x' | x, a) - p(x' | x, a)| \leq \sqrt{\frac{2\mathbf{L}(n_k(x, a))}{n_k(x, a)}} \right\} \\ \mathcal{A}^{\text{var, val}} &:= \left\{ \forall k, h, x, a : \left| \|\mathbf{V}_h^*\|_{2, \widehat{p}(x, a)} - \|\mathbf{V}_h^*\|_{2, p(x, a)} \right| \leq H \sqrt{\frac{2\mathbf{L}n_k(x, a)}{n_k(x, a) - 1}} \right\} \\ \mathcal{A}^{\text{var, rw}} &:= \left\{ \forall k, h, x, a : \left| \sqrt{\widehat{\text{Var}}(R(x, a))} - \sqrt{\text{Var}(R(x, a))} \right| \leq \sqrt{\frac{2\mathbf{L}(n_k(x, a))}{n_k(x, a) - 1}} \right\} \end{aligned}$$

Proof. The proof of these the first four events follows from standard applications of Bernstein's and Hoeffding's inequality, and the last two from Theorem 10 in [10]. Similar proofs can be found in [16, 1, 5]. As in those works, the only subtlety is to use the appropriate concentration inequality with respect to an appropriate filtration to attain bounds that depend on $\mathbf{L}(n_k(x, a))$, rather than on $\mathbf{L}(T)$.

Let's prove \mathcal{A}^{rw} as an example. We it suffices to only consider rounds for which $n_k(x, a) \geq 1$, for otherwise the bound is vacuous. Fix an action (x, a) , and let $\tau_i \in \{1, 2, \dots\} \cup \{\infty\}$ denote the round $k + 1$ immediately after the i -th round k at which a pair (x, a) is observed at least once during the rollout, and define a sub-filtration $\{\mathcal{G}_i\}$ via $\mathcal{G}_i = \mathcal{F}_{\tau_i}$. Then, for any given i , a martingale analogue of Bernstein's inequality yields

$$\mathbb{P} \left[|\widehat{r}_{\tau_i}(x, a) - r(x, a)| \mathbb{I}(\tau_n < \infty) \geq \sqrt{\frac{2\text{Var}[R(x, a)] \log(2/\eta)}{n_{\tau_i}(x, a)}} + \frac{2 \log(2/\eta)}{3n_{\tau_i}(x, a)} \right] \leq \eta.$$

Now fix an $i \geq 1$. Since $\widehat{r}_k(x, a)$ and $n_k(x, a)$ are constant for $k \in \{\tau_i, \dots, \tau_{i+1} - 1\}$, we have

$$\forall n, \mathbb{P} \left[\exists k \in \{\tau_i, \dots, \tau_{i+1} - 1\} : |\widehat{r}_k(x, a) - r(x, a)| \geq \sqrt{\frac{2\text{Var}[R(x, a)] \log(2/\eta)}{n_k(x, a)}} + \frac{2 \log(2/\eta)}{3n_k(x, a)} \right] \leq \eta,$$

Applying the above with $\eta \leftarrow 2\eta/i^2$ and union bounding over n , we have

$$\mathbb{P} \left[\exists i, k : k \in \{\tau_i, \dots, \tau_{i+1} - 1\}, |\widehat{r}_k(x, a) - r(x, a)| \geq \sqrt{\frac{2\text{Var}[R(x, a)] \log(4i^2/\eta)}{n_k(x, a)}} + \frac{2 \log(4i^2/\eta)}{3n_k(x, a)} \right] \leq \eta.$$

Since $n_k(x, a)$ increments by at least one for each τ_i , we have $i \leq n_k(x, a)$ for $k \in \{\tau_i, \dots, \tau_{i+1} - 1\}$. Thus,

$$\mathbb{P} \left[\exists i, k : k \in \{\tau_i, \dots, \tau_{i+1} - 1\}, |\widehat{r}_k(x, a) - r(x, a)| \geq \sqrt{\frac{2\text{Var}[R(x, a)] \log(4n_k(x, a)^2/\eta)}{n_k(x, a)}} + \frac{2 \log(4n_k(x, a)^2/\eta)}{3n_k(x, a)} \right] \leq \eta.$$

Lastly, since for any k , there always exist some i for which $k \in \{\tau_i, \dots, \tau_{i+1} - 1\}$, we have

$$\mathbb{P} \left[k : |\widehat{r}_k(x, a) - r(x, a)| \geq \sqrt{\frac{2\text{Var}[R(x, a)] \log(4n_k(x, a)^2/\eta)}{n_k(x, a)}} + \frac{2 \log(4n_k(x, a)^2/\eta)}{3n_k(x, a)} \right] \leq \eta.$$

We then conclude by union bounding over SA , and letting $\eta = \delta/12SA$, yielding the following log factor: $\log(48SAN_k(x, a)^2/\delta) \leq \mathbf{L}(n_k(x, a))$, where we recall $\mathbf{L}(u) = \sqrt{2 \log(10M^2 \max\{u, 1\}/\delta)}$ for $M = SAH$. The proof for $\mathcal{A}^{\text{prob}}$ is analogous, the proof for \mathcal{A}^{val} requires union bounding over states x' , incurring a log factor $\log(4S^2An_k(x, a)^2/\delta) \leq \mathbf{L}(n_k(x, a))$. \square

Proof of Lemma F.2. We prove the bound $|(\widehat{p}(x, a) - p(x, a))^\top \mathbf{V}_{h+1}^*| \leq \mathbf{b}_{k,h}^{\text{prob}}(x, a)$; the analogous bounds for rewards is similar. Note that since $\widehat{p}(x, a)^\top \mathbf{V}_{h+1}^* \in [0, H]$ and $p(x, a)^\top \mathbf{V}_{h+1}^* \in [0, H]$, $|(\widehat{p}(x, a) - p(x, a))^\top \mathbf{V}_{h+1}^*| \in [0, H]$. This takes care of the first ' $H \wedge$ ' in $\mathbf{b}_{k,h}^{\text{prob}}(x, a)$. Next, on \mathcal{A}^{val} and $\mathcal{A}^{\text{var, val}}$,

$$\begin{aligned} & |(\widehat{p}(x, a) - p(x, a))^\top \mathbf{V}_{h+1}^*| \\ & \leq \sqrt{\frac{2\text{Var}_{p(x, a)}[\mathbf{V}_{h+1}^*] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{2H\mathbf{L}(n_k(x, a))}{3n_k(x, a)} && \text{(on } \mathcal{A}^{\text{val}}\text{)} \\ & \leq \sqrt{\frac{2\text{Var}_{\widehat{p}(x, a)}[\mathbf{V}_{h+1}^*] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{8H\mathbf{L}(n_k(x, a))}{3(n_k(x, a) - 1)} && \text{(on } \mathcal{A}^{\text{var, val}}\text{)} \\ & = \sqrt{\frac{2\text{Var}_{\widehat{p}(x, a)}[\overline{\mathbf{V}}_{k, h+1}] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{8H\mathbf{L}(n_k(x, a))}{3(n_k(x, a) - 1)} \\ & + (\sqrt{\text{Var}_{\widehat{p}(x, a)}[\mathbf{V}_{h+1}^*]} - \sqrt{\text{Var}_{\widehat{p}(x, a)}[\widehat{p}(x, a)^\top \overline{\mathbf{V}}_{k, h+1}]}) \sqrt{\frac{2\mathbf{L}(n_k(x, a))}{n_k(x, a)}}. \end{aligned}$$

Lastly, by Lemma F.5, we have the bound

$$\left| \sqrt{\text{Var}_{\widehat{p}(x, a)}[\mathbf{V}_{h+1}^*]} - \sqrt{\text{Var}_{\widehat{p}(x, a)}[\overline{\mathbf{V}}_{k, h+1}]} \right| \leq \|\underline{\mathbf{V}}_{k, h+1} - \overline{\mathbf{V}}_{k, h+1}\|_{2, \widehat{p}(x, a)}.$$

\square

Proof of Lemma F.3. Summing up the condition of event $\mathcal{A}^{\text{prob}}$ over states $x' \in \mathcal{S}$, and then applying event $\mathcal{A}^{\text{var,prob}}$ to control $|p(x'|x, a) - \widehat{p}(x'|x, a)|$:

$$\begin{aligned}
(\widehat{p}(x, a) - p(x, a))^\top (\overline{\mathbf{V}}_{k, h+1} - \mathbf{V}_{h+1}^*) &\leq \sum_{x'} \sqrt{2 \frac{\mathbf{L}(n_k(x, a)) p(x'|x, a) (1 - p(x'|x, a))}{n_k(x, a)}} |V_2(x') - V_1(x')| \\
&\quad + \frac{2}{3} \sum_{x'} \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)} |V_2(x') - V_1(x')|, \\
&\stackrel{(i)}{\leq} \sum_{x'} \sqrt{2 \frac{\mathbf{L}(n_k(x, a)) p(x'|x, a)}{n_k(x, a)}} |V_2(x') - V_1(x')| \\
&\quad + \frac{2}{3} \sum_{x'} \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)} |V_2(x') - V_1(x')|, \\
&\stackrel{(ii)}{\leq} \sum_{x'} \sqrt{2 \frac{\mathbf{L}(n_k(x, a)) \widehat{p}(x'|x, a)}{n_k(x, a)}} |V_2(x') - V_1(x')| \\
&\quad + \underbrace{\frac{8}{3} \sum_{x'} \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)} |V_2(x') - V_1(x')|}_{\leq \frac{8}{3} \frac{H S \mathbf{L}(n_k(x, a))}{n_k(x, a)}},
\end{aligned}$$

where (i) uses $p(x'|x, a)(1 - p(x'|x, a)) \leq p(x'|x, a)$, (ii) uses event $\mathcal{A}^{\text{var,prob}}$, and where bound in the bracket is because there $|\overline{\mathbf{V}}_{k, t+1}(x') - \mathbf{V}_{h+1}^*(x')| \leq H$ by Proposition F.1 part (b), and there are at most S terms in the summation. To bound the first term, we have

$$\begin{aligned}
&\sum_{x'} \sqrt{2 \frac{\mathbf{L}(n_k(x, a)) \widehat{p}(x'|x, a)}{n_k(x, a)}} |V_2(x') - V_1(x')| \\
&= \sqrt{2 \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)}} \sum_{x'} \sqrt{\widehat{p}(x'|x, a)} |V_2(x') - V_1(x')| \\
&\stackrel{(i)}{\leq} \sqrt{2 \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)}} \sqrt{S \|V_2 - V_1\|_{2, \widehat{p}}} \\
&\stackrel{(ii)}{\leq} \|\overline{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, \widehat{p}} \sqrt{2 \frac{S \mathbf{L}(n_k(x, a))}{n_k(x, a)}}
\end{aligned}$$

(i) bounds uses Cauchy-Schwartz, and (ii) uses Proposition F.1 part (c) to bound $\|V_2 - V_1\|_{2, p} \leq \|\overline{\mathbf{V}}_{k, h+1} - \underline{\mathbf{V}}_{k, h+1}\|_{2, \widehat{p}}$ for $\underline{\mathbf{V}}_{k, h+1} \leq V_1 \leq V_2 \leq \overline{\mathbf{V}}_{k, h+1}$, in light of Fact F.6. \square

Proof of Lemma F.4. Under the event $\mathcal{A}^{\text{conc}}$ we have

$$\begin{aligned}
&\mathbf{b}_k^{\text{rw}}(x, a) \\
&\lesssim 1 \wedge \left(\sqrt{\frac{\widehat{\text{Var}}[R(x, a)] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a) - 1} \right) \quad (\text{definition}) \\
&\leq 1 \wedge \left(\sqrt{\frac{\widehat{\text{Var}}[R(x, a)] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \left| \sqrt{\widehat{\text{Var}}[R(x, a)]} - \sqrt{\widehat{\text{Var}}[R(x, a)]} \right| \sqrt{\frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a) - 1} \right) \\
&\lesssim 1 \wedge \left(\sqrt{\frac{\widehat{\text{Var}}[R(x, a)] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a) - 1} \right) \lesssim \sqrt{\frac{\widehat{\text{Var}}[R(x, a)] \mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{\mathbf{L}(n_k(x, a))}{n_k(x, a)},
\end{aligned}$$

where in the second-to-last inequality, we used the event $\mathcal{A}^{\text{var, rw}}$ to control $\left| \sqrt{\widehat{\text{Var}}[R(x, a)]} - \sqrt{\widehat{\text{Var}}[R(x, a)]} \right| \lesssim \sqrt{\frac{\mathbf{L}(n_k(x, a))}{n_k(x, a) - 1}}$. \square

Proof of Lemma F.8. Let $a = \pi_{k,h}(x)$. Then, by definition of $\bar{\mathbf{V}}_{k,h}(x), \mathbf{V}_{k,h}(x)$

$$\begin{aligned}
\bar{\mathbf{V}}_{k,h}(x) - \mathbf{V}_{k,h}(x) &= (H - h + 1) \wedge (\hat{p}_{k,h}(x, a)^\top \bar{\mathbf{V}}_{k,h+1} + \mathbf{b}_k^{\text{rw}}(x, a) + \mathbf{b}_{k,h}^{\text{prob}}(x, a) + \mathbf{b}_{k,h}^{\text{str}}(x, a)) \\
&\quad - 0 \vee (\hat{p}_{k,h}(x, a)^\top \mathbf{V}_{k,h+1} - \mathbf{b}_k^{\text{rw}}(x, a) - \mathbf{b}_{k,h}^{\text{prob}}(x, a)) \\
&\leq H \wedge \left\{ \hat{p}_{k,h}(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}) + 2\mathbf{b}_k^{\text{rw}}(x, a) + 2\mathbf{b}_{k,h}^{\text{prob}}(x, a) + \mathbf{b}_{k,h}^{\text{str}}(x, a) \right\} \\
&= H \wedge \left\{ p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}) + (\hat{p}(x, a) - p(x, a))^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}) \right. \\
&\quad \left. + 2\mathbf{b}_k^{\text{rw}}(x, a) + 2\mathbf{b}_{k,h}^{\text{prob}}(x, a) + \mathbf{b}_{k,h}^{\text{str}}(x, a) \right\} \\
&\leq H \wedge \left\{ p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}) + 2\mathbf{b}_k^{\text{rw}}(x, a) + 2\mathbf{b}_{k,h}^{\text{prob}}(x, a) + 2\mathbf{b}_{k,h}^{\text{str}}(x, a) \right\}, \\
&\hspace{15em} \text{(Lemma F.3)} \\
&\leq p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}) + H \wedge \left\{ 2\mathbf{b}_k^{\text{rw}}(x, a) + 2\mathbf{b}_{k,h}^{\text{prob}}(x, a) + 2\mathbf{b}_{k,h}^{\text{str}}(x, a) \right\},
\end{aligned}$$

where the last line uses the fact that $p(x, a)^\top (\bar{\mathbf{V}}_{k,h+1} - \mathbf{V}_{k,h+1}) \geq 0$ on $\mathcal{A}^{\text{conc}}$ (Proposition F.1, part (c)). Unfolding the above expression inductively, we then find that

$$\bar{\mathbf{V}}_{k,h}(x) - \mathbf{V}_{k,h}(x) \leq \mathbb{E}^{\pi_k} \left[\sum_{t=h}^H H \wedge \left\{ 2\mathbf{b}_k^{\text{rw}}(x_t, a_t) + 2\mathbf{b}_{k,h}^{\text{prob}}(x_t, a_t) + 2\mathbf{b}_{k,h}^{\text{str}}(x_t, a_t) \right\} \mid x_t = x \right].$$

To conclude, it suffices to check that $H \wedge \left\{ 2\mathbf{b}_k^{\text{rw}}(x, a) + 2\mathbf{b}_{k,h}^{\text{prob}}(x, a) + 2\mathbf{b}_{k,h}^{\text{str}}(x, a) \right\} \lesssim \sqrt{\mathbf{Z}_k(x, a)}$, for any triple x, a, h . To check that this bound holds, we have from (19) that

$$\begin{aligned}
&2\mathbf{b}_k^{\text{rw}}(x_t, a_t) + 2\mathbf{b}_{k,t}^{\text{prob}}(x_t, a_t) + 2\mathbf{b}_{k,t}^{\text{str}}(x_t, a_t) \\
&\lesssim \sqrt{\frac{\text{Var}_{t,x,a}^* \mathbf{L}(n_k(x_t, a_t))}{n_k(x_t, a_t)}} + \frac{SH\mathbf{L}(n_k(x_t, a_t))}{n_k(x_t, a_t)} + \sqrt{\frac{S \|\bar{\mathbf{V}}_{k,t+1} - \mathbf{V}_{k,t+1}\|_{2,(\hat{p}_k+p)(x_t, a_t)}^2 \mathbf{L}(n_k(x_t, a_t))}{n_k(x_t, a_t)}} \\
&\lesssim \sqrt{\frac{H\mathbf{L}(n_k(x_t, a_t))}{n_k(x_t, a_t)}} + \frac{SH\mathbf{L}(n_k(x_t, a_t))}{n_k(x_t, a_t)} + H \sqrt{\frac{S\mathbf{L}(n_k(x_t, a_t))}{n_k(x_t, a_t)}},
\end{aligned}$$

where we recall the notation $\|V\|_{2, \hat{p}+p} = \sqrt{\|V\|_{2,p}^2 + \|V\|_{2, \hat{p}}^2}$, and thus the final bound holds since $\text{Var}_{t,x,a}^* \leq H$ implying that $\|\bar{\mathbf{V}}_{k,t+1} - \mathbf{V}_{k,t+1}\|_{2,(\hat{p}_k+p)(x_t, a_t)}^2 \leq 4H$ for $0 \leq \mathbf{V}_{k,t+1} \leq \bar{\mathbf{V}}_{k,t+1} \leq H$. Consolidating the terms, we have $2\mathbf{b}_k^{\text{rw}}(x_t, a_t) + 2\mathbf{b}_{k,t}^{\text{prob}}(x_t, a_t) + 2\mathbf{b}_{k,t}^{\text{str}}(x_t, a_t)$ is at most $\lesssim \left(H \sqrt{\frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)} \right)$, and thus $H \wedge 2\mathbf{b}_k^{\text{rw}}(x_t, a_t) + 2\mathbf{b}_{k,t}^{\text{prob}}(x_t, a_t) + 2\mathbf{b}_{k,t}^{\text{str}}(x_t, a_t)$ is $\lesssim H \wedge \left(H \sqrt{\frac{S\mathbf{L}(n_k(x, a))}{n_k(x, a)}} + \frac{SH\mathbf{L}(n_k(x, a))}{n_k(x, a)} \right) := \sqrt{\mathbf{Z}_k(x, a)}$. \square

Part III

Lower Bounds

G Min-Gap Lower Bound for Optimistic Algorithms (Theorem 2.3)

G.1 Formal Statement

We begin a formal version of the lower bound, Theorem 2.3.

Theorem G.1. *Let c_1, c_2, c_3 be absolute constants that may depend on the constants defined in Section G.2. Let Alg denote an algorithm in the class described in Section G.2 run with confidence parameter $\delta \in (0, 1/8)$. For any $S \geq 1$ and $\epsilon \leq 1/\lceil c_1 S \log(S/\delta) \rceil$, fix any MDP in the class described in Section G.3 so that $|\mathcal{S}| = 2S + 1$, $|\mathcal{A}| = 2$, $H = 2$, and exactly one state has a sub-optimality gap of $\text{gap}_{\min} = \epsilon$ and all other states have a minimum sub-optimality gap of at least $1/2$. Then $\sum_{h,x,a:\text{gap}_h(x,a)>0} \frac{1}{\text{gap}_h(x,a)} \lesssim S + \frac{1}{\text{gap}_{\min}}$ but Alg for all sufficiently large K suffers a regret*

$$\text{Regret}_K \geq \frac{c_2 S}{\text{gap}_{\min}} \log(1/\delta) \gtrsim \sum_{h,x,a:\text{gap}_h(x,a)>0} \frac{1}{\text{gap}_h(x,a)} + \frac{S}{\text{gap}_{\min}}$$

with probability at least $1 - c_2 S \epsilon^{-2} \log(1/\delta) e^{-c_3 S} - 3\delta$.

In particular, for any $\epsilon \in (0, c)$ for some constant c , if $\log(\epsilon^{-1}/\delta) \lesssim S \lesssim \epsilon^{-1}/\log(\epsilon^{-1}/\delta)$ then the above regret lower bound holds with probability $1 - O(\delta)$.

G.2 Algorithm Class

Optimistic Q-functions: We consider algorithms where the optimistic Q-function is constructed as follows: given a reward bonus function $\mathbf{b}_k^{\text{rw}}(x, a) \geq 0$ and an additional nonnegative stage-dependent bonus $\mathbf{b}_{k,h}(x, a)$, and empirical estimates $\hat{r}_k(x, a)$ of the reward and $\hat{p}_k(x, a) = (\hat{p}(x'|x, a))$ of the transition probabilities. We set the Q-function at stage H as $\bar{\mathbf{Q}}_{k,H}(x, a) = \hat{r}_k(x, a) + \mathbf{b}_k^{\text{rw}}(x, a)$, where $\hat{r}_k(x, a)$, and for $h \in \{1, \dots, H-1\}$,

$$\begin{aligned} \bar{\mathbf{V}}_{k,h+1}(x) &:= \max_a \bar{\mathbf{Q}}_{k,h+1}(x', a) \\ \bar{\mathbf{Q}}_{k,h}(x, a) &:= \hat{r}_k(x, a) + \mathbf{b}_k^{\text{rw}}(x, a) + \hat{p}_k(x, a)^\top \bar{\mathbf{V}}_{k,h+1} + \mathbf{b}_{k,h}(x, a) \end{aligned} \quad (23)$$

Lastly, suppose that $\mathbf{b}_k^{\text{rw}}(x, a)$ depends only on rewards collected when the state (x, a) is visited.

Note that this template subsumes the model-based approaches of [1, 16, 6], and if $\mathbf{b}^{\text{rw}}(x, a)$ is made to be time dependent, captures the approach of [5] as well. For the specific lower bound instance we consider, each stage $x \in \mathcal{S}$ can only be visited at a single stages $h \in [2]$, so \mathbf{b}^{rw} may be chosen to be time dependent without loss of generality. In order to capture the ‘‘model-free’’ methods based on Q-learning due to [9], we can instead mandate that

$$\bar{\mathbf{Q}}_{k,h}(x, a) := \hat{r}_k(x, a) + \mathbf{b}_k^{\text{rw}}(x, a) + (p(x, a)^\top \widehat{\mathbf{V}}_{h+1}^*) + \mathbf{b}_{k,h}(x, a),$$

where $(p(x, a)^\top \widehat{\mathbf{V}}_{h+1}^*)$ is a generalized estimate of $p(x, a)^\top \mathbf{V}_{h+1}^*$, and such that $(p(x, a)^\top \widehat{\mathbf{V}}_{h+1}^*)$ is nonnegative. In Lemma 4.2 in [9], one can see that we can take

$$(p(x, a)^\top \widehat{\mathbf{V}}_{h+1}^*) = \sum_{s=1}^{n_k(x,a)} \alpha_s \widehat{\mathbb{P}}_{k_s,h}(x, a)^\top \bar{\mathbf{V}}_{k_s,h+1}(x, a),$$

where k_s is the round at which (x, a) was selected for the s -th time, α_s is an appropriate weight, $\widehat{\mathbb{P}}_{k_s,h}(x, a)$ is the empirical probability estimate $\widehat{\mathbb{P}}_{k_s}(x, a)[x'] = \mathbb{I}(x' = x_{k_s,h+1})$ equal to indicator at the state $x_{k_s,h+1}$ visited after playing a at x at round k_s , and where $\bar{\mathbf{V}}_{k_s,h+1}$ is an optimistic estimate of \mathbf{V}_{h+1}^* at round k_s .

For simplicity, we shall work with the model based formulation (23), though the lower bound can be extended to this more general class.

Confidence Interval Assumptions: Our class of algorithms takes in a confidence parameter $\delta \in (0, 1/8)$. We shall also assume that there exists constants $c_{\text{bon}}, \bar{c}_{\text{bon}}$ such that, when the algorithm is run with parameter δ , the bonuses \mathbf{b}^{rw} and \mathbf{b}_k^{rw} satisfy⁷

$$\mathbf{b}_{k,h}(x, a) \geq \frac{c_{\text{bon}}}{1 \vee n_k(x, a)}, \quad c_{\text{bon}} \sqrt{\frac{\text{Var}[R(x, a)] \log(1/\delta)}{1 \vee n_k(x, a)}} \leq \mathbf{b}_k^{\text{rw}}(x, a) \leq \bar{c}_{\text{bon}} \sqrt{\frac{\log(M(1 \vee n_k(x, a))/\delta)}{1 \vee n_k(x, a)}}$$

We further assume that $\mathbf{b}^{\text{rw}}(x, a)$ is δ -correct, in the sense that,

$$\mathbb{P}[\forall x, a, k : \mathbf{b}_k^{\text{rw}}(x, a) + \hat{r}_k(x, a) \geq r(x, a)] \geq 1 - \delta.$$

Lastly, we shall assume that the optimistic overestimate is *consistent* in the sense that for any MDP \mathcal{M} with optimal value $\mathbf{V}_0^{*, \mathcal{M}}$, for any $\epsilon, \delta > 0$ there exists a function $f_{\mathcal{M}}$ such that

$$\mathbb{P}[\forall k \geq f_{\mathcal{M}}(\epsilon, \delta), \bar{\mathbf{V}}_{k,0} - \mathbf{V}_0^{*, \mathcal{M}} \leq \epsilon] \geq 1 - \delta.$$

Intuitively, this condition states that with high probability, the optimistic over-estimate of the value estimate approaches the expected reward under the optimal policy. Note that this does *not* assume uniform convergence of the entire value function itself, just the expected reward with respect to the initial state distribution p_0 on the optimal policy.

Remark G.1. Note that we do not require that our algorithm's confidence intervals are "inflated", in the sense that, with high probability, $\hat{r}_k(x, a) + \mathbf{b}_k^{\text{rw}}(x, a) - r(x, a) \geq c \mathbf{b}_k^{\text{rw}}(x, a)$, for a universal constant c . With this stronger assumption, we note that the proof of the lower bound can be simplified, and some restrictions on S, ϵ removed. In the interest of generality, we refrain from making this assumption.

G.3 Formal Lower Bound Instance

Consider the following simple game with $H = 2$, $\mathcal{A} = \{-1, +1\}$ and $\mathcal{S} = \{-S, \dots, -1, 0, 1, \dots, S\} = \mathcal{S}_- \cup \{0\} \cup \mathcal{S}_+$, where $\mathcal{S}_- = -[S]$ and $\mathcal{S}_+ = [S]$ (note $|\mathcal{S}| = 2S + 1$). The game always begins at state $x_1 = 0$ with two available actions, $a \in \{-1, +1\}$. Then, $x_2 | (x_1 = 0, a_1 = +1) \stackrel{\text{unif}}{\sim} \mathcal{S}_+$, and $x_2 | (x_1 = 0, a_1 = -1) \stackrel{\text{unif}}{\sim} \mathcal{S}_-$. Lastly, let \mathcal{D} denote any symmetric distribution on $[-1, 1]$ with $\Omega(1)$ variance. For $\epsilon \in (0, 1/8)$, we formally define the reward distributions

$$R(x, a) \sim \begin{cases} 0 & x = 0 \text{ or } a = -1 \\ \frac{1}{2} + \epsilon + \frac{1}{4} \mathcal{D} & (x, a) = (s, 1), s \in [S] \\ \frac{1}{2} + \frac{1}{4} \mathcal{D} & (x, a) = (-s, 1), s \in [S] \end{cases}.$$

It is straightforward to verify the following fact

Fact G.2. *The optimal action is always $a = 1$. Moreover, $\text{gap}_1(0, -1) = \text{gap}_{\min} = \epsilon$, whereas $\text{gap}_2(x, -1) \geq \frac{1}{2}$ for $x \neq 0$.*

In other words, all the gaps for suboptimal arms are $\Omega(1)$, except for the gap at state $x = 0$, which means for this instance with $H = 2$ and $A = 2$ we have $\sum_{x,a,h} \frac{1}{\text{gap}_h(x,a)} \approx S + \frac{1}{\epsilon}$. Nevertheless, we shall show that any algorithm in the class above suffers regret

$$\gtrsim \frac{S}{\epsilon} \log(1/\delta) = \frac{S}{\text{gap}_{\min}} \log(1/\delta).$$

G.4 The Lower Bound:

The Lower Bound: We first show that the optimistic Q-function relative to the optimal value at $(0, 1)$ decays at a rate of at least $\sqrt{S \log(1/\delta) / n_k(0, 1)}$. This will ultimately lead to incurring a regret of $\frac{S \log(1/\delta)}{\epsilon}$, despite the fact that all but one of the Q-function gaps are $\Omega(1)$.

⁷The quantity $\text{Var}[R(x, a)]$ below can also be replaced with an empirical variance, but we choose the true variance for simplicity.

Proposition G.3. *Let Alg denote an algorithm in the class described in Section G.2 run with confidence parameter $\delta \in (0, 1/8)$. Then there exists constants c_1, c_2, c_3 , depending only on the constants described in Section G.2, such that the following holds. For any $\epsilon \leq 1/\lceil c_1 S \log(S/\delta) \rceil$ and for $N = \lfloor c_2 S \log(1/\delta)/\epsilon^2 \rfloor$,*

$$\mathbb{P} [\forall k : n_k(0, -1) \leq N, \bar{\mathbf{Q}}_{k,1}(0, -1) - \mathbf{V}_1^*(0) \geq \epsilon] \geq 1 - Ne^{-c_3 S} - 2\delta$$

We now use Proposition G.3 to prove Theorem G.1. Note that $\mathbf{V}_1^*(0) = \mathbf{V}_0^*$. By assumption, with probability $1 - \delta$, $\bar{\mathbf{V}}_0 \leq \mathbf{V}_0^* + \eta$ after $f(\eta, \delta)$ rounds. Fix an appropriate ϵ and N in Proposition G.3 and let $K \geq f_{\mathcal{M}}(\epsilon/2, \delta) + N$. If $n_K(0, -1) > N$ times, then we have

$$\text{Regret}_K > \epsilon N \gtrsim \frac{S \log(1/\delta)}{\epsilon}$$

and the theorem is proved. Thus, suppose not so that $n_K(0, -1) \leq N$. Then by Proposition G.3 we have with high probability that

$$\bar{\mathbf{V}}_0 - \mathbf{V}_1^*(0) = \max_{a \in \{-1, 1\}} \bar{\mathbf{Q}}_{k,1}(0, a) - \mathbf{V}_1^*(0) \geq \bar{\mathbf{Q}}_{k,1}(0, -1) - \mathbf{V}_1^*(0) \geq \epsilon$$

However, by assumption $K \geq f_{\mathcal{M}}(\epsilon/2, \delta)$ which means that on an event that holds with probability at least $1 - \delta$, we have $\bar{\mathbf{V}}_0 - \mathbf{V}_1^*(0) = \max_{a \in \{-1, 1\}} \bar{\mathbf{Q}}_{k,1}(0, a) - \mathbf{V}_1^*(0) \leq \epsilon/2$, a contradiction.

G.4.1 Proof of Proposition G.3

Throughout, we will use upper case C_1, C_2, \dots to denote possibly changing numerical constants that depend on the constants in the definition of Alg, as set in Section G.2. The lower cast constants c_1, c_2 will be coincide with those in Proposition G.3.

Since $\mathbf{Q}_1^*(0, 1) = \frac{1}{2} + \epsilon$, it suffices to show that

$$\mathbb{P} \left[\forall k : n_k(0, -1) \leq N, \bar{\mathbf{Q}}_{k,1}(0, -1) - \frac{1}{2} \geq 2\epsilon \right] \geq 1 - Ne^{-C_3 S} - \delta$$

Fix an $n_0 = \lceil c_1 S / \log(S/\delta) \rceil$ for a constant c_1 be specified later, and let

$$\mathcal{E}^{\text{opt}} := \left\{ \forall k \geq 1, x \in \mathcal{S}_-, \hat{r}_k(x, 1) + \mathbf{b}_k^{\text{rw}}(x, a) \geq r(x, 1) = \frac{1}{2} \right\}.$$

By the optimism assumption, \mathcal{E}^{opt} holds with probability at least $1 - \delta$. First we verify that $\bar{\mathbf{Q}}_{k,1}(0, -1) - \frac{1}{2} \geq 2\epsilon$ for $0 \leq n_k(0, -1) \leq n_0$, provided that ϵ is sufficiently small:

Claim G.4. *Suppose that $\epsilon \leq \frac{c_{\text{bon}}}{2n_0}$. Then, with probability $1 - \delta$, $\bar{\mathbf{Q}}_{k,1}(0, -1) - \frac{1}{2} \geq 2\epsilon$ whenever $0 \leq n_k(0, -1) \leq n_0$:*

Proof. We have that

$$\bar{\mathbf{Q}}_{k,1}(0, -1) = \mathbf{b}_{k,1}(x, a) + \sum_{x' \in \mathcal{S}} \hat{p}_k(x'|0, -1) \bar{\mathbf{V}}_{k,2}(x').$$

Since $p(x|0, -1) = 0$ for $x \notin \mathcal{S}_-$, the empirical probability $\hat{p}(x|0, -1)$ is also 0, and thus

$$\begin{aligned} \bar{\mathbf{Q}}_{k,1}(0, -1) - \frac{1}{2} &= \mathbf{b}_{k,1}(x, a) + \sum_{x' \in \mathcal{S}_-} \hat{p}_k(x'|0, -1) (\bar{\mathbf{V}}_{k,2}(x') - \frac{1}{2}) \\ &\geq \mathbf{b}_{k,1}(x, a) + \min_{x' \in \mathcal{S}_-} (\bar{\mathbf{V}}_{k,2}(x') - \frac{1}{2}) \geq \mathbf{b}_{k,1}(x, a), \end{aligned} \quad (24)$$

where the first equality and first inequality use $\sum_{x' \in \mathcal{S}_-} \hat{p}(x'|0, -1) = 1$, and the second uses the optimistic event \mathcal{E}^{opt} to show that $\bar{\mathbf{V}}_{k,2}(x') \geq \hat{r}_k(x', 1) + \mathbf{b}_k^{\text{rw}}(x', 1) \geq r(x', 1) = \frac{1}{2}$ for $x' \in \mathcal{S}_-$. Using the assumption that $\mathbf{b}_{k,1}(x, a) \geq \frac{c_{\text{bon}}}{1 \vee n_k(x, a)}$, we see that if $n_k(x, a) \leq n_0$ and $\epsilon \leq \frac{c_{\text{bon}}}{2n_0}$, then $\mathbf{b}_{k,1}(x, a) \geq \frac{c_{\text{bon}}}{n_0} \geq 2\epsilon$, as needed. \square

Now, we turn to the case where $n_k(x, a) \in \{n_0, \dots, N\}$ for some $N = \lfloor c_2 S \log(1/\delta) / \epsilon^2 \rfloor$. It light of (24), it suffices to show that for $n_k \leq N$,

$$\sum_{x' \in \mathcal{S}_-} \hat{p}_k(x'|0, -1) (\bar{\mathbf{V}}_{k,2}(x') - \frac{1}{2}) \geq 2\epsilon. \quad (25)$$

By the definition of our algorithm class, the optimistic Q-function at stage $h = 2$ and pair (x, a) depend only at rewards collected at (x, a) , and the construction of our MDP, pairs (x, a) for $x \in \mathcal{S}_-$ are only accessible by playing $(0, -1)$. Hence, to analyze $\bar{\mathbf{Q}}_{k,1}(0, -1)$, for $n_0 \leq n_k(0, -1) \leq N$, it suffices to prove our described lower bound on $\bar{\mathbf{Q}}_{k,1}(0, -1)$ in the simplified game, where at each round $k = 1, 2, \dots$, the algorithm *always* selects $(0, -1)$, and show that for this algorithm

$$\Delta_0(k) := \sum_{x \in \mathcal{S}_-} \hat{p}_k(x|0, -1) (\bar{\mathbf{V}}_{k,2}(x) - \frac{1}{2}) \geq 2\epsilon, \forall k \in \{n_0, \dots, N\}.$$

Turning our attention to this simplified game, for $x \in \mathcal{S}_-$ let $n_k(x)$ denote the number of times x has been visited up to round k , and recall $n_k(x, a)$ is the number of times action a is played at stage s . Further, set

$$\Delta(x, k) := \bar{\mathbf{V}}_{k,2}(x) - \frac{1}{2}$$

We now make a couple of observations

- (a) The vector $(n_k(x))_{x \in \mathcal{S}_-}$ is a uniform multinomial on the states in \mathcal{S}_- .
- (b) Conditioned on $(n_k(x))_{x \in \mathcal{S}_-}$, we can see that the values of $\bar{\mathbf{V}}_{k,2}(x)$ are independent, because for each $x \in \mathcal{S}_-$, the game decouples into $n_k(x)$ rounds of a two arm bandit game on actions $a \in \{-1, 1\}$.

Using these observations, we prove the following claim:

Claim G.5. *There exists constants C_1, C_2 such that for any $x \in \mathcal{S}_-$, if $\delta \leq 1/8$ and $n_k(x) \geq C_1 \log(M/\delta)$, then conditioned on the history $(n_j(x'))_{x' \in \mathcal{S}_-, j \geq 1}$, the following event holds with probability at least $1/4$:*

$$\mathcal{E}_k^\Delta(x) := \left\{ \Delta(k, x) := \bar{\mathbf{V}}_{k,2}(x) - \frac{1}{2} \geq C_2 \sqrt{\log(1/\delta) / n_k(x)} \right\},$$

and the events $\{\mathcal{E}_j^\Delta(x) : x \in \mathcal{S}_-\}$ are mutually independent (again, given $(n_j(x'))_{x' \in \mathcal{S}_-, j \geq 1}$).

Therefore, on the optimistic event \mathcal{E}^{opt} , where $\{\Delta(k, x) \geq 0\}$, we can lower bound (again, in the simplified game where we always select action $(0, -1)$),

$$\begin{aligned} \Delta_0(k) &\geq \sum_{x \in \mathcal{S}_-} \hat{p}(x'|0, -1) \Delta(x, k) \\ &\geq \sum_{x \in \mathcal{S}_-} \hat{p}(x'|0, -1) \mathbb{I}(\mathcal{E}_k^\Delta(x)) C_2 \sqrt{\frac{\log(1/\delta)}{n_k(x)}} \\ &\stackrel{(i)}{=} \sum_{x \in \mathcal{S}_-} \frac{n_k(x)}{k} \mathbb{I}(\mathcal{E}_k^\Delta(x)) C_2 \sqrt{\frac{\log(1/\delta)}{n_k(x)}} \\ &= \frac{C_2 \sqrt{\log(1/\delta)}}{k} \sum_{x \in \mathcal{S}_-} \mathbb{I}(\mathcal{E}_k^\Delta(x)) \sqrt{n_k(x)} \end{aligned}$$

where (i) uses the fact that for $x \in \mathcal{S}_-$ is only accessible through $(0, -1)$, and that $(0, -1)$ is always selected in the simplified game. Next, observe that in the simplified game, $\bar{n}_k(x) = k/S$, so that if $n_0/S \geq C_3 \log(1/\delta)$ for some constant C_3 , it holds by an argument similar to Lemma B.7 that with

probability $1 - \delta$, the event $\mathcal{E}_1 := \{\forall x \in \mathcal{S}_-, \forall k \geq n_0, n_k(x) \geq \bar{n}_k(x)/4 = k/4S\}$ holds, yielding

$$\begin{aligned} \Delta_0(k) &\geq \frac{C_2 \sqrt{\log(1/\delta)}}{2k} \sqrt{k/S} \cdot \sum_{x \in \mathcal{S}_-} \mathbb{I}(\mathcal{E}_k^\Delta) \\ &= \frac{C_2}{2} \sqrt{S \log(1/\delta)/k} \cdot \left(\frac{1}{S} \sum_{x \in \mathcal{S}_-} \mathbb{I}(\mathcal{E}_k^\Delta) \right). \end{aligned}$$

Finally, if in addition $n_0/4S \geq C_1 \log(1/\delta)$, where C_1 is the constant from claim G.5, then on \mathcal{E}_1 , it holds that for $k \geq n_0$, $n_k(x) \geq C_1 \log(1/\delta)$. We then set the constant c_1 so that $n_0/S \geq C_3 \log(1/\delta)$ and $n_0/4S \geq C_1 \log(1/\delta)$ hold.

Lastly, since (a) \mathcal{E}_1 is measurable with respect to the counts $(n_j(x'))_{x' \in \mathcal{S}_-, j \geq 1}$, (b) since $\mathcal{E}_k^\Delta(x)$ are independent given these counts, and (c) $\mathbb{E}[\mathbb{I}(\mathcal{E}_k^\Delta)] \geq 1/4$, a Chernoff bound shows that for $k \geq n_0$, the event $\mathcal{E}_2(k) := \left\{ \left(\frac{1}{S} \sum_{x \in \mathcal{S}_-} \mathbb{I}(\mathcal{E}_k^\Delta) \right) \geq 1/8 \right\}$ holds with probability at least $e^{-C_5 S}$ conditioned on \mathcal{E}_1 . Hence, on $\mathcal{E}^{\text{opt}} \cap \mathcal{E}_1 \cap \bigcup_{k=n_0}^N \mathcal{E}_2(k)$, we have

$$\Delta_0(k) \geq \frac{C_2}{16} \sqrt{S \log(1/\delta)/k} \geq \frac{C_2}{16} \sqrt{S \log(1/\delta)/N}, \quad \forall k \in \{n_0, \dots, N\}.$$

Hence, if $N \leq c_2 \frac{S \log(1/\delta)}{\epsilon^2}$ for some constant c_2 , we see that $\Delta_0(k) \geq 2\epsilon$ for all $k \in \{n_0, \dots, N\}$. Lastly, we see that

$$\begin{aligned} \mathbb{P}[(\mathcal{E}^{\text{opt}} \cap \mathcal{E}_1 \cap \bigcup_{k=n_0}^N \mathcal{E}_2(k))^c] &\leq \mathbb{P}[(\mathcal{E}^{\text{opt}})^c] + \mathbb{P}[\mathcal{E}_1^c] + \mathbb{P}[(\bigcup_{k=n_0}^N \mathcal{E}_2(k))^c \wedge \mathcal{E}_1] \\ &\leq \mathbb{P}[(\mathcal{E}^{\text{opt}})^c] + \mathbb{P}[\mathcal{E}_1^c] + N \max_{k \geq n_0} \mathbb{P}[\mathcal{E}_2(k)^c \mid \mathcal{E}_1] \leq 2\delta + N e^{-C_4 S}. \end{aligned}$$

Translating to the non-simplified game, we have therefore established that

$$\mathbb{P}[\forall k : n_0 \leq n_k(0, -1) \leq N, \bar{\mathbf{Q}}_{k,1}(-1, 1) - \frac{1}{2} \geq 2\epsilon] \geq 1 - 2\delta + N e^{-C_4 S}.$$

Combining with the additional probability of error δ for the case $n_k(0, -1) \leq n_0$ concludes the proof.

G.5 Proof of Claim G.5

We observe that conditioned on the vector $(n_j(x'))_{x' \in \mathcal{S}_-, j \geq 1}$, the games at states x and round k are equivalent to S independent two-arm bandit games with $n_k(x)$ rounds. Note moreover that $\Delta(x, k) = \bar{\mathbf{V}}_{k,2}(x) - \frac{1}{2} \geq \mathbf{b}_k^{\text{rw}}(x, 1) + \hat{r}_k(x, 1) - \frac{1}{2}$. Hence, restricting to a single state x (and dropping the dependence on x for simplicity), it suffices to show that for k rounds of an appropriate two-arm bandit game with $a \in \{-1, 1\}$ with empirical rewards $\hat{r}_k(a)$ and bonuses $\mathbf{b}_k^{\text{rw}}(a)$, $R(-1) = 0$ and $R(1) \sim \frac{1}{2} + \frac{1}{4}\mathcal{D}$, that

$$\forall k \geq C_1 \log(S/\delta), \quad \mathbb{P}[\mathbf{b}_k^{\text{rw}}(1) + \hat{r}_k(1) - \frac{1}{2} \geq \sqrt{\log(1/\delta)/k}] \geq \frac{1}{4}$$

where we have dropped the dependence on x for simplicity. Throughout, we will also use the notation C_1, C_2, C_3 to denote constants specific to the proof of Claim G.5, and reserve C_1, C_2 for the constants in the claim statement.

If $\delta \leq 1/8$, then a standard argument shows that for some constant C_1 (depending on \bar{c}_{bon}), $n_k(-1) \leq C_1 \log(S/\delta)$. Indeed, define the event $\mathcal{E}_0 := \{\forall k \geq 1 : \mathbf{b}_k^{\text{rw}}(1) + \hat{r}_k(1) \geq r(1) = \frac{1}{2}\}$; by assumption on our confidence intervals, complement of this event occurs with probability at most $\delta \leq 1/8$. Note also that on \mathcal{E}_0 , since $R(-1) = 0$ with probability 1, it holds that for any $j \leq k$ with $n_j(-1) \geq C_1 \log(S/\delta)$

$$\begin{aligned} \hat{r}_j(-1) + \mathbf{b}_j^{\text{rw}}(x, a) &= \mathbf{b}_j^{\text{rw}}(x, a) \stackrel{(i)}{\leq} \bar{c}_{\text{bon}} \sqrt{\frac{\log(S n_j(-1)/\delta)}{n_j(-1)}} \\ &\stackrel{(ii)}{\leq} \frac{1}{2} = r(1) \leq \hat{r}_j(1) + \mathbf{b}_j^{\text{rw}}(1), \end{aligned}$$

where in (i) we used the definition of the confidence interval with $M \lesssim S$, and in (ii) we used $n_j(-1) \geq C_1 \log(S/\delta)$ for an appropriately tuned constant C_1 . Since $a_j := \arg \max_a \hat{r}_j(a) + \mathbf{b}_j^{\text{rw}}(a)$, we have $a_j = 1$. This implies that $n_k(-1) \leq \max_{j \geq 1} n_j(-1) \leq C_1 \log(M/\delta)$.

Next, set $k_0 = C_1 \log(M/\delta)$. We wish to show that for $k \geq k_0$,

$$\hat{r}_k(1) + \mathbf{b}_k^{\text{rw}}(1) \gtrsim \frac{\log(1/\delta)}{k}$$

There are two technical challenges: first, the confidence interval $\mathbf{b}_k^{\text{rw}}(1)$ might be nearly tight, so that we cannot show that with high probability, $\hat{r}_k(1) + \mathbf{b}_k^{\text{rw}}(1) \gtrsim \mathbf{b}_k^{\text{rw}}(1)$. Second, because the algorithm adaptively chooses to sample actions $a \in \{-1, 1\}$, $\hat{r}_k(1)$ *does not* have the distribution of $n_k(1)$ i.i.d. samples from $R(1)$.

We can get around this as follows. We can imagine all rewards sampled from action 1 as being drawn at the start of the game, and constituting a sequence $R^{(1)}(1), R^{(2)}(1), \dots$ and so on. Then, $\hat{r}_k(1)$ is the average of the samples $1, \dots, n_k(1)$, where $n_k(1) \leq k$. Therefore

$$\begin{aligned} n_k(1)(\hat{r}_k(1) - \frac{1}{2}) &= \sum_{i=1}^{n_k(1)} (R^{(i)}(1) - \frac{1}{2}) = \sum_{i=1}^k (R^{(i)}(1) - \frac{1}{2}) - \sum_{i=n_k(1)+1}^k (R^{(i)}(1) - \frac{1}{2}). \\ &= \sum_{i=1}^k (R^{(i)}(1) - \frac{1}{2}) - \sum_{i=k-n_k(-1)+1}^k (R^{(i)}(1) - \frac{1}{2}), \end{aligned}$$

where the last line uses $n_k(1) + n_k(-1) = k$.

Now consider the event $\mathcal{E}_1(\delta) := \{n_k(-1) \leq k_0\}$, where we recall $k_0 = C_1 \log(M/\delta)$ was our $1 - \delta$ -probability upper bound on $n_k(-1)$. On $\mathcal{E}_1(\delta)$, $n_k(-1) = j$ for some $j \in \{0, 1, \dots, k_0\}$, and we can lower bound the above expression by

$$\geq \sum_{i=1}^k (R^{(i)}(1) - \frac{1}{2}) - \max_{j=0, \dots, k_0} \sum_{i=k-j+1}^k (R^{(i)}(1) - \frac{1}{2}).$$

Observe now that we have lower bounded $n_k(1)(\hat{r}_k(1) - \frac{1}{2})$ in terms of quantities depending only on the i.i.d. reward sequence $(R^{(i)}(1))$, and *not* on the quantities $n_k(-1), n_k(1)$.

Moreover, a standard maximal inequality implies that the following event $\mathcal{E}_2(\delta)$ holds for an appropriate constant C_2 with probability $1 - \delta$:

$$\mathcal{E}_2(\delta) := \left\{ \max_{j=0, \dots, k_0} \sum_{i=k-j+1}^k (R^{(i)}(1) - \frac{1}{2}) \leq C_2 \sqrt{k_0 \log(1/\delta)} \right\} \quad (26)$$

Lastly, since $R^{(i)}$ is symmetric, we have that the following event \mathcal{E}_3 holds with probability $1/2$:

$$\mathcal{E}_3 := \left\{ \sum_{i=1}^k (R^{(i)}(1) - \frac{1}{2}) \geq 0 \right\}.$$

Hence, on $\mathcal{E}_1(\delta) \cap \mathcal{E}_2(\delta) \cap \mathcal{E}_3$,

$$\begin{aligned} n_k(1)(\hat{r}_k(1) - \frac{1}{2}) &\geq \underbrace{\sum_{i=1}^k (R^{(i)}(1) - \frac{1}{2})}_{\geq 0} - \max_{j=0, \dots, k_0} \sum_{i=k-j+1}^k (R^{(i)}(1) - \frac{1}{2}) \\ &\geq -C_2 \sqrt{\log(1/\delta) k_0}. \end{aligned}$$

If we further assume that $k \geq 2C_1 \log(M/\delta)$, then $n_k(-1) \leq k_0 \leq k/2$, so that $\mathcal{E}_1(\delta)$ implies $n_k(1) \geq k/2$. Dividing both sides of the above by k and bringing $1/k$ into the square root yields (again on $\mathcal{E}_1(\delta) \cap \mathcal{E}_2(\delta) \cap \mathcal{E}_3$)

$$(\hat{r}_k(1) - \frac{1}{2}) \geq -\frac{C_2}{2} \sqrt{\frac{\log(1/\eta)}{k} \cdot \frac{k_0}{k}}. \quad (27)$$

Moreover, by the lower bound assumption on \mathbf{b}^{rw} and the fact that $R(1)$ has $\Omega(1)$ variance, there exists some constant C_3 such that

$$\mathbf{b}^{\text{rw}}(n_k(1)) \geq c_{\text{bon}} \sqrt{\frac{\text{Var}[R(x, a)] \log(1/\delta)}{n_k(1)}} \geq C_3 \sqrt{\frac{\log(1/\delta)}{k}},$$

where again we use $n_k(1) \leq k$. Combining with (27), we have on $\mathcal{E}_1(\delta) \cap \mathcal{E}_2(\delta) \cap \mathcal{E}_3$ that

$$(\hat{r}_k(1) - \frac{1}{2}) + \mathbf{b}^{\text{rw}}(n_k(1)) \geq C_3 \sqrt{\frac{\log(1/\delta)}{k}} - \frac{C_2}{2} \sqrt{\frac{\log(1/\delta)}{k}} \cdot \frac{k_0}{k}.$$

Hence, if $k_0/k \leq (C_3/C_2)^2$, or equivalently if $k \geq C_1 (C_3/C_2)^{-2} \log(M/\delta)$, then

$$(\hat{r}_k(1) - \frac{1}{2}) + \mathbf{b}^{\text{rw}}(n_k(1)) \geq \frac{C_3}{2} \sqrt{\frac{\log(1/\delta)}{k}} \geq \frac{C_3}{2} \sqrt{\frac{\log(1/\delta)}{2n_k(1)}}$$

on the event $\mathcal{E}_1(\delta) \cap \mathcal{E}_2(\delta) \cap \mathcal{E}_3$. Lastly, for $\delta \leq 1/8$, we note $\mathbb{P}[\mathcal{E}_1(\delta) \cap \mathcal{E}_2(\delta) \cap \mathcal{E}_3] \geq \frac{1}{2} - 2\delta \geq 1/4$. Recalling our earlier condition $k \geq 2C_1 \log(M/\delta)$, the claim now holds with by setting the constant C_1 in the claim statement to be $C_1 \max\{2, (C_3/C_2)^{-2}\}$, and C_2 to be $\frac{C_3}{2\sqrt{2}}$.

H Information Theoretic Lower Bound (Proposition 2.2)

In this section we construct give a proof of the information theoretic lower bound Proposition 2.2, as well as a non-asymptotic bound that holds even for non-uniformly good algorithms.

H.1 Construction of the hard instance

Our construction mirrors the lower bounds due to [4], but with specific and non-uniform gaps. We define \mathcal{M} as an MDP on state space $\mathcal{S} = [S + 2]$, with actions $\mathcal{A} = [A]$, and horizon $[H]$. We will first state the construction for $H \geq 2$, and then remark on the modification for $H = 1$ at the end of the section. For $a \in [A]$, $x \in [S]$, we set

$$p(x' = S + 1|x, a) = \frac{3}{4} - \frac{2}{H-1}\Delta_{x,a}, \quad p(x' = S + 2|x, a) = 1 - p(x' = S + 1|x, a).$$

Furthermore, we set the initial state to have the distribution $x_1 \stackrel{unif}{\sim} [S]$, and set

$$p(x' = S + 1|x = S + 1, a) = 1, \quad p(x' = S + 2|x = S + 2, a) = 1 \forall a \in [A].$$

Finally, the rewards are set deterministically as

$$R(x, a) := \begin{cases} 0 & x \in [S] \\ 0 & x \in \{S + 1, S + 2\}, a > 1 \\ 1 & (x, a) = (S + 1, 1) \\ \frac{1}{2} & (x, a) = (S + 2, 1) \end{cases}$$

We may then verify that $\mathbf{V}_h^*(S + 1) = (H - h + 1)$ and $\mathbf{V}_h^*(S + 2) = (h - H + 1)/2$, which implies that that for $x \in [S]$,

$$\begin{aligned} \text{gap}_h(x, a) &= \left(\max_{a'} \sum_{x'} p(x'|x, a') \mathbf{V}_{h+1}^*(x') \right) \\ &= \max_{a'} (p(S + 1|x, a') - p(S + 1|x, a))(H - h) + (p(S + 2|x, a') - p(S + 2|x, a)) \frac{(H - h)}{2} \\ &= \max_{a'} (p(S + 1|x, a') - p(S + 1|x, a))(H - h) - (p(S + 1|x, a') - p(S + 1|x, a)) \frac{(H - h)}{2} \\ &= \max_{a'} (p(S + 1|x, a') - p(S + 1|x, a)) \frac{H - h}{2} \\ &= \frac{2\Delta_{x,a}}{(H - 1)} \cdot \frac{H - h}{2}, \end{aligned}$$

and in particular that $\text{gap}_1(x, a) = \Delta_{x,a}$. For $H = 1$, the construction is modified so that $\mathcal{S} = [S]$, and

$$R(x, a) \sim \text{Bernoulli}\left(\frac{3}{4} - \Delta_{x,a}\right), \text{ and } x_1 \stackrel{unif}{\sim} [S].$$

Then, we see that $\text{gap}_1(x, a) = \Delta_{x,a}$. In what follows, we will address the $H \geq 2$ case; the case $H = 1$ will follow from similar, but simpler arguments.

H.2 Regret Lower Bound Decomposition

We can now lower bound the expected regret as

$$\begin{aligned}
\mathbb{E}^{\mathcal{M}}[\text{Regret}_K] &:= \mathbb{E}^{\mathcal{M}}\left[\sum_{k=1}^K \mathbf{V}_0^* - \mathbf{V}_0^{\pi_k}\right] \\
&:= \mathbb{E}\left[\sum_{k=1}^K \sum_x p(x_1 = x) \{\mathbf{V}_1^*(x) - \mathbf{V}_1^{\pi_k}(x)\}\right] \\
&\stackrel{(i)}{\geq} \mathbb{E}^{\mathcal{M}}\left[\sum_{k=1}^K \sum_x p(x_1 = x) \{\mathbf{V}_1^*(x) - \mathbf{Q}_1^*(x, \pi_{k,1}(x))\}\right] \\
&= \mathbb{E}^{\mathcal{M}}\left[\sum_{k=1}^K \sum_x p(x_1 = x) \text{gap}_1(x, \pi_{k,1}(x))\right] \\
&= \mathbb{E}^{\mathcal{M}}\left[\sum_{k=1}^K \sum_{x,a} p(x_1 = x) \mathbb{I}(\pi_{k,1}(x) = a) \text{gap}_1(x, a)\right] \\
&= \sum_{x,a} \mathbb{E}^{\mathcal{M}}[\bar{n}_K(x, a)] \text{gap}_1(x, a), \tag{28}
\end{aligned}$$

where inequality (i) follows since $\mathbf{V}_1^{\pi_k}(x) = \mathbf{Q}_1^{\pi_k}(x, \pi_{k,1}(x)) \leq \mathbf{Q}_1^*(x, \pi_{k,1}(x))$. We now show that for all sufficiently large $K \geq K_0(\mathcal{M})$, any uniformly correct algorithm must have

$$\begin{aligned}
\forall(x, a) : x \in [S], \text{gap}_1(x, a) > 0, \forall K \geq K_0(\mathcal{M}) \\
\mathbb{E}^{\mathcal{M}}[\bar{n}_K(x, a)] \gtrsim \left(\frac{2}{H-1} \Delta_{x,a}\right)^{-2} \log K \gtrsim \frac{H^2}{\Delta_{x,a}^2} \log K = \frac{H^2}{\text{gap}_1(x, a)^2} \log K, \tag{29}
\end{aligned}$$

which concludes the proof since

$$\mathbb{E}^{\mathcal{M}}[\text{Regret}_K] \gtrsim \sum_{x,a:\text{gap}_1(x,a)>0} \frac{H^2}{\text{gap}_1(x, a)^2} \log K \cdot \text{gap}_1(x, a) = \sum_{x,a:\text{gap}_1(x,a)>0} \frac{H^2}{\text{gap}_1(x, a)} \log K.$$

We further note that this argument can also show that, for all K sufficiently large and all $h \in [H-1]$

$$\mathbb{E}^{\mathcal{M}}[\text{Regret}_K] \gtrsim \sum_{x,a:\text{gap}_h(x,a)>0} \frac{(H-h)^2}{\text{gap}_h(x, a)} \log K. \tag{30}$$

as well.

H.3 Proof of Equation (29)

Throughout, we fix a state $x \in [S]$, and an action $a : \text{gap}_1(x, a) > 0$. We shall further introduce the shorthand

$$\Delta_{x,a} := \frac{2\Delta_{x,a}}{H-1} \in (0, 1/2), \tag{31}$$

where the bound on $\Delta_{x,a}$ follows from $\Delta_{x,a} \in (0, H/8)$.

To lower bound Equation (29), we follow steps analogues to standard information theoretic lower bounds. Our exposition will follow [7]. First, we state a lemma which is the MDP analogue of Equation (6) in [7]. Its proof is analogous, and omitted for the sake of brevity:

Lemma H.1. *Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, H, r, p^{\mathcal{M}}, p_0, R^{\mathcal{M}})$ and $\mathcal{M}' = (\mathcal{S}, \mathcal{A}, H, r, p^{\mathcal{M}'}, p_0, R^{\mathcal{M}'})$ denote two episodic MDPs with the same state space \mathcal{S} , action space \mathcal{A} and horizon h , and initial state distribution p_0 . For any $(x, a) \in \mathcal{S} \times \mathcal{A}$, let $\nu^{\mathcal{M}}(x, a)$ denote the law of the joint distribution of (X', R) where $X' \sim p^{\mathcal{M}}(\cdot|x, a)$ and $R \sim R^{\mathcal{M}}(x, a)$; define the law $\nu^{\mathcal{M}'}(x, a)$ analogously.*

Finally, fix a horizon $K \geq 1$, and let \mathcal{F}_K denote the filtration generated by all rollouts up to episode K . Then, for any \mathcal{F}_K -measurable random variable $Z \in [0, 1]$,

$$\text{kl}(\mathbb{E}^{\mathcal{M}}[Z], \mathbb{E}^{\mathcal{M}'}[Z]) \leq \sum_{x,a} \mathbb{E}^{\mathcal{M}}[\bar{n}_K(x,a)] \text{KL}(\nu^{\mathcal{M}}(x,a), \nu^{\mathcal{M}'}(x,a)),$$

where $\text{kl}(x,y) = x \log \frac{x}{y} + (1-x) \log \frac{1-x}{1-y}$ denotes the binary KL-divergence, and $\text{KL}(\cdot, \cdot)$ denotes the KL-divergence between two probability laws.

We apply the above lemma as follows. For our fixed pair (x, a) , define an alternate \mathcal{M}' to be the MDP which coincides with \mathcal{M} except that

$$p(x'|x, a) = \frac{3}{4} + \eta, \quad \eta = \min\{7/8, \frac{3}{4} + \Delta_{x,a}\}$$

By construction, \mathcal{M} and \mathcal{M}' differ only at their law at (x, a) . Thus,

$$\text{kl}(\mathbb{E}^{\mathcal{M}}[Z], \mathbb{E}^{\mathcal{M}'}[Z]) \leq \mathbb{E}^{\mathcal{M}}[\bar{n}_K(x, a)] \text{KL}(\nu^{\mathcal{M}}(x, a), \nu^{\mathcal{M}'}(x, a)).$$

The following lower bound controls the KL divergence between the laws $\nu^{\mathcal{M}}(x, a), \nu^{\mathcal{M}'}(x, a)$:

Claim H.2. *There exists a universal constant c such that*

$$\text{KL}(\nu^{\mathcal{M}}(x, a), \nu^{\mathcal{M}'}(x, a)) \lesssim c \Delta_{x,a}^2.$$

Proof. At (x, a) , $R(x, a) = 0$ with probability one under both $\mathcal{M}, \mathcal{M}'$. Moreover, recall that under \mathcal{M} , (x, a) transition to state $S+1$ with probability $\frac{3}{4} - \Delta_{x,a}$, and to $S+2$ with probability $1 - (\frac{3}{4} - \Delta_{x,a})$. On the other hand, \mathcal{M}' transition to $S+1$ with probability $\frac{3}{4} + \eta$, and $S+2$ with probability $1 - (\frac{3}{4} + \eta)$. Consequently both laws are equivalent to Bernoulli distributions with parameters $\frac{3}{4} - \Delta_{x,a}$ and $\frac{3}{4} + \eta$, respectively. Since $\text{kl}(x, y)$ is precisely $\text{KL}(\text{Bernoulli}(x), \text{Bernoulli}(y))$ for $x, y \in (0, 1)$,

$$\text{KL}(\nu^{\mathcal{M}}(x, a), \nu^{\mathcal{M}'}(x, a)) = \text{kl}\left(\frac{3}{4} - \Delta_{x,a}, \frac{3}{4} + \eta\right).$$

Lastly, set $x = \frac{3}{4} - \Delta_{x,a}$ and $y = \frac{3}{4} + \min\{\frac{7}{8}, \Delta_{x,a}\}$. We $y - x \leq 2\Delta_{x,a}$, and by assumption on $\Delta_{x,a} \leq 1/2$, Thus, $1/4 \leq x \leq y \leq 7/8$. Hence, a standard Taylor expansion (e.g. **(author?)** [13, Lemma E.1]) shows that there exists a universal constant c such that $\text{kl}(x, y) \leq \frac{c}{(x-y)^2} \leq \frac{4c}{\Delta_{x,a}^2}$, as needed. \square

As a consequence, we see that for any \mathcal{F}_K -measurable $Z \in [0, 1]$, we find

$$\mathbb{E}^{\mathcal{M}}[\bar{n}_K(x, a)] \gtrsim \Delta_{x,a}^{-2} \text{kl}(\mathbb{E}^{\mathcal{M}}[Z], \mathbb{E}^{\mathcal{M}'}[Z]) \gtrsim \frac{H^2}{\Delta_{x,a}^2} \text{kl}(\mathbb{E}^{\mathcal{M}}[Z], \mathbb{E}^{\mathcal{M}'}[Z]),$$

where the last inequality uses that $\Delta_{x,a} \lesssim \Delta_{x,a}/H$.

To conclude, it suffices to exhibit a random variable Z_K such that, for K sufficiently large,

$$\text{kl}(\mathbb{E}^{\mathcal{M}}[Z], \mathbb{E}^{\mathcal{M}'}[Z]) \gtrsim (1 - \alpha) \log K.$$

To this end, consider $Z_K = \frac{S\bar{n}_K(x,a)}{K}$. Note that since x is only visited with probability at most $1/S$ at stage $h = 1$, and with probability 0 for stages $h \geq 2$, we have

$$\bar{n}_K(x, a) = \sum_{k=1}^K \mathbb{E}^{\mathcal{M}}[\mathbb{P}(x_1 = x) \mathbb{I}(\pi_{k,1}(x_1) = 1)] = \frac{1}{S} \sum_{k=1}^K \mathbb{E}^{\mathcal{M}}[\mathbb{I}(\pi_{k,1}(x_1) = 1)] \leq K/S,$$

which implies that, $Z_K \in [0, 1]$ with probability one. Moreover, note that by an argument similar to that of (28), that under the MDP \mathcal{M}' ,

$$\mathbb{E}^{\mathcal{M}'}[\text{Regret}_K] \geq \eta \mathbb{E}^{\mathcal{M}'}\left[\sum_{a' \neq a} \bar{n}_K(x, a')\right] = \eta \left(\frac{K}{S} - \mathbb{E}^{\mathcal{M}'}[\bar{n}_K(x, a')]\right) = \frac{\eta K}{S} (1 - \mathbb{E}^{\mathcal{M}'}[Z_K]).$$

Hence, if Alg is α -uniformly good, then there exists a constant $C_{\mathcal{M}'}$ such that

$$1 - \mathbb{E}^{\mathcal{M}'}[Z_K] \leq \frac{C_{\mathcal{M}'}\eta}{S} K^{\alpha-1}.$$

By the same token, there exists a constant $C_{\mathcal{M}}$ such that

$$C_{\mathcal{M}}K^\alpha \geq \mathbb{E}^{\mathcal{M}}[\text{Regret}_K] \geq \text{gap}_1(x, a)\mathbb{E}^{\mathcal{M}}[\bar{n}_K(x, a)] = \frac{K\text{gap}_1(x, a)}{S}\mathbb{E}^{\mathcal{M}}[Z_K]\text{gap}_1(x, a)/S.$$

which implies that $\mathbb{E}^{\mathcal{M}}[Z_K] \leq \frac{SC_{\mathcal{M}}K^{\alpha-1}}{\text{gap}_1(x, a)}$. Furthermore, by Inequality (11) in [7], it holds that

$$\text{kl}(x, y) \geq (1 - x) \log \frac{1}{1 - y} - \log 2$$

which implies that for K sufficiently large,

$$\text{kl}(\mathbb{E}^{\mathcal{M}}[Z_K], \mathbb{E}^{\mathcal{M}'}[Z_K]) \geq \left(1 - \frac{SC_{\mathcal{M}}K^{\alpha-1}}{\text{gap}_1(x, a)}\right) \left\{ (1 - \alpha) \log K - \log \frac{C_{\mathcal{M}'}\eta}{2S} \right\} \gtrsim (1 - \alpha) \log K.$$