

Softmax Policy Gradient Methods Can Take Exponential Time to Converge

Gen Li*
UPenn

Yuting Wei*
UPenn

Yuejie Chi†
CMU

Yuxin Chen*
UPenn

February 2021; Revised: December 2022

Abstract

The softmax policy gradient (PG) method, which performs gradient ascent under softmax policy parameterization, is arguably one of the de facto implementations of policy optimization in modern reinforcement learning. For γ -discounted infinite-horizon tabular Markov decision processes (MDPs), remarkable progress has recently been achieved towards establishing global convergence of softmax PG methods in finding a near-optimal policy. However, prior results fall short of delineating clear dependencies of convergence rates on salient parameters such as the cardinality of the state space \mathcal{S} and the effective horizon $\frac{1}{1-\gamma}$, both of which could be excessively large. In this paper, we deliver a pessimistic message regarding the iteration complexity of softmax PG methods, despite assuming access to exact gradient computation. Specifically, we demonstrate that the softmax PG method with stepsize η can take

$$\frac{1}{\eta} |\mathcal{S}|^{2^{\Omega\left(\frac{1}{1-\gamma}\right)}} \text{ iterations}$$

to converge, even in the presence of a benign policy initialization and an initial state distribution amenable to exploration (so that the distribution mismatch coefficient is not exceedingly large). This is accomplished by characterizing the algorithmic dynamics over a carefully-constructed MDP containing only three actions. Our exponential lower bound hints at the necessity of carefully adjusting update rules or enforcing proper regularization in accelerating PG methods.

Keywords: policy gradient methods, exponential lower bounds, softmax parameterization, discounted infinite-horizon MDPs

Contents

1	Introduction	2
1.1	Main result	3
1.2	Other related works	5
1.3	Paper organization	6
2	Background	6
3	Construction of a hard MDP	7
4	Analysis: proof outline	10
4.1	Preparation: crossing times and choice of constants	11
4.2	A high-level picture	12
4.3	Proof outline	13

This work was presented in part at the Conference on Learning Theory (COLT) 2021.

*Department of Statistics and Data Science, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA.

†Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

5 Discussion	20
A Preliminary facts	24
A.1 Basic properties of the constructed MDP	24
A.2 A type of recursive relations	26
A.3 Proof of Lemma 1	27
A.4 Proof of Lemma 8	28
A.5 Proof of Lemma 11	30
B Discounted state visitation probability (Lemmas 2-3)	31
B.1 Lower bounds: proof of Lemma 2	31
B.2 Upper bounds: proof of Lemma 3	32
C Crossing times of the first few states (Lemma 4)	38
C.1 Crossing times for the buffer states in \mathcal{S}_1 and \mathcal{S}_2	38
C.2 Crossing times for the adjoint state $\bar{1}$	41
C.3 Auxiliary facts	44
D Analysis for the initial stage (Lemma 5)	45
D.1 Two key properties	46
D.2 Proof of the properties (133) and (134)	46
D.3 Auxiliary facts	49
E Analysis for the intermediate stage (Lemma 6)	51
E.1 Main steps	51
E.2 Proof of Lemma 17	52
E.3 Proof of Lemma 18	57
F Analysis for the blowing-up lemma (Lemma 7)	59
F.1 Which reference point t_{ref} shall we choose?	59
F.2 Stage I: the duration where $\theta^{(t)}(s, a_2) < \theta^{(t)}(s, a_0)$	60
F.3 Stage II: the duration where $\theta^{(t)}(s, a_2) \geq \theta^{(t)}(s, a_0)$	65
F.4 Proof of the claims (45a) and (45b)	68
F.5 Proof of the claim (45c)	68

1 Introduction

Despite their remarkable empirical popularity in modern reinforcement learning (Mnih et al., 2015; Silver et al., 2016), theoretical underpinnings of policy gradient (PG) methods and their variants (Kakade, 2002; Konda and Tsitsiklis, 2000; Peters and Schaal, 2008; Sutton et al., 2000; Williams, 1992) remain severely obscured. Due to the nonconcave nature of value function maximization induced by complicated dynamics of the environments, it is in general highly challenging to pinpoint the computational efficacy of PG methods in finding a near-optimal policy. Motivated by their practical importance, a recent strand of work sought to make progress towards demystifying the effectiveness of policy gradient type methods (e.g., Agarwal et al. (2021); Bhandari (2020); Bhandari and Russo (2019, 2021); Cen et al. (2022a); Fazel et al. (2018); Khodadadian et al. (2022); Lan (2022a); Liu et al. (2020); Mei et al. (2021, 2020b); Shani et al. (2020); Yang et al. (2020); Zhan et al. (2021); Zhang et al. (2021a,b, 2020c)), focusing primarily on canonical settings such as tabular Markov decision processes (MDPs) for discrete-state problems and linear quadratic regulators for continuous-state problems.

The current paper studies PG methods with softmax parameterization — commonly referred to as *softmax policy gradient* methods — which are among the de facto implementations of PG methods in practice. An intriguing theoretical result was recently obtained by the work Agarwal et al. (2021), which established asymptotic global convergence of softmax PG methods for infinite-horizon γ -discounted tabular MDPs. Subsequently, Mei et al. (2020b) strengthened the theory by demonstrating that softmax PG methods are

capable of finding an ε -optimal policy with an iteration complexity proportional to $1/\varepsilon$ (see Table 1 for the precise form). While these results take an important step towards understanding the effectiveness of softmax PG methods, caution needs to be exercised before declaring fast convergence of the algorithms. In particular, the iteration complexity derived by Mei et al. (2020b) falls short of delineating clear dependencies on important salient parameters of the MDP, such as the dimension of the state space \mathcal{S} and the effective horizon $1/(1 - \gamma)$. These parameters are, more often than not, enormous in contemporary RL applications, and might play a pivotal role in determining the scalability of softmax PG methods.

Additionally, it is worth noting that existing literature largely concentrated on developing algorithm-dependent upper bounds on the iteration complexities. Nevertheless, we recommend caution when directly comparing computational upper bounds for distinct algorithms: the superiority of the computational upper bound for one algorithm does not necessarily imply this algorithm outperforms others, unless we can certify the tightness of all upper bounds being compared. As a more concrete example, it is of practical interest to benchmark softmax PG methods against natural policy gradient (NPG) methods with softmax parameterization, the latter of which is a variant of policy optimization lying underneath several mainstream RL algorithms such as *proximal policy optimization* (PPO) (Schulman et al., 2017) and *trust region policy optimization* (TRPO) (Schulman et al., 2015). While it is tempting to claim superiority of NPG methods over softmax PG methods — given the appealing convergence properties of NPG methods (Agarwal et al., 2021) (see Table 1) — existing theory fell short to reach such a conclusion, due to the absence of convergence lower bounds for softmax PG methods in prior literature.

The above considerations thus lead to a natural question that we aim to address in the present paper:

Can we develop a lower bound on the iteration complexity of softmax PG methods that reflects explicit dependency on salient parameters of the MDP of interest?

1.1 Main result

As an attempt to address the question posed above, our investigation delivers a somewhat surprising message that can be described in words as follows:

Softmax PG methods can take (super-)exponential time to converge, even in the presence of a benign initialization and an initial state distribution amenable to exploration.

Our finding, which is concerned with a discounted infinite-horizon tabular MDP, is formally stated in the following theorem. Here and throughout, $|\mathcal{S}|$ denotes the size of the state space \mathcal{S} , $0 < \gamma < 1$ stands for the discount factor, V^* indicates the optimal value function, $\eta > 0$ is the learning rate or stepsize, whereas $V^{(t)}$ represents the value function estimate of softmax PG methods in the t -th iteration. All immediate rewards are assumed to fall within $[-1, 1]$. See Section 2 for formal descriptions.

Theorem 1. *Assume that the softmax PG method adopts a uniform initial state distribution, a uniform policy initialization, and has access to exact gradient computation. Suppose that $0 < \eta < (1 - \gamma)^2/5$, then there exist universal constants $c_1, c_2, c_3 > 0$ such that: for any $0.96 < \gamma < 1$ and $|\mathcal{S}| \geq c_3(1 - \gamma)^{-6}$, one can find a γ -discounted MDP with state space \mathcal{S} that takes the softmax PG method at least*

$$\frac{c_1}{\eta} |\mathcal{S}|^{2\frac{c_2}{1-\gamma}} \text{ iterations} \tag{1}$$

to reach

$$\frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} [V^*(s) - V^{(t)}(s)] \leq 0.07. \tag{2}$$

Remark 1 (Action space). The MDP we construct contains at most three actions for each state.

Remark 2 (Stepsize range). Our lower bound operates under the assumption that $\eta < (1 - \gamma)^2/5$. In comparison, prior convergence guarantees for PG-type methods with softmax parameterization (e.g., Agarwal et al. (2021, Theorem 5.1) and Mei et al. (2020b, Theorem 6)) required $\eta < (1 - \gamma)^3/8$, a range of stepsizes fully covered by our theorem. In fact, prior works could only guarantee monotonicity of softmax PG methods (in terms of the value function) within the range $\eta < (1 - \gamma)^2/5$ (see Agarwal et al. (2021, Lemma C.2)).

algorithm	iteration complexity	reference
softmax PG upper bound	asymptotic	Agarwal et al. (2021, Thm. 5.1)
softmax PG upper bound	$O\left(\mathcal{C}_{\text{spg}}^2(\mathcal{M}) \left\ \frac{d_\mu^\pi}{\mu} \right\ _\infty^2 \left\ \frac{1}{\mu} \right\ _\infty \frac{ \mathcal{S} }{(1-\gamma)^6 \varepsilon}\right)$	Mei et al. (2020b, Thm. 4)
softmax NPG upper bound	$O\left(\frac{1}{(1-\gamma)^2 \varepsilon}\right)$	Agarwal et al. (2021, Thm. 5.3)
softmax PG lower bound	$\frac{(1-\gamma)^5 \Delta_\star^2}{12\varepsilon}$	Mei et al. (2020b, Thm. 10)
softmax PG lower bound	$ \mathcal{S} ^{2^{\Omega(\frac{1}{1-\gamma})}}$	this work

Table 1: Upper and lower bounds on the iteration complexities of PG and NPG methods with softmax parameterization in finding an ε -optimal policy obeying $\|V^\star - V^{(t)}\|_\infty \leq \varepsilon \leq 0.15$. We assume exact gradient evaluation, and hide the dependencies that are logarithmic on problem parameters. Here, μ denotes the initial state distribution, $\|d_\mu^\pi/\mu\|_\infty$ is the distribution mismatch coefficient, $a^\star(s)$ is the optimal action at state s according to π^\star , $\mathcal{C}_{\text{spg}}(\mathcal{M}) := (\inf_{s \in \mathcal{S}} \inf_{t \geq 1} \pi^{(t)}(a^\star(s) | s))^{-1}$ is a quantity depending on both the PG trajectory and salient MDP parameters, whereas $\Delta_\star := \min_{s \in \mathcal{S}, a \neq a^\star(s)} \{Q^\star(s, a^\star(s)) - Q^\star(s, a)\}$ is the optimality gap w.r.t. the optimal Q-function Q^\star .

Remark 3. While we can also provide explicit numbers for the constants $c_1, c_2, c_3 > 0$, these numbers are not informative, and hence we omit explicit numbers here to streamline the proof a bit.

For simplicity of presentation, Theorem 2 is stated for the long-effective-horizon regime where $\gamma > 0.96$; it continues to hold when $\gamma > c_0$ for some smaller constant $c_0 > 0$. Our result is obtained by exhibiting a hard MDP instance — which is a properly augmented chain-like MDP — for which softmax PG methods converge extremely slowly even when perfect model specification is available. Several remarks and implications of our result are in order.

Comparisons with prior results. Table 1 provides an extensive comparison of the iteration complexities — including both upper and lower bounds — of PG and NPG methods under softmax parameterization. As suggested by our result, the iteration complexity $O(\mathcal{C}_{\text{spg}}^2(\mathcal{M}) \frac{1}{\varepsilon})$ derived in Mei et al. (2020b) (see Table 1) might not be as rosy as it seems for problems with large state space and long effective horizon; in fact, the crucial quantity $\mathcal{C}_{\text{spg}}(\mathcal{M})$ therein could scale in a prohibitive manner with both $|\mathcal{S}|$ and $\frac{1}{1-\gamma}$. Mei et al. (2020b) also developed a lower bound on the iteration complexity of softmax PG methods, which falls short of capturing the influence of the state space dimension and might become smaller than 1 unless ε is very small (e.g., $\varepsilon \lesssim (1-\gamma)^3$) for problems with long effective horizons. In addition, Mei et al. (2020a) provided some interesting evidence that a poorly-initialized softmax PG algorithm can get stuck at suboptimal policies for a single-state MDP (i.e., the bandit problem). This result, however, fell short of providing a complete runtime analysis and did not look into the influence of a large state space. By contrast, our theory reveals that softmax PG methods can take exponential time to reach even a moderate accuracy level.

Slow convergence even with benign distribution mismatch. Existing computational complexities for policy gradient type methods (e.g., Agarwal et al. (2021); Mei et al. (2020b)) typically scale polynomially in the so-called *distribution mismatch coefficient*¹ $\| \frac{d_\mu^\pi}{\mu} \|_\infty$, where d_μ^π stands for a certain discounted state visitation distribution (see (13) in Section 2), and μ denotes the distribution over initial states. It is thus natural to wonder whether the exponential lower bound in Theorem 1 is a consequence of an exceedingly large distribution mismatch coefficient. This, however, is not the case; in fact, our theory chooses μ to be a benign uniform distribution so that $\| \frac{d_\mu^\pi}{\mu} \|_\infty \leq \| \frac{1}{\mu} \|_\infty \leq |\mathcal{S}|$, which scales at most linearly in $|\mathcal{S}|$.

¹Here and throughout, the division of two vectors represents componentwise division.

Benchmarking with softmax NPG methods. Our algorithm-specific lower bound suggests that softmax PG methods — in their vanilla form — might take a prohibitively long time to converge when the state space and effective horizon are large. This is in stark contrast to the convergence rate of NPG type methods, whose iteration complexity is dimension-free and scales only polynomially with the effective horizon (Agarwal et al., 2021; Cen et al., 2022a). Consequently, our results shed light on the practical superiority of NPG-based algorithms such as PPO (Schulman et al., 2017) and TRPO (Schulman et al., 2015).

Crux of our design. As we shall elucidate momentarily in Section 3, our exponential lower bound is obtained through analyzing the trajectory of softmax PG methods on a carefully-designed MDP instance with no more than 3 actions per state, when a uniform initialization scheme and a uniform initial state distribution are adopted. Our construction underscores the critical challenge of *credit assignments* (Sutton, 1984) in RL compounded by the presence of delayed rewards, long horizon, and intertwined interactions across states. While it is difficult to elucidate the source of exponential lower bound without presenting our MDP construction, we take a moment to point out some critical properties that underlie our designs. To be specific, we seek to design a chain-like MDP containing $H = O(\frac{1}{1-\gamma})$ key primary states $\{1, \dots, H\}$ (each coupled with many auxiliary states), for which the softmax PG method satisfies the following properties.

- For the two key primary states, we have

$$\min \{ \text{convergence-time}(\text{state } 1), \text{convergence-time}(\text{state } 2) \} \geq \frac{|\mathcal{S}|}{\eta}. \quad (3)$$

- (*A blowing-up phenomenon*) For each key primary state $3 \leq s \leq H = O(\frac{1}{1-\gamma})$, one has

$$\text{convergence-time}(\text{state } s) \gtrsim (\text{convergence-time}(\text{state } s-2))^{1.5}, \quad 3 \leq s \leq H. \quad (4)$$

Here, it is understood that $\text{convergence-time}(\text{state } s)$ represents informally the time taken for the value function of state s to be sufficiently close to its optimal value. The blowing-up phenomenon described above is precisely the source of our (super)-exponential lower bound.

1.2 Other related works

Non-asymptotic analysis of (natural) policy gradient methods. Moving beyond tabular MDPs, finite-time convergence guarantees of PG / NPG methods and their variants have recently been studied for control problems (e.g., Fazel et al. (2018); Jansch-Porto et al. (2020); Tu and Recht (2019); Zhang et al. (2020b)), regularized MDPs (e.g., Cen et al. (2022a); Lan (2022a); Zhan et al. (2021)), constrained MDPs (e.g., Ding et al. (2020); Xu et al. (2020a)), robust MDPs (e.g., Li et al. (2022b); Zhang et al. (2021c)), MDPs with function approximation (e.g., Agarwal et al. (2021); Agazzi and Lu (2021); Cai et al. (2020); Lan (2022b); Liu et al. (2019); Wang et al. (2019)), Markov games (e.g., Cen et al. (2021); Daskalakis et al. (2020); Wei et al. (2021); Xie et al. (2020); Zhao et al. (2021)), and their use in actor-critic methods (e.g., Alacaoglu et al. (2022); Cen et al. (2022b); Wu et al. (2020); Xu et al. (2020b)).

Other policy parameterizations. In addition to softmax parameterization, several other policy parameterization schemes have also been investigated in the context of policy optimization and reinforcement learning at large. For example, Agarwal et al. (2021); Lan (2022a); Zhan et al. (2021); Zhang et al. (2020a) studied the convergence of projected PG methods and policy mirror descent with direct parameterization, Asadi and Littman (2017) introduced the so-called mallow parameterization, while Mei et al. (2020a) studied the escort parameterization. Part of these parameterizations were proposed in response to the ineffectiveness of softmax parameterization observed in practice.

Lower bounds. Establishing information-theoretic or algorithmic-specific lower bounds on the statistical and computational complexities of RL algorithms — often achieved by constructing hard MDP instances — plays an instrumental role in understanding the bottlenecks of RL algorithms. To give a few examples, Azar et al. (2013); Domingues et al. (2021); Li et al. (2022a); Yan et al. (2022) developed information-theoretic

lower bounds on the sample complexity of RL under multiple sampling mechanisms (e.g., sampling with a generative model, online RL, and offline/batch RL), [Li et al. \(2021\)](#) established an algorithm-dependent lower bound on the sample complexity of Q-learning, whereas [Khamaru et al. \(2021\)](#); [Pananjady and Wainwright \(2020\)](#) developed instance-dependent lower bounds for policy evaluation. Additionally, [Agarwal et al. \(2021\)](#) constructed a chain-like MDP whose value function under direct parameterization might contain very flat saddle points under a certain initial state distribution, highlighting the role of distribution mismatch coefficients in policy optimization. Finally, exponential-time convergence of gradient descent has been observed in other nonconvex problems as well (e.g., [Du et al. \(2017\)](#)) despite its asymptotic convergence ([Lee et al., 2016](#)), although the context and analysis therein are drastically different from what happens in RL settings.

1.3 Paper organization

The rest of this paper is organized as follows. In Section 2, we introduce the basics of Markov decision processes, and describe the softmax policy gradient method along with several key functions/quantities. Section 3 constructs a chain-like MDP, which is the hard MDP instance underlying our computational lower bound for PG methods. In Section 4, we outline the proof of Theorem 1, starting with the proof of a weaker version before establishing Theorem 1. The proof of all technical lemmas are deferred to the appendix. We conclude the paper in Section 5 with a summary of our findings.

2 Background

In this section, we introduce the basics of MDPs, and formally describe the softmax PG method. Here and throughout, we denote by $\Delta(\mathcal{X})$ the probability simplex over a set \mathcal{X} , and let $|\mathcal{X}|$ represent the cardinality of the set \mathcal{X} . Given two probability distributions p and q over \mathcal{S} , we adopt the notation $\| \frac{p}{q} \|_\infty = \max_{s \in \mathcal{S}} \frac{p(s)}{q(s)}$ and $\| \frac{1}{q} \|_\infty = \max_{s \in \mathcal{S}} \frac{1}{q(s)}$. Throughout this paper, the notation $f(\mathcal{M}) \gtrsim g(\mathcal{M})$ (resp. $f(\mathcal{M}) \lesssim g(\mathcal{M})$) means there exist some universal constants $c > 0$ independent of the parameters of the MDP \mathcal{M} such that $f(\mathcal{M}) \geq cg(\mathcal{M})$ (resp. $f(\mathcal{M}) \leq cg(\mathcal{M})$), while the notation $f(\mathcal{M}) \asymp g(\mathcal{M})$ means that $f(\mathcal{M}) \gtrsim g(\mathcal{M})$ and $f(\mathcal{M}) \lesssim g(\mathcal{M})$ hold simultaneously.

Infinite-horizon discounted MDP. Let $\mathcal{M} = (\mathcal{S}, \{\mathcal{A}_s\}_{s \in \mathcal{S}}, P, r, \gamma)$ be an infinite-horizon discounted MDP. Here, \mathcal{S} represents the state space, \mathcal{A}_s denotes the action space associated with state $s \in \mathcal{S}$, $\gamma \in (0, 1)$ indicates the discount factor, P is the probability transition kernel (namely, for each state-action pair (s, a) , $P(\cdot | s, a) \in \Delta(\mathcal{S})$ denotes the transition probability from state s to the next state when action a is taken), and r stands for a deterministic reward function (namely, $r(s, a)$ is the immediate reward received in state s upon executing action a). Throughout this paper, we assume normalized rewards such that $-1 \leq r(s, a) \leq 1$ for any state-action pair (s, a) . In addition, we concentrate on the scenario where γ is quite close to 1, and often refer to $\frac{1}{1-\gamma}$ as the effective horizon of the MDP.

Policy, value function, Q-function and advantage function. The agent operates by adopting a policy π , which is a (randomized) action selection rule based solely on the current state of the MDP. More precisely, for any state $s \in \mathcal{S}$, we use $\pi(\cdot | s) \in \Delta(\mathcal{A}_s)$ to specify a probability distribution, with $\pi(a | s)$ denoting the probability of executing action $a \in \mathcal{A}_s$ when in state s . The value function $V^\pi : \mathcal{S} \rightarrow \mathbb{R}$ of a policy π — which indicates the expected discounted cumulative reward induced by policy π — is defined as

$$\forall s \in \mathcal{S} : \quad V^\pi(s) := \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r(s^k, a^k) \mid s^0 = s \right]. \quad (5)$$

Here, the expectation is taken over the randomness of the MDP trajectory $\{(s^k, a^k)\}_{k \geq 0}$ and the policy, where $s^0 = s$ and, for all $k \geq 0$, $a^k \sim \pi(\cdot | s^k)$ follows the policy π and $s^{k+1} \sim P(\cdot | s^k, a^k)$ is generated by the transition kernel P . Analogously, we shall also define the value function $V^\pi(\mu)$ of a policy π when the initial state is drawn from a distribution μ over \mathcal{S} , namely,

$$V^\pi(\mu) := \mathbb{E}_{s \sim \mu} [V^\pi(s)]. \quad (6)$$

Additionally, the Q-function Q^π of a policy π — namely, the expected discounted cumulative reward under policy π given an initial state-action pair $(s^0, a^0) = (s, a)$ — is formally defined by

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}: \quad Q^\pi(s, a) := \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r(s^k, a^k) \mid s^0 = s, a^0 = a \right], \quad (7)$$

where the expectation is again over the randomness of the MDP trajectory $\{(s^k, a^k)\}_{k \geq 1}$ when policy π is adopted. In addition, the advantage function of policy π is defined as

$$A^\pi(s, a) := Q^\pi(s, a) - V^\pi(s) \quad (8)$$

for every state-action pair (s, a) .

A major goal is to find a policy that optimizes the value function and the Q-function. Throughout this paper, we denote respectively by V^* and Q^* the optimal value function and optimal Q-function, namely,

$$V^*(s) := \max_{\pi} V^\pi(s), \quad Q^*(s, a) := \max_{\pi} Q^\pi(s, a), \quad \text{for all } s \in \mathcal{S} \text{ and } a \in \mathcal{A}_s. \quad (9)$$

Softmax parameterization and policy gradient methods. The family of policy optimization algorithms attempts to identify optimal policies by resorting to optimization-based algorithms. To facilitate differentiable optimization, a widely adopted scheme is to parameterize policies using softmax mappings. Specifically, for any real-valued parameter $\theta = [\theta(s, a)]_{s \in \mathcal{S}, a \in \mathcal{A}_s}$, the corresponding softmax policy $\pi_\theta := \text{softmax}(\theta)$ is defined such that

$$\forall s \in \mathcal{S} \text{ and } a \in \mathcal{A}_s: \quad \pi_\theta(a \mid s) := \frac{\exp(\theta(s, a))}{\sum_{a' \in \mathcal{A}_s} \exp(\theta(s, a'))}. \quad (10)$$

With the aim of maximizing the value function under softmax parameterization, namely,

$$\text{maximize}_{\theta} \quad V^{\pi_\theta}(\mu), \quad (11)$$

the softmax PG method proceeds by adopting gradient ascent update rules w.r.t. θ :

$$\theta^{(t+1)} = \theta^{(t)} + \eta \nabla_{\theta} V^{\pi_\theta}(\mu), \quad t = 0, 1, \dots \quad (12a)$$

Here and throughout, we let $V^{(t)} = V^{\pi^{(t)}}$ and $Q^{(t)} = Q^{\pi^{(t)}}$ abbreviate respectively the value function and Q-function of the policy iterate $\pi^{(t)} := \pi_{\theta^{(t)}}$ in the t -th iteration, and $\eta > 0$ denotes the stepsize or learning rate. Interestingly, the gradient $\nabla_{\theta} V^{\pi_\theta}$ under softmax parameterization (10) admits a closed-form expression (Agarwal et al., 2021), that is, for any state-action pair (s, a) ,

$$\frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta(s, a)} = \frac{1}{1 - \gamma} d_{\mu}^{\pi_\theta}(s) \pi_\theta(a \mid s) A^{\pi_\theta}(s, a). \quad (12b)$$

Here, $d_{\mu}^{\pi_\theta}(s)$ represents the *discounted state visitation distribution* of a policy π given the initial state $s^0 \sim \mu$:

$$\forall s \in \mathcal{S}: \quad d_{\mu}^{\pi}(s) := (1 - \gamma) \mathbb{E}_{s^0 \sim \mu} \left[\sum_{k=0}^{\infty} \gamma^k \mathbb{P}(s^k = s \mid s^0) \right], \quad (13)$$

with the expectation taken over the randomness of the MDP trajectory $\{(s^k, a^k)\}_{k \geq 0}$ under the policy π and the initial state distribution μ . In words, $d_{\mu}^{\pi}(s)$ measures — starting from an initial distribution μ — how frequently state s will be visited in a properly discounted fashion. Throughout this paper, we shall denote $A^{(t)} := A^{\pi^{(t)}}$ and $d_{\mu}^{(t)}(s) := d_{\mu}^{\pi^{(t)}}(s)$ for notation simplicity.

3 Construction of a hard MDP

This section constructs a discounted infinite-horizon MDP $\mathcal{M} = \{\mathcal{S}, \{\mathcal{A}_s\}_{s \in \mathcal{S}}, r, P, \gamma\}$, as depicted in Fig. 1, which forms the basis of the exponential lower bound claimed in this paper. In addition to the basic notation already introduced in Section 2, we remark on the action space as follows.

- For each state $s \in \mathcal{S}$, we have $\mathcal{A}_s \subseteq \{a_0, a_1, a_2\}$. For convenience of presentation, we allow the action space to vary with $s \in \mathcal{S}$, but it always comprises no more than 3 actions.




Figure 1: An illustration of the constructed MDP.

State space partitioning. The states of our MDP exhibit certain group structure. To be precise, we partition the state space \mathcal{S} into a few *disjoint* subsets

$$\mathcal{S} = \{0\} \cup \mathcal{S}_{\text{primary}} \cup \mathcal{S}_{\text{adj}} \cup \mathcal{S}_1 \cup \mathcal{S}_2 \cup \widehat{\mathcal{S}}_1 \cup \dots \cup \widehat{\mathcal{S}}_H \cup \widehat{\mathcal{S}}_{\bar{1}} \cup \dots \cup \widehat{\mathcal{S}}_{\bar{H}}, \quad (14)$$

which entails:

- state 0 (an absorbing state);
- two key “buffer” state subsets \mathcal{S}_1 and \mathcal{S}_2 ;
- a set of $H - 2$ key primary states $\mathcal{S}_{\text{primary}} := \{3, \dots, H\}$;²
- a set of H key adjoint states $\mathcal{S}_{\text{adj}} := \{\bar{1}, \bar{2}, \dots, \bar{H}\}$;
- $2H$ “booster” state subsets $\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_H, \widehat{\mathcal{S}}_{\bar{1}}, \dots, \widehat{\mathcal{S}}_{\bar{H}}$.

Remark 4. Our subsequent analysis largely concentrates on the subsets \mathcal{S}_1 , \mathcal{S}_2 , $\mathcal{S}_{\text{primary}}$ and \mathcal{S}_{adj} . In particular, each state $s \in \{3, \dots, H\}$ is paired with what we call an adjoint state \bar{s} , whose role will be elucidated shortly. In addition, state $\bar{1}$ (resp. state $\bar{2}$) can be viewed as the adjoint state of the set \mathcal{S}_1 (resp. \mathcal{S}_2). The sets $\mathcal{S}_{\text{primary}}$ and \mathcal{S}_{adj} comprise a total number of $2H - 2$ states; in comparison, \mathcal{S}_1 and \mathcal{S}_2 are chosen to be much larger and contain a number of replicated states, a crucial design component that helps ensure the property (3) under a uniform initial state distribution. As we shall make clear momentarily, the “booster” state sets are introduced mainly to help boost the discounted visitation distribution of the states in \mathcal{S}_1 , \mathcal{S}_2 , $\mathcal{S}_{\text{primary}}$, and \mathcal{S}_{adj} at the initial stage.

We shall also specify below the size of these state subsets as well as some key parameters, where the choices of the quantities $c_h, c_{b,1}, c_{b,2}, c_m \asymp 1$ will be made clear in the analysis (cf. (35)).

- H is taken to be on the same order as the “effective horizon” of this discounted MDP, namely,

$$H = \frac{c_h}{1 - \gamma}. \quad (15)$$

- The two buffer state subsets \mathcal{S}_1 and \mathcal{S}_2 have size

$$|\mathcal{S}_1| = c_{b,1}(1 - \gamma)|\mathcal{S}| \quad \text{and} \quad |\mathcal{S}_2| = c_{b,2}(1 - \gamma)|\mathcal{S}|. \quad (16)$$

- The booster state sets are of the same size, namely,

$$|\widehat{\mathcal{S}}_1| = \dots |\widehat{\mathcal{S}}_H| = |\widehat{\mathcal{S}}_{\bar{1}}| = \dots = |\widehat{\mathcal{S}}_{\bar{H}}| = c_m(1 - \gamma)|\mathcal{S}|. \quad (17)$$

²While we do not include states 1 and 2 here, any state in \mathcal{S}_1 (resp. \mathcal{S}_2) can essentially be viewed as a (replicated) copy of state 1 (resp. state 2).

Probability transition kernel and reward function. We now describe the probability transition kernel and the reward function for each state subset. Before continuing, we find it helpful to isolate a few key parameters that will be used frequently in our construction:

$$\tau_s := 0.5\gamma^{\frac{2s}{3}}, \quad (18a)$$

$$p := c_p(1 - \gamma), \quad (18b)$$

$$r_s := 0.5\gamma^{\frac{2s}{3} + \frac{5}{6}}, \quad (18c)$$

where $s \in \{1, 2, \dots, H\}$, and $c_p > 0$ is some small constant that shall be specified later (see (35)). To facilitate understanding, we shall often treat τ_s and r_s ($s \leq H$) as quantities that are all fairly close to 0.5 (which would happen if γ is close to 1 and $H = \frac{c_h}{1-\gamma}$ for c_h sufficiently small).

We are now positioned to make precise descriptions of both P and r as follows.

- *Absorbing state 0:* singleton action space $\{a_0\}$,

$$P(0 | 0, a_0) = 1, \quad r(0, a_0) = 0. \quad (19)$$

This is an absorbing state, namely, the MDP will stay in this state permanently once entered. As we shall see below, taking action a_0 in an arbitrary state will enter state 0 immediately.

- *Key primary states $s \in \{3, \dots, H\}$:* action space $\{a_0, a_1, a_2\}$,

$$P(0 | s, a_0) = 1, \quad r(s, a_0) = r_s + \gamma^2 p \tau_{s-2}, \quad (20a)$$

$$P(\overline{s-1} | s, a_1) = 1, \quad r(s, a_1) = 0, \quad (20b)$$

$$P(0 | s, a_2) = 1 - p, \quad r(s, a_2) = r_s, \quad (20c)$$

$$P(\overline{s-2} | s, a_2) = p, \quad (20d)$$

where p , τ_s and r_s are all defined in (18).

- *Key adjoint states $\bar{s} \in \{\bar{3}, \dots, \bar{H}\}$:* action space $\{a_0, a_1\}$,

$$P(0 | \bar{s}, a_0) = 1, \quad r(\bar{s}, a_0) = \gamma \tau_s, \quad (21a)$$

$$P(s | \bar{s}, a_1) = 1, \quad r(\bar{s}, a_1) = 0, \quad (21b)$$

where τ_s is defined in (18a).

- *Key buffer state subsets \mathcal{S}_1 and \mathcal{S}_2 :* action space $\{a_0, a_1\}$,

$$\forall s_1 \in \mathcal{S}_1 : \quad P(0 | s_1, a_0) = 1, \quad r(s_1, a_0) = -\gamma^2, \quad (22a)$$

$$P(0 | s_1, a_1) = 1, \quad r(s_1, a_1) = \gamma^2, \quad (22b)$$

$$\forall s_2 \in \mathcal{S}_2 : \quad P(0 | s_2, a_0) = 1, \quad r(s_2, a_0) = -\gamma^4, \quad (22c)$$

$$P(0 | s_2, a_1) = 1, \quad r(s_2, a_1) = \gamma^4. \quad (22d)$$

Given the homogeneity of the states in \mathcal{S}_1 (resp. \mathcal{S}_2), we shall often use the shorthand notation $P(\cdot | 1, a)$ (resp. $P(\cdot | 2, a)$) to abbreviate $P(\cdot | s_1, a)$ (resp. $P(\cdot | s_2, a)$) for any $s_1 \in \mathcal{S}_1$ (resp. $s_2 \in \mathcal{S}_2$) for the sake of convenience.

- *Other adjoint states $\bar{1}$ and $\bar{2}$:* action space $\{a_0, a_1\}$,

$$P(0 | \bar{1}, a_0) = 1, \quad r(\bar{1}, a_0) = \gamma \tau_1, \quad P(s_1 | \bar{1}, a_1) = \frac{1}{|\mathcal{S}_1|}, \quad \forall s_1 \in \mathcal{S}_1, \quad r(\bar{1}, a_1) = 0, \quad (23a)$$

$$P(0 | \bar{2}, a_0) = 1, \quad r(\bar{2}, a_0) = \gamma \tau_2, \quad P(s_2 | \bar{2}, a_1) = \frac{1}{|\mathcal{S}_2|}, \quad \forall s_2 \in \mathcal{S}_2, \quad r(\bar{2}, a_1) = 0, \quad (23b)$$

where τ_1 and τ_2 are defined in (18a).

- *Booster state subsets* $\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_H, \widehat{\mathcal{S}}_{\overline{1}}, \dots, \widehat{\mathcal{S}}_{\overline{H}}$: singleton action space $\{a_1\}$,

$$\forall s' \in \widehat{\mathcal{S}}_1, s \in \mathcal{S}_1 : \quad P(s | s', a_1) = 1/|\mathcal{S}_1|, \quad (24a)$$

$$\forall s' \in \widehat{\mathcal{S}}_2, s \in \mathcal{S}_2 : \quad P(s | s', a_1) = 1/|\mathcal{S}_2|; \quad (24b)$$

for any $s \in \{3, \dots, H\}$,

$$\forall s' \in \widehat{\mathcal{S}}_s, : \quad P(s | s', a_1) = 1, \quad (24c)$$

and for any $\bar{s} \in \{\overline{1}, \dots, \overline{H}\}$,

$$\forall s' \in \widehat{\mathcal{S}}_{\bar{s}}, : \quad P(\bar{s} | s', a_1) = 1. \quad (24d)$$

The rewards in all these cases are set to be 0 (in fact, they will not even appear in the analysis). In addition, any transition probability that has not been specified is equal to zero.

Convenient notation for buffer state subsets \mathcal{S}_1 and \mathcal{S}_2 . By construction, it is easily seen that the states in \mathcal{S}_1 (resp. \mathcal{S}_2) have identical characteristics; in fact, all states in \mathcal{S}_1 (resp. \mathcal{S}_2) share exactly the same value functions and Q-functions throughout the execution of the softmax PG method. As a result, we introduce the following convenient notation whenever it is clear from the context:

$$V^\pi(s_1) =: V^\pi(1), \quad Q^\pi(s_1, a) =: Q^\pi(1, a), \quad A^\pi(s_1) =: A^\pi(1) \quad \text{for all } s_1 \in \mathcal{S}_1; \quad (25a)$$

$$V^\pi(s_2) =: V^\pi(2), \quad Q^\pi(s_2, a) =: Q^\pi(2, a), \quad A^\pi(s_2) =: A^\pi(2) \quad \text{for all } s_2 \in \mathcal{S}_2; \quad (25b)$$

$$d_\mu^\pi(s_1) =: d_\mu^\pi(1), \quad \pi(a | s_1) =: \pi(a | 1), \quad \theta(s_1, a) =: \theta(1, a) \quad \text{for all } s_1 \in \mathcal{S}_1; \quad (25c)$$

$$d_\mu^\pi(s_2) =: d_\mu^\pi(2), \quad \pi(a | s_2) =: \pi(a | 2), \quad \theta(s_2, a) =: \theta(2, a) \quad \text{for all } s_2 \in \mathcal{S}_2. \quad (25d)$$

Optimal values and optimal actions of the constructed MDP. Before concluding this section, we find it convenient to determine the optimal value functions and the optimal actions of the constructed MDP, which would be particularly instrumental when presenting our analysis. This is summarized in the lemma below, whose proof can be found in Appendix A.3.

Lemma 1. *Suppose that $\gamma^{2H} \geq 2/3$ and $H \geq 2$. Then one has*

$$V^*(0) = 0, \quad V^*(s) = Q^*(s, a_1) = \gamma^{2s}, \quad 1 \leq s \leq H, \quad (26a)$$

$$V^*(\bar{s}) = Q^*(\bar{s}, a_1) = \gamma^{2s+1}, \quad 1 \leq s \leq H, \quad (26b)$$

and the optimal policy is to take action a_1 in all non-absorbing states. In addition, for any policy π and any state-action pair (s, a) , one has $Q^\pi(s, a) \geq -\gamma^2$.

Lemma 1 tells us that for this MDP, the optimal policy for all non-absorbing states takes a simple form: sticking to action a_1 . In particular, when $\gamma \approx 1$ and $\gamma^H \approx 1$, Lemma 1 reveals that the optimal values of all non-absorbing major states are fairly close to 1, namely,

$$V^*(s) \approx 1 \quad \text{for all } s \in \{1, \dots, H\} \cup \{\overline{1}, \dots, \overline{H}\}. \quad (27)$$

Additionally, the above lemma directly implies that the Q-function (and hence the value function) is always bounded below by -1 , a property that will be used several times in our analysis.

4 Analysis: proof outline

In this section, we present the main steps for establishing our computational lower bound in Theorem 1. Before doing so, we find it convenient to start by presenting and proving a weaker version as follows.

Theorem 2. Consider the MDP \mathcal{M} constructed in Section 3 (and Fig. 1). Assume that the softmax PG method adopts a uniform initial state distribution, a uniform policy initialization, and has access to exact gradient computation. Suppose that $0 < \eta < (1 - \gamma)^2/5$. There exist universal constants $c_1, c_2, c_3 > 0$ such that: for any $0.96 < \gamma < 1$ and $|\mathcal{S}| \geq c_3(1 - \gamma)^{-6}$, one has

$$V^*(s) - V^{(t)}(s) > 0.15, \quad \text{for all primary states } 0.1H < s < H, \quad (28)$$

provided that the iteration number satisfies

$$t < \frac{c_1}{\eta} |\mathcal{S}|^{2\frac{c_2}{1-\gamma}}. \quad (29)$$

In what follows, we shall concentrate on establishing Theorem 2, on the basis of the MDP instance constructed in Section 3. Once this theorem is established, we shall revisit Theorem 1 (towards the end of Section 4.3) and describe how the proof of Theorem 2 can be adapted to prove Theorem 1.

4.1 Preparation: crossing times and choice of constants

Crossing times. To investigate how long it takes for softmax PG methods to converge to the optimal policy, we shall pay particular attention to a family of key quantities: the number of iterations needed for $V^{(t)}(s)$ to surpass a prescribed threshold τ ($\tau < 1$) *before* it reaches its optimal value. To be precise, for each $s \in \{3, \dots, H\} \cup \{\bar{1}, \dots, \bar{H}\}$ and any given threshold $\tau > 0$, we introduce the following *crossing time*:

$$t_s(\tau) := \arg \min \{t \mid V^{(t)}(s) \geq \tau\}. \quad (30)$$

When it comes to the buffer state subsets \mathcal{S}_1 and \mathcal{S}_2 , we define the crossing times analogously as follows

$$t_1(\tau) := \arg \min \{t \mid V^{(t)}(1) \geq \tau\} \quad \text{and} \quad t_2(\tau) := \arg \min \{t \mid V^{(t)}(2) \geq \tau\}, \quad (31)$$

where we recall the notation $V^{(t)}(1)$ and $V^{(t)}(2)$ introduced in (25).

Monotonicity of crossing times. Recalling the definition (30) of the crossing time $t_s(\cdot)$, we know that

$$V^{(t)}(s) < \tau_s \quad \text{for all } t < t_s(\tau_s), \quad (32)$$

with τ_s defined in expression (18a). We immediately make note of the following crucial monotonicity property that will be justified later in Remark 8:

$$t_2(\tau_2) \leq t_3(\tau_3) \leq \dots \leq t_H(\tau_H). \quad (33)$$

It will also be shown in Lemma 4 that $t_1(\tau_1) \leq t_2(\tau_2)$ when the constants $c_{b,1}, c_{b,2}$ and c_m are properly chosen.

Remark 5. As we shall see shortly (i.e., Part (iii) of Lemma 8), one has $t_{\bar{s}}(\gamma\tau_s) = t_s(\tau_s)$ for any $\bar{s} \in \{\bar{1}, \dots, \bar{H}\}$, which combined with (33) leads to

$$t_1(\tau_1) = t_{\bar{1}}(\gamma\tau_1) \leq t_2(\tau_2) = t_{\bar{1}}(\gamma\tau_1) \leq t_3(\tau_3) = t_{\bar{3}}(\gamma\tau_3) \leq \dots \leq t_H(\tau_H) = t_{\bar{H}}(\gamma\tau_H). \quad (34)$$

Choice of parameters. We assume the following choice of parameters throughout the proof:

$$\gamma > 0.96, \quad c_m < 1, \quad c_h < 0.19, \quad \eta < \frac{(1 - \gamma)^2}{5}, \quad \frac{c_{b,1}}{c_m} \leq \frac{1}{79776}, \quad 8 \leq \frac{c_{b,2}}{c_m} \leq 15, \quad c_p < \frac{1}{2016}. \quad (35)$$

In the sequel, we outline the key steps that underlie the proof of our main results, with the proofs of the key lemmas postponed to the appendix.

4.2 A high-level picture

While our proof is highly technical, it is prudent to point out some key features that help paint a high-level picture about the slow convergence of the algorithm. Recall that a_1 is the optimal action in the constructed MDP. The chain-like structure of our MDP underscores a sort of *sequential dependency*: the dynamic of any primary state $s \in \{3, \dots, H\}$ depends heavily on what happens in those states prior to s — particularly state $s-1$, state $s-2$ as well as the associated adjoint states. By carefully designing the immediate rewards, we can ensure that for any $s \in \{3, \dots, H\}$, the iterate $\pi^{(t)}(a_1 | s)$ corresponding to the optimal action a_1 keeps decreasing before $\pi^{(t)}(a_1 | s-2)$ gets reasonably close to 1. As illustrated in Figure 2, this feature implies that the time taken for $\pi^{(t)}(a_1 | s)$ to get close to 1 grows (at least) geometrically as s increases, as will be formalized in (46).




Figure 2: An illustration of the dynamics of $\pi^{(t)}(a_1 | s)$ vs. the iteration count t . The yellow line, the middle red line, and the dark red line illustrate the dynamics of $\pi^{(t)}(a_1 | 1)$, $\pi^{(t)}(a_1 | 3)$ and $\pi^{(t)}(a_1 | 5)$, respectively.

Furthermore, we summarize below the typical dynamics of the iterates $\theta^{(t)}(s, a)$ before they converge, which are helpful for the reader to understand the proof. We start with the key buffer state sets \mathcal{S}_1 and \mathcal{S}_2 , which are the easiest to describe.

Dynamics of $\theta^{(t)}(s, a)$ (for key buffer state sets \mathcal{S}_1 and \mathcal{S}_2):

1. Initialization: $\theta^{(0)}(1, a_0) = \theta^{(0)}(1, a_1) = 0$ and $\theta^{(0)}(2, a_0) = \theta^{(0)}(2, a_1) = 0$
2. All iterations (Lemma 4):
 - $\theta^{(t)}(1, a_1)$ and $\theta^{(t)}(2, a_1)$ keep increasing and remains the largest
 - $\theta^{(t)}(1, a_0)$ and $\theta^{(t)}(2, a_0)$ keep decreasing and remains the smallest

Next, the dynamics of $\theta^{(t)}(s, a)$ for the key primary states $3 \leq s \leq H$ are much more complicated, and rely heavily on the status of several prior states $s-1$, $s-2$ and $s-3$. This motivates us to divide the dynamics into several stages based on the crossing times of these prior states, which are illustrated in Figure 3 as well. Here, we remind the reader of the definition of τ_s in (18).




Figure 3: An illustration of the dynamics of $\{\theta^{(t)}(s, a)\}_{a \in \{a_0, a_1, a_2\}}$ vs. iteration number t . The blue, red and yellow lines represent the dynamics of $\theta^{(t)}(s, a_0)$, $\theta^{(t)}(s, a_1)$ and $\theta^{(t)}(s, a_2)$, respectively. Here, we use solid lines to emphasize the time ranges for which the dynamics of $\theta^{(t)}(s, a)$ play the most crucial roles in our lower bound analysis.

Dynamics of $\theta^{(t)}(s, a)$ (for key primary states $3 \leq s \leq H$):

1. Initialization: $\theta^{(0)}(s, a_0) = \theta^{(0)}(s, a_1) = \theta^{(0)}(s, a_2) = 0$
2. Initial stage: $t < t_{s-2}(\tau_{s-2})$ (Lemma 5)
 - $\theta^{(t)}(s, a_1)$ keeps decreasing and remains the smallest
 - $\theta^{(t)}(s, a_0)$ keeps increasing and remains the largest
 - $\theta^{(t)}(s, a_2)$ keeps increasing
3. Intermediate stage: $t_{s-2}(\tau_{s-2}) \leq t \leq t_{s-1}(\tau_s)$ (Lemma 6)
 - $\theta^{(t)}(s, a_1)$ keeps decreasing and remains the smallest
 - $\theta^{(t)}(s, a_2)$ keeps increasing
4. Final stage (part 1): $t_{s-1}(\tau_s) < t < t_{\text{ref}}$ (Lemma 7)
 - $\theta^{(t)}(s, a_1)$ increases a little
 - $\theta^{(t)}(s, a_0)$ keeps decreasing and approaches $\theta^{(t)}(s, a_1)$
 - $\theta^{(t)}(s, a_2)$ keeps increasing and becomes the largest
5. Final stage (part 2): $t \geq t_{\text{ref}}$ (Lemma 7)
 - $\theta^{(t)}(s, a_1)$ keeps increasing and becomes the largest
 - $\theta^{(t)}(s, a_2)$ decreases a lot

4.3 Proof outline

We are now in a position to outline the main steps of the proof of Theorem 1 and Theorem 2, with details deferred to the appendix. In the following, Steps 1-6 are devoted to analyzing the dynamics of softmax PG methods when applied to the constructed MDP \mathcal{M} , which in turn establish Theorem 2. Step 7 describes how these can be easily adapted to prove Theorem 1, by slightly modifying the MDP construction.

Step 1: bounding the discounted state visitation distributions

In view of the PG update rule (12), the size of the policy gradient relies heavily on the discounted state visitation distribution $d_\mu^{(t)}(s)$. In light of this observation, this step aims to quantify the magnitudes of $d_\mu^{(t)}(s)$, for which we start with several universal lower bounds regardless of the policy in use.

Lemma 2. *For any policy π , the following lower bounds hold true:*

$$d_\mu^\pi(s) \geq c_m \gamma (1 - \gamma)^2, \quad \text{if } s \in \{3, \dots, H\}, \quad (36a)$$

$$d_\mu^\pi(\bar{s}) \geq c_m \gamma (1 - \gamma)^2, \quad \text{if } \bar{s} \in \{\bar{1}, \dots, \bar{H}\}, \quad (36b)$$

$$d_\mu^\pi(1) \geq \frac{c_m \gamma (1 - \gamma)^2}{|\mathcal{S}_1|} = \gamma (1 - \gamma) \frac{c_m}{c_{b,1}} \cdot \frac{1}{|\mathcal{S}|}, \quad (36c)$$

$$d_\mu^\pi(2) \geq \frac{c_m \gamma (1 - \gamma)^2}{|\mathcal{S}_2|} = \gamma (1 - \gamma) \frac{c_m}{c_{b,2}} \cdot \frac{1}{|\mathcal{S}|}. \quad (36d)$$

As it turns out, the above lower bounds are order-wise tight estimates prior to certain crucial crossing times. This is formalized in the following lemma, where we recall the definition of τ_s in (18).

Lemma 3. *Under the assumption (35), the following results hold:*

$$\forall 3 \leq s \leq H, t \leq t_s(\tau_s) : \quad d_\mu^{(t)}(s) \leq 14c_m(1 - \gamma)^2, \quad (37a)$$

$$\forall 2 \leq s \leq H, t \leq t_s(\tau_s) : \quad d_\mu^{(t)}(\bar{s}) \leq 14c_m(1 - \gamma)^2, \quad (37b)$$

$$\forall t \leq t_2(\tau_2) : \quad d_\mu^{(t)}(2) \leq \frac{1 - \gamma}{|\mathcal{S}|} \left(1 + \frac{8c_m}{c_{b,2}} \right), \quad (37c)$$

$$\forall t \leq t_2(\tau_2) : \quad d_\mu^{(t)}(\bar{1}) \leq 14c_m(1 - \gamma)^2, \quad (37d)$$

$$\forall t \leq \min\{t_1(\tau_1), t_2(\tau_2)\} : \quad d_\mu^{(t)}(1) \leq \frac{1 - \gamma}{|\mathcal{S}|} \left(1 + \frac{17c_m}{c_{b,1}} \right). \quad (37e)$$

Remark 6. As will be demonstrated in Lemma 4, one has $t_1(\tau_1) \leq t_2(\tau_2)$ for properly chosen constants $c_{b,1}$, $c_{b,2}$ and c_m . Therefore, we shall bear in mind that the properties (37d) and (37e) hold for any $t \leq t_1(\tau_1)$.

The proofs of these two lemmas are deferred to Appendix B. The sets of booster states, whose cardinality is controlled by c_m , play an important role in sandwiching the initial distribution of the states in \mathcal{S}_1 , \mathcal{S}_2 , $\mathcal{S}_{\text{primary}}$, and \mathcal{S}_{adj} . Combining these bounds, we uncover the following properties happening before $V^{(t)}(s)$ exceeds τ_s :

- For any key primary state $s \in \{3, \dots, H\}$ or any adjoint state $s \in \{\bar{1}, \dots, \bar{H}\}$, one has

$$d_\mu^{(t)}(s) \asymp (1 - \gamma)^2.$$

- For any state s contained in the buffer state subsets \mathcal{S}_1 and \mathcal{S}_2 , we have

$$d_\mu^{(t)}(1) \asymp \frac{(1 - \gamma)^2}{|\mathcal{S}_1|} \quad \text{and} \quad d_\mu^{(t)}(2) \asymp \frac{(1 - \gamma)^2}{|\mathcal{S}_2|},$$

where we recall the size of \mathcal{S}_1 and \mathcal{S}_2 in (16). In other words, the discounted state visitation probability of any buffer state is substantially smaller than that of any key primary state $3, \dots, H$ or adjoint state. In principle, the size of each buffer state subset plays a crucial role in determining the associated $d_\mu^{(t)}(s)$ — the larger the size of the buffer state subset, the smaller the resulting state visitation probability.

- Further, the aggregate discounted state visitation probability of the above states is no more than the order of

$$(1 - \gamma)^2 \cdot H \asymp 1 - \gamma = o(1),$$

which is vanishingly small. In fact, state 0 and the booster states account for the dominant fraction of state visitations at the initial stage of the algorithm.

Step 2: characterizing the crossing times for the first few states (\mathcal{S}_1 , \mathcal{S}_2 , and $\bar{1}$)

Armed with the bounds on $d_\mu^{(t)}$ developed in Step 1, we can move forward to study the crossing times for the key states. In this step, we pay attention to the crossing times for the buffer states $\mathcal{S}_1, \mathcal{S}_2$ as well as the first adjoint state $\bar{1}$, which forms a crucial starting point towards understanding the behavior of the subsequent states. Specifically, the following lemma develops lower and upper bounds regarding these quantities, whose proof can be found in Appendix C.

Lemma 4. *Suppose that (35) holds. If $|\mathcal{S}| \geq 1/(1-\gamma)^2$, then the crossing times satisfy*

$$\frac{\log 3}{1 + 17c_m/c_{b,1}} \frac{|\mathcal{S}|}{\eta} \leq t_1(\tau_1) \leq t_1(\gamma^2 - 1/4) \leq t_2(\tau_2) \leq t_2(\gamma^4 - 1/4) \leq \frac{15c_{b,2}}{c_m} \frac{|\mathcal{S}|}{\eta}. \quad (38a)$$

In addition, if $|\mathcal{S}| \geq \frac{320\gamma^3}{c_m(1-\gamma)^2}$, then one has

$$t_2(\tau_2) > t_{\bar{1}}(\gamma^3 - 1/4). \quad (38b)$$

For properly chosen constants $c_{b,1}$, $c_{b,2}$ and c_m , Lemma 4 delivers the following important messages:

- The cross times of these first few states are already fairly large; for instance,

$$t_1(\tau_1) \asymp t_2(\tau_2) \asymp \frac{|\mathcal{S}|}{\eta}, \quad (39)$$

which scale linearly with the state space dimension. As we shall see momentarily, while $t_1(\tau_1)$ and $t_2(\tau_2)$ remain polynomially large, these play a pivotal role in ensuring rapid explosion of the crossing times of the states that follow (namely, the states $\{3, \dots, H\}$).

- We can guarantee a strict ordering such that the crossing time of state 2 is at least as large as that of both state 1 and state $\bar{1}$. This property is helpful as well for subsequent analysis.

Step 3: understanding the dynamics of $\theta^{(t)}(s, a)$ before $t_{s-2}(\tau_{s-2})$

With the above characterization of the crossing times for the first few states, we are ready to investigate the dynamics of $\theta^{(t)}(s, a)$ ($3 \leq s \leq H$) at the initial stage, that is, the duration prior to the threshold $t_{s-2}(\tau_{s-2})$. Our finding for this stage is summarized in the following lemma, with the proof deferred to Appendix D.

Lemma 5. *Suppose that (35) holds. For any $3 \leq s \leq H$ and any $0 \leq t \leq t_{s-2}(\tau_{s-2})$, one has*

$$\theta^{(t)}(s, a_1) \leq -\frac{1}{2} \log \left(1 + \frac{c_m \gamma}{35} \eta (1 - \gamma)^2 t \right) \quad (40)$$

$$\text{and } \theta^{(t)}(s, a_0) \geq \theta^{(t)}(s, a_2) \geq 0. \quad (41)$$

Lemma 5 makes clear the behavior of $\theta^{(t)}(s, a)$ during this initial stage:

- The iterate $\theta^{(t)}(s, a_1)$ associated with the optimal action a_1 keeps dropping at a rate of $\log(O(\frac{1}{\sqrt{t}}))$, and remains the smallest compared to the ones with other actions (since $\theta^{(t)}(s, a_1) \leq 0 \leq \theta^{(t)}(s, a_2) \leq \theta^{(t)}(s, a_0)$).
- The other two iterates $\theta^{(t)}(s, a_0)$ and $\theta^{(t)}(s, a_2)$ stay non-negative throughout this stage, with a_0 being perceived as more favorable than the other two actions.
- In fact, a closer inspection of the proof in Appendix D reveals that $\theta^{(t)}(s, a_2)$ remains increasing — even though at a rate slower than that of $\theta^{(t)}(s, a_0)$ — throughout this stage (see (134) and the gradient expression (12b)).

In particular, around the threshold $t_{s-2}(\tau_{s-2})$, the iterate $\theta^{(t)}(s, a_1)$ becomes as small as

$$\exp\left(\theta^{(t)}(s, a_1)\right) \leq O\left(\frac{1}{\sqrt{\eta(1-\gamma)^2 t_{s-2}(\tau_{s-2})}}\right).$$

In fact, an inspection of the proof of this lemma reveals that

$$\pi^{(t)}(a_1 | s) \leq O\left(\frac{1}{\eta(1-\gamma)^2 t_{s-2}(\tau_{s-2})}\right) \quad \text{when } t = t_{s-2}(\tau_{s-2}).$$

This means that $\pi^{(t)}(a_1 | s)$ becomes smaller for a larger $t_{s-2}(\tau_{s-2})$, making it more difficult to return/converge to 1 afterward.

Step 4: understanding the dynamics of $\theta^{(t)}(s, a)$ between $t_{s-2}(\tau_{s-2})$ and $t_{s-1}(\tau_s)$

Next, we investigate, for any $3 \leq s \leq H$, the behavior of the iterates during an “intermediate” stage, namely, the duration when the iteration count t is between $t_{s-2}(\tau_{s-2})$ and $t_{s-1}(\tau_s)$. This is summarized in the following lemma, whose proof can be found in Appendix E.

Lemma 6. *Consider any $3 \leq s \leq H$. Assume that (35) holds. Suppose that*

$$t_{s-1}(\tau_{s-1}) > t_{s-2}(\tau_{s-1}) + \frac{2444s}{c_m \gamma \eta (1-\gamma)^2}, \quad (42a)$$

$$t_3(\tau_3) > t_2(\gamma^4 - 1/4). \quad (42b)$$

Then one has

$$\theta^{(t_{s-1}(\tau_s))}(s, a_1) \leq \theta^{(t_{s-2}(\tau_{s-2}))}(s, a_1) \quad \text{and} \quad \theta^{(t_{s-1}(\tau_s))}(s, a_2) \geq 0. \quad (43)$$

In particular, when $s = 3$, the results in (43) hold true without requiring the assumption (42).

Remark 7. Condition (42a) only requires $t_{s-1}(\tau_{s-1})$ to be slightly larger than $t_{s-2}(\tau_{s-1})$, which will be justified using an induction argument when proving the main theorem.

As revealed by the claim (43) of Lemma 6, the iterate $\theta^{(t)}(s, a_2)$ remains sufficiently large during this intermediate stage. In the meantime, Lemma 6 guarantees that during this stage, $\theta^{(t)}(s, a_1)$ lies below the level of $\theta^{(t_{s-2}(\tau_{s-2}))}(s, a_1)$ that has been pinned down in Lemma 5 (which has been shown to be quite small). Both of these properties make clear that the iterates $\theta^{(t)}(s, a)$ remain far from optimal at the end of this intermediate stage.

Step 5: establishing a blowing-up phenomenon

The next lemma, which plays a pivotal role in developing the desired exponential convergence lower bound, demonstrates that the cross times explode at a super fast rate. The proof is postponed to Appendix F.

Lemma 7. *Consider any $3 \leq s \leq H$. Suppose that (35) holds and*

$$t_{s-2}(\tau_{s-2}) \geq \left(\frac{6300e}{c_p(1-\gamma)}\right)^4 \frac{1}{\frac{c_m \gamma}{35} \eta (1-\gamma)^2}. \quad (44)$$

Then there exists a time instance t_{ref} obeying $t_{s-1}(\tau_s) \leq t_{\text{ref}} < t_s(\tau_s)$ such that

$$\theta^{(t_{\text{ref}})}(s, a_0) \leq \theta^{(t_{\text{ref}})}(s, a_1) - \log\left(\frac{c_p}{16128}(1-\gamma)\right), \quad (45a)$$

$$\theta^{(t_{\text{ref}})}(s, a_1) \leq -\frac{1}{2} \log\left(1 + \frac{c_m \gamma}{35} \eta (1-\gamma)^2 t_{s-2}(\tau_{s-2})\right) + 1, \quad (45b)$$

and at the same time,

$$t_s(\tau_s) - t_{\text{ref}} \geq 10^{-10} c_p c_m^{0.5} \eta^{0.5} (1-\gamma)^2 \left(t_{s-2}(\tau_{s-2})\right)^{1.5}. \quad (45c)$$

The most important message of Lemma 7 lies in property (45c). In a nutshell, this property uncovers that the crossing time $t_s(\tau_s)$ is substantially larger than $t_{s-2}(\tau_{s-2})$, namely,

$$t_s(\tau_s) \gtrsim \eta^{0.5}(1-\gamma)^2 \left(t_{s-2}(\tau_{s-2}) \right)^{1.5}, \quad (46)$$

thus leading to explosion at a super-linear rate. By contrast, the other two properties unveil some important features happening between $t_{s-1}(\tau_s)$ and $t_s(\tau_s)$ that in turn lead to property (45c). In words, property (45a) requires $\theta^{(t_{\text{ref}})}(s, a_0)$ to be not much larger than $\theta^{(t_{\text{ref}})}(s, a_1)$; property (45b) indicates that: when $t_{s-2}(\tau_{s-2})$ is large, both $\theta^{(t_{\text{ref}})}(s, a_1)$ and $\theta^{(t_{\text{ref}})}(s, a_0)$ are fairly small, with $\theta^{(t_{\text{ref}})}(s, a_2)$ being the dominant one (due to the fact $\sum_a \theta^{(t_{\text{ref}})}(s, a) = 0$ as will be seen in Part (vii) of Lemma 8).

The reader might naturally wonder what the above results imply about $\pi^{(t_{\text{ref}})}(a_1 | s)$ (as opposed to $\theta^{(t_{\text{ref}})}(s, a_1)$). Towards this end, we make the observation that

$$\begin{aligned} \pi^{(t_{\text{ref}})}(a_1 | s) &= \frac{\exp(\theta^{(t_{\text{ref}})}(s, a_1))}{\sum_a \exp(\theta^{(t_{\text{ref}})}(s, a))} \leq \frac{\exp(\theta^{(t_{\text{ref}})}(s, a_1))}{\exp(\theta^{(t_{\text{ref}})}(s, a_2))} \stackrel{(i)}{=} \exp(2\theta^{(t_{\text{ref}})}(s, a_1) + \theta^{(t_{\text{ref}})}(s, a_0)) \\ &\stackrel{(ii)}{\lesssim} \frac{1}{(1-\gamma) \left(\eta(1-\gamma)^2 t_{s-2}(\tau_{s-2}) \right)^{1.5}} \asymp \frac{1}{\eta^{1.5}(1-\gamma)^4 (t_{s-2}(\tau_{s-2}))^{1.5}}, \end{aligned} \quad (47)$$

where (i) holds true since $\sum_a \theta^{(t_{\text{ref}})}(s, a) = 0$ (a well-known property of policy gradient methods as recorded in Lemma 8(vii)), and (ii) follows from the properties (45a) and (45b). In other words, $\pi^{(t_{\text{ref}})}(s, a_1)$ is inversely proportional to $(t_{s-2}(\tau_{s-2}))^{3/2}$. As we shall see, the time taken for $\pi^{(t_{\text{ref}})}(a_1 | s)$ to converge to 1 is proportional to the inverse policy iterate $(\pi^{(t)}(s, a_1))^{-1}$, meaning that it is expected to take an order of $(t_{s-2}(\tau_{s-2}))^{3/2}$ iterations to increase from $\pi^{(t_{\text{ref}})}(s, a_1)$ to 1.

Step 6: putting all this together to establish Theorem 2

With the above steps in place, we are ready to combine them to establish the following result. As can be easily seen, Theorem 2 is an immediate consequence of Theorem 3.

Theorem 3. *Suppose that (35) holds. There exist some universal constants $c_1, c_2, c_3 > 0$ such that*

$$t_s(0.5) \geq c_1 \frac{|\mathcal{S}|^{\frac{2}{3}}}{\eta} \left(c_2 |\mathcal{S}| \right)^{\frac{1}{3} \cdot 1.5^{\lfloor s/2 \rfloor}}, \quad (48)$$

provided that

$$|\mathcal{S}| \geq \frac{c_3}{(1-\gamma)^6}. \quad (49)$$

Proof of Theorem 3. Let us define two universal constants $C_1 := \frac{\log 3}{1+17c_m/c_{b,1}}$ and $C_2 := \frac{10^{-20}c_p^2c_m \log 3}{1+17c_m/c_{b,1}}$. We claim that if one can show that

$$t_s(\tau_s) \geq C_1 \frac{|\mathcal{S}|}{\eta} \left(C_2 (1-\gamma)^4 |\mathcal{S}| \right)^{1.5^{\lfloor (s-1)/2 \rfloor - 1}}, \quad (50)$$

then the desired bound (48) holds true directly. In order to see this, recall that $\tau_s \leq 1/2$ by definition, and therefore,

$$t_s(0.5) \geq t_s(\tau_s) \stackrel{(i)}{\geq} C_1 \frac{|\mathcal{S}|}{\eta} \left(C_2 \sqrt[3]{|\mathcal{S}|} \right)^{1.5^{\lfloor (s-1)/2 \rfloor - 1}} \stackrel{(ii)}{\geq} c_1 \frac{|\mathcal{S}|^{\frac{2}{3}}}{\eta} \left(c_2 |\mathcal{S}| \right)^{\frac{1}{3} \cdot 1.5^{\lfloor s/2 \rfloor}}.$$

Here, (i) follows from (50) in conjunction with the assumption (49), whereas (ii) holds true by setting $c_1 = C_1/C_2$ and $c_2 = C_2^3$.

It is then sufficient to prove the inequality (50), towards which we shall resort to mathematical induction in conjunction with the following induction hypothesis

$$t_s(\tau_s) > t_{s-1}(\tau_s) + \frac{2444(s+1)}{c_m \gamma \eta (1-\gamma)^2}, \quad \text{for } s \geq 3. \quad (51)$$

- We start with the cases with $s = 1, 2, 3$. It follows from Lemma 4 that

$$t_2(\tau_2) \geq t_1(\tau_1) \geq \frac{\log 3}{1 + 17c_m/c_{b,1}} \frac{|\mathcal{S}|}{\eta} = C_1 \frac{|\mathcal{S}|}{\eta}, \quad (52)$$

which validates the above claim (50) for $s = 1$ and $s = 2$. In addition, Lemma 7 ensures that

$$\begin{aligned} t_3(\tau_3) - \max \left\{ t_1(\gamma^3 - 1/4), t_2(\tau_3) \right\} &\geq 10^{-10} c_p c_m^{0.5} \eta^{0.5} (1 - \gamma)^2 \left(t_1(\tau_1) \right)^{1.5} \\ &\geq \frac{9776}{c_m \gamma \eta (1 - \gamma)^2}, \end{aligned} \quad (53)$$

where the last inequality is guaranteed by (52) and the assumption $|\mathcal{S}| \geq \max \left\{ \frac{4888}{C_1 c_m \gamma (1 - \gamma)^2}, \frac{4}{C_2 (1 - \gamma)^4} \right\}$. This implies that the inequality (51) is satisfied when $s = 3$.

- Next, suppose that the inequality (50) holds true up to state $s - 1$ and the inequality (51) holds up to s for some $3 \leq s \leq H$. To invoke the induction argument, it suffices to show that the inequality (50) continues to hold for state s and the inequality (51) remains valid for $s + 1$. This will be accomplished by taking advantage of Lemma 7.

Given that the inequality (50) holds true for every state up to $s - 1$, one has

$$t_{s-1}(\tau_{s-1}) \geq t_{s-2}(\tau_{s-2}) \geq C_1 \frac{|\mathcal{S}|}{\eta} \left(C_2 (1 - \gamma)^4 |\mathcal{S}| \right)^{1.5 \lfloor (s-3)/2 \rfloor - 1} \geq \left(\frac{6300e}{c_p (1 - \gamma)} \right)^4 \frac{1}{\frac{c_m \gamma}{35} \eta (1 - \gamma)^2},$$

where the last inequality is satisfied provided that $|\mathcal{S}| > \max \left\{ \left(\frac{6300e}{c_p} \right)^4 \frac{35}{C_1 c_m \gamma (1 - \gamma)^6}, \frac{4}{C_2 (1 - \gamma)^4} \right\}$. Therefore, Lemma 7 is applicable for both s and $s + 1$, thus leading to

$$\begin{aligned} t_s(\tau_s) - t_{s-1}(\tau_s) &\geq 10^{-10} c_p c_m^{0.5} \eta^{0.5} (1 - \gamma)^2 \left(t_{s-2}(\tau_{s-2}) \right)^{1.5} \\ &\geq 10^{-10} c_p c_m^{0.5} \eta^{0.5} (1 - \gamma)^2 \left(C_1 \frac{|\mathcal{S}|}{\eta} \left(C_2 (1 - \gamma)^4 |\mathcal{S}| \right)^{1.5 \lfloor (s-3)/2 \rfloor - 1} \right)^{1.5} \\ &\geq C_1 \frac{|\mathcal{S}|}{\eta} \left(C_2 (1 - \gamma)^4 |\mathcal{S}| \right)^{1.5 \lfloor (s-1)/2 \rfloor - 1}. \end{aligned}$$

Here, the last step relies on the condition $10^{-10} c_p c_m^{0.5} \eta^{0.5} (1 - \gamma)^2 \left(C_1 \frac{|\mathcal{S}|}{\eta} \right)^{0.5} \geq 1$. This in turn establishes the property (50) for state s (given that $t_{s-1}(\tau_s) \geq 0$). In addition, Lemma 7 — when applied to $s + 1$ — gives

$$\begin{aligned} t_{s+1}(\tau_{s+1}) - t_s(\tau_{s+1}) &\geq 10^{-10} c_p c_m^{0.5} \eta^{0.5} (1 - \gamma)^2 \left(t_{s-1}(\tau_{s-1}) \right)^{1.5} \\ &\geq 10^{-10} c_p c_m^{0.5} \eta^{0.5} (1 - \gamma)^2 \left(C_1 \frac{|\mathcal{S}|}{\eta} \left(C_2 (1 - \gamma)^4 |\mathcal{S}| \right)^{1.5 \lfloor (s-2)/2 \rfloor - 1} \right)^{1.5} \\ &\geq C_1 \frac{|\mathcal{S}|}{\eta} \left(C_2 (1 - \gamma)^4 |\mathcal{S}| \right)^{1.5 \lfloor s/2 \rfloor - 1} \\ &\geq \frac{2444(s+2)}{c_m \gamma \eta (1 - \gamma)^2}, \end{aligned}$$

where the last step follows as long as $|\mathcal{S}| > \max \left\{ \frac{4888}{C_1 c_m \gamma (1 - \gamma)^2}, \frac{4}{C_2 (1 - \gamma)^4} \right\}$. We have thus established the property (51) for state $s + 1$.

Putting all the above pieces together, we arrive at the inequality (50), thus establishing Theorem 3. \square

Step 7: adapting the proof to establish Theorem 1

Thus far, we have established Theorem 2, and are well equipped to return to the proof of Theorem 1. As a remark, Theorem 2 and its analysis posits that for a large fraction of the key primary states (as well as their associated adjoint states), softmax PG methods can take a prohibitively large number of iterations to converge. The issue, however, is that there are in total only $O(H)$ key primary states and adjoint states, accounting for a vanishingly small fraction of all $|\mathcal{S}|$ states. In order to extend Theorem 2 to Theorem 1 (the latter of which is concerned with the error averaged over the entire state space), we would need to show that the value functions associated with those booster states — which account for a large fraction of the state space — also converge slowly.

In the MDP instance constructed in Section 3, however, the action space associated with the booster states is a singleton set, meaning that the action is always optimal. As a result, we would first need to modify/augment the action space of booster states, so as to ensure that their learned actions remain suboptimal before the algorithm converges for the associated key primary states and adjoint states.

A modified MDP instance. We now augment the action space for all booster states in the MDP \mathcal{M} constructed in Section 3, leading to a *slightly modified MDP* denoted by $\mathcal{M}_{\text{modified}}$:

- for any key primary state $s \in \{3, \dots, H\}$ and any associated booster state $\hat{s} \in \hat{\mathcal{S}}_s$, take the action space of \hat{s} to be $\{a_0, a_1\}$ and let

$$P(0 | \hat{s}, a_0) = 0.9, \quad P(s | \hat{s}, a_0) = 0.1, \quad r(\hat{s}, a_0) = 0.9\gamma\tau_s, \quad P(s | \hat{s}, a_1) = 1, \quad r(\hat{s}, a_1) = 0; \quad (54)$$

- for any key adjoint state $\bar{s} \in \{\bar{1}, \dots, \bar{H}\}$ and any associated booster state $\hat{s} \in \hat{\mathcal{S}}_{\bar{s}}$, take the action space of \hat{s} to be $\{a_0, a_1\}$ and let

$$P(0 | \hat{s}, a_0) = 0.9, \quad P(\bar{s} | \hat{s}, a_0) = 0.1, \quad r(\hat{s}, a_0) = 0.9\gamma^2\tau_s, \quad P(\bar{s} | \hat{s}, a_1) = 1, \quad r(\hat{s}, a_1) = 0; \quad (55)$$

- all other components of $\mathcal{M}_{\text{modified}}$ remain identical to those of the original \mathcal{M} .

Analysis for the new booster states. Given that the dynamics of non-booster states are unaffected by the booster states, it suffices to perform analysis for the booster states. Let us first consider any key primary state s and any associated booster state \hat{s} .

- As can be easily seen,

$$\begin{aligned} Q^{(t)}(\hat{s}, a_0) &= r(\hat{s}, a_0) + \gamma P(s | \hat{s}, a_0) V^{(t)}(s) + \gamma P(0 | \hat{s}, a_0) V^{(t)}(0) = 0.9\gamma\tau_s + 0.1\gamma V^{(t)}(s), \\ Q^{(t)}(\hat{s}, a_1) &= r(\hat{s}, a_1) + \gamma P(s | \hat{s}, a_0) V^{(t)}(s) = \gamma V^{(t)}(s), \end{aligned}$$

where we have used the basic fact $V^{(t)}(0) = 0$ (see (73) in Lemma 8). Given that $V^{(t)}(\hat{s})$ is a convex combination of $Q^{(t)}(\hat{s}, a_0)$ and $Q^{(t)}(\hat{s}, a_1)$, one can easily see that: if $V^{(t)}(s) < \tau_s$, then one necessarily has $V^{(t)}(\hat{s}) < \gamma\tau_s$.

- Similarly, the optimal Q-function w.r.t. \hat{s} is given by

$$\begin{aligned} Q^*(\hat{s}, a_0) &= r(\hat{s}, a_0) + \gamma P(s | \hat{s}, a_0) V^*(s) + \gamma P(0 | \hat{s}, a_0) V^*(0) = 0.9\gamma\tau_s + 0.1\gamma V^*(s), \\ Q^*(\hat{s}, a_1) &= r(\hat{s}, a_1) + \gamma P(s | \hat{s}, a_0) V^*(s) = \gamma V^*(s), \end{aligned}$$

which together with Lemma 1 and the definition (18) of τ_s indicates that $V^*(\hat{s}) = Q^*(\hat{s}, a_1) = \gamma^{2s+1}$.

- The above facts taken collectively imply that: if $V^{(t)}(s) < \tau_s$, then

$$V^*(\hat{s}) - V^{(t)}(\hat{s}) > \gamma^{2s+1} - \gamma\tau_s = \gamma(\gamma^{2s} - 0.5\gamma^{\frac{2s}{3}}) > 0.22, \quad (56)$$

provided that γ is sufficiently large (which is satisfied under the condition (35)).

Similarly, for any key adjoint state \bar{s} and any associated booster state \hat{s} , if $V^{(t)}(\bar{s}) < \gamma\tau_s$, then one must have

$$V^*(\hat{s}) - V^{(t)}(\hat{s}) > 0.22. \quad (57)$$

Repeating the same proof as for Theorem 2, one can easily show that (with slight adjustment of the universal constants)

$$t_s(\tau_s) = t_{\bar{s}}(\gamma\tau_s) \geq \frac{1}{\eta} |\mathcal{S}|^{2^{\Omega(\frac{1}{1-\gamma})}} \quad \text{for all } s > 0.1H. \quad (58)$$

This taken together with the above analysis suffices to establish Theorem 1, given the following two simple facts: (i) there are $2Hc_m(1-\gamma)|\mathcal{S}| = 2c_m c_h |\mathcal{S}|$ booster states, and (ii) more than 90% of them need a prohibitively large number of iterations (cf. (58)) to reach 0.22-optimality. Here, we can take $c_m c_h > 0.18$ which satisfies (35). The proof is thus complete.

5 Discussion

This paper has developed an algorithm-specific lower bound on the iteration complexity of the softmax policy gradient method, obtained by analyzing its trajectory on a carefully-designed hard MDP instance. We have shown that the iteration complexity of softmax PG methods can scale pessimistically, in fact (super-)exponentially, with the dimension of the state space and the effective horizon of the discounted MDP of interest. Our finding makes apparent the potential *inefficiency* of softmax PG methods in solving large-dimensional and long-horizon problems. In turn, this suggests the necessity of carefully adjusting update rules and/or enforcing proper regularization in accelerating policy gradient methods.

Our work relies heavily on proper exploitation of the structural properties of the MDP in algorithm-dependent analysis, which might shed light on lower bound construction for other algorithms as well. For instance, if the objective function (i.e., the value function) is augmented by a regularization term, how does the choice of regularization affect the global convergence behavior? While Agarwal et al. (2021) demonstrated polynomial-time convergence of PG methods in the presence of log-barrier regularization, non-asymptotic analysis of PG methods with other popular regularization — particularly entropy regularization — remains unavailable in existing literature. How to understand the (in)-effectiveness of entropy-regularized PG methods is of fundamental importance in the theory of policy optimization. Additionally, the current paper concentrates on the use of constant learning rates; it falls short of accommodating more adaptive learning rates, which might be a potential solution to accelerate vanilla PG methods. Furthermore, our strategy for lower bound construction might be extended to unveil algorithmic bottlenecks of policy optimization in multi-agent Markov games as well. All this is worthy of future investigation.

Acknowledgements

Y. Wei is supported in part by the the NSF grants CCF-2106778 and CAREER award DMS-2143215. Y. Chi is supported in part by the grants ONR N00014-18-1-2142 and N00014-19-1-2404, ARO W911NF-18-1-0303, NSF CCF-1806154, CCF-2007911 and CCF-2106778. Y. Chen is supported in part by the Alfred P. Sloan Research Fellowship, the Google Research Scholar Award, the AFOSR grant FA9550-22-1-0198, the ONR grant N00014-22-1-2354, and the NSF grants CCF-2221009, CCF-1907661, IIS-2218713 and IIS-2218773.

References

- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. (2021). On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98):1–76.
- Agazzi, A. and Lu, J. (2021). Global optimality of softmax policy gradient with single hidden layer neural networks in the mean-field regime. *International Conference on Learning Representations (ICLR)*.

- Alacaoglu, A., Viano, L., He, N., and Cevher, V. (2022). A natural actor-critic framework for zero-sum Markov games. In *International Conference on Machine Learning*, pages 307–366. PMLR.
- Asadi, K. and Littman, M. L. (2017). An alternative softmax operator for reinforcement learning. In *International Conference on Machine Learning*, pages 243–252.
- Azar, M. G., Munos, R., and Kappen, H. J. (2013). Minimax PAC bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91(3):325–349.
- Beck, A. (2017). *First-order methods in optimization*. SIAM.
- Bhandari, J. (2020). *Optimization Foundations of Reinforcement Learning*. PhD thesis, Columbia University.
- Bhandari, J. and Russo, D. (2019). Global optimality guarantees for policy gradient methods. *arXiv preprint arXiv:1906.01786*.
- Bhandari, J. and Russo, D. (2021). On the linear convergence of policy gradient methods for finite MDPs. In *International Conference on Artificial Intelligence and Statistics*, pages 2386–2394. PMLR.
- Cai, Q., Yang, Z., Jin, C., and Wang, Z. (2020). Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pages 1283–1294.
- Cen, S., Cheng, C., Chen, Y., Wei, Y., and Chi, Y. (2022a). Fast global convergence of natural policy gradient methods with entropy regularization. *Operations Research*, 70(4):2563–2578.
- Cen, S., Chi, Y., Du, S. S., and Xiao, L. (2022b). Faster last-iterate convergence of policy optimization in zero-sum Markov games. *arXiv preprint arXiv:2210.01050*.
- Cen, S., Wei, Y., and Chi, Y. (2021). Fast policy extragradient methods for competitive games with entropy regularization. *Advances in Neural Information Processing Systems*, 34:27952–27964.
- Daskalakis, C., Foster, D. J., and Golowich, N. (2020). Independent policy gradient methods for competitive reinforcement learning. *Advances in neural information processing systems*, 33:5527–5540.
- Ding, D., Zhang, K., Basar, T., and Jovanovic, M. (2020). Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33.
- Domingues, O. D., Ménard, P., Kaufmann, E., and Valko, M. (2021). Episodic reinforcement learning in finite MDPs: Minimax lower bounds revisited. In *Algorithmic Learning Theory*, pages 578–598.
- Du, S., Jin, C., Jordan, M., Póczos, B., Singh, A., and Lee, J. (2017). Gradient descent can take exponential time to escape saddle points. In *Advances in Neural Information Processing Systems*, pages 1068–1078.
- Fazel, M., Ge, R., Kakade, S., and Mesbahi, M. (2018). Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476.
- Jansch-Porto, J. P., Hu, B., and Dullerud, G. E. (2020). Convergence guarantees of policy optimization methods for Markovian jump linear systems. In *American Control Conference*, pages 2882–2887. IEEE.
- Kakade, S. M. (2002). A natural policy gradient. In *Advances in neural information processing systems*, pages 1531–1538.
- Khamaru, K., Pananjady, A., Ruan, F., Wainwright, M. J., and Jordan, M. I. (2021). Is temporal difference learning optimal? an instance-dependent analysis. *SIAM Journal on Mathematics of Data Science*, 3(4):1013–1040.
- Khodadadian, S., Doan, T. T., Romberg, J., and Maguluri, S. T. (2022). Finite sample analysis of two-time-scale natural actor-critic algorithm. *IEEE Transactions on Automatic Control*.
- Konda, V. R. and Tsitsiklis, J. N. (2000). Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014.

- Lan, G. (2022a). Policy mirror descent for reinforcement learning: Linear convergence, new sampling complexity, and generalized problem classes. *Mathematical programming*, pages 1–48.
- Lan, G. (2022b). Policy optimization over general state and action spaces. *arXiv preprint arXiv:2211.16715*.
- Lee, J. D., Simchowitz, M., Jordan, M. I., and Recht, B. (2016). Gradient descent only converges to minimizers. In *Conference on learning theory*, pages 1246–1257.
- Li, G., Cai, C., Chen, Y., Wei, Y., and Chi, Y. (2021). Is Q-learning minimax optimal? a tight sample complexity analysis. *arXiv preprint arXiv:2102.06548*.
- Li, G., Shi, L., Chen, Y., Chi, Y., and Wei, Y. (2022a). Settling the sample complexity of model-based offline reinforcement learning. *arXiv preprint arXiv:2204.05275*.
- Li, Y., Zhao, T., and Lan, G. (2022b). First-order policy optimization for robust Markov decision process. *arXiv preprint arXiv:2209.10579*.
- Liu, B., Cai, Q., Yang, Z., and Wang, Z. (2019). Neural proximal/trust region policy optimization attains globally optimal policy. In *Neural Information Processing Systems*, pages 10565–10576.
- Liu, Y., Zhang, K., Basar, T., and Yin, W. (2020). An improved analysis of (variance-reduced) policy gradient and natural policy gradient methods. *Advances in Neural Information Processing Systems*, 33.
- Mei, J., Gao, Y., Dai, B., Szepesvari, C., and Schuurmans, D. (2021). Leveraging non-uniformity in first-order non-convex optimization. In *International Conference on Machine Learning*, pages 7555–7564.
- Mei, J., Xiao, C., Dai, B., Li, L., Szepesvári, C., and Schuurmans, D. (2020a). Escaping the gravitational pull of softmax. *Advances in Neural Information Processing Systems*, 33.
- Mei, J., Xiao, C., Szepesvari, C., and Schuurmans, D. (2020b). On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pages 6820–6829.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.
- Pananjady, A. and Wainwright, M. J. (2020). Instance-dependent ℓ_∞ -bounds for policy evaluation in tabular reinforcement learning. *IEEE Transactions on Information Theory*, 67(1):566–585.
- Peters, J. and Schaal, S. (2008). Natural actor-critic. *Neurocomputing*, 71(7-9):1180–1190.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shani, L., Efroni, Y., and Mannor, S. (2020). Adaptive trust region policy optimization: Global convergence and faster rates for regularized MDPs. In *AAAI Conference on Artificial Intelligence*, volume 34, pages 5668–5675.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587):484–489.
- Sutton, R. S. (1984). Temporal credit assignment in reinforcement learning. *PhD thesis, University of Massachusetts*.
- Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063.

- Tu, S. and Recht, B. (2019). The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint. In *Conference on Learning Theory*, pages 3036–3083.
- Wang, L., Cai, Q., Yang, Z., and Wang, Z. (2019). Neural policy gradient methods: Global optimality and rates of convergence. In *International Conference on Learning Representations*.
- Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. (2021). Last-iterate convergence of decentralized optimistic gradient descent/ascent in infinite-horizon competitive Markov games. In *Conference on learning theory*, pages 4259–4299. PMLR.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Wu, Y. F., Zhang, W., Xu, P., and Gu, Q. (2020). A finite-time analysis of two time-scale actor-critic methods. *Advances in Neural Information Processing Systems*, 33:17617–17628.
- Xie, Q., Yang, Z., Wang, Z., and Minca, A. (2020). Provable fictitious play for general mean-field games. *arXiv preprint arXiv:2010.04211*.
- Xu, T., Liang, Y., and Lan, G. (2020a). A primal approach to constrained policy optimization: Global optimality and finite-time analysis. *arXiv preprint arXiv:2011.05869*.
- Xu, T., Wang, Z., and Liang, Y. (2020b). Non-asymptotic convergence analysis of two time-scale (natural) actor-critic algorithms. *arXiv preprint arXiv:2005.03557*.
- Yan, Y., Li, G., Chen, Y., and Fan, J. (2022). Model-based reinforcement learning is minimax-optimal for offline zero-sum Markov games. *arXiv preprint arXiv:2206.04044*.
- Yang, W., Li, X., Xie, G., and Zhang, Z. (2020). Finding the near optimal policy via adaptive reduced regularization in MDPs. *arXiv preprint arXiv:2011.00213*.
- Zhan, W., Cen, S., Huang, B., Chen, Y., Lee, J. D., and Chi, Y. (2021). Policy mirror descent for regularized reinforcement learning: A generalized framework with linear convergence. *arXiv preprint arXiv:2105.11066*.
- Zhang, J., Kim, J., O’Donoghue, B., and Boyd, S. (2021a). Sample efficient reinforcement learning with REINFORCE. In *AAAI Conference on Artificial Intelligence*, volume 35, pages 10887–10895.
- Zhang, J., Koppel, A., Bedi, A. S., Szepesvari, C., and Wang, M. (2020a). Variational policy gradient method for reinforcement learning with general utilities. *Advances in Neural Information Processing Systems*, 33.
- Zhang, J., Ni, C., Szepesvari, C., and Wang, M. (2021b). On the convergence and sample efficiency of variance-reduced policy gradient method. *Advances in Neural Information Processing Systems*, 34:2228–2240.
- Zhang, K., Hu, B., and Basar, T. (2020b). Policy optimization for \mathcal{H}_2 linear control with \mathcal{H}_∞ robustness guarantee: Implicit regularization and global convergence. In *Learning for Dynamics and Control*, pages 179–190.
- Zhang, K., Koppel, A., Zhu, H., and Basar, T. (2020c). Global convergence of policy gradient methods to (almost) locally optimal policies. *SIAM Journal on Control and Optimization*, 58(6):3586–3612.
- Zhang, X., Chen, Y., Zhu, X., and Sun, W. (2021c). Robust policy gradient against strong data corruption. In *International Conference on Machine Learning*, pages 12391–12401.
- Zhao, Y., Tian, Y., Lee, J., and Du, S. (2021). Provably efficient policy gradient methods for two-player zero-sum Markov games. *International Conference on Artificial Intelligence and Statistics*.

$\mathcal{S}, \mathcal{A}_s, \gamma$	state space, action space associated with state s , discount factor
$P(s' s, a)$	probability of transitioning from state s to state s' upon execution of action a
$r(s, a)$	immediate reward gained in state s when action a is taken; $r(s, a) \in [-1, 1]$
η	stepsize or learning rate
$\pi^{(t)}, \theta^{(t)}$	policy estimate and its associated parameterization in the t -th iteration
$V^\pi, V^{(t)}, V^*$	value function of π , value function of $\pi^{(t)}$, optimal value function
$Q^\pi, Q^{(t)}, Q^*$	Q-function of π , Q-function of $\pi^{(t)}$, optimal Q-function
$A^\pi, A^{(t)}$	advantage function of π , advantage function of $\pi^{(t)}$
μ	initial state distribution (used in defining the objective function (11))
$d_\mu^\pi, d_\mu^{(t)}$	discounted state visitation distribution of π and $\pi^{(t)}$ from initial state distribution μ
τ_s, p, r_s	useful quantities: $\tau_s = 0.5\gamma^{\frac{2s}{3}}$, $p = c_p(1 - \gamma)$ and $r_s = 0.5\gamma^{\frac{2s}{3} + \frac{5}{6}}$
H	number of key primary states: $H = \frac{c_b}{1-\gamma}$
$\mathcal{S}_1, \mathcal{S}_2$	buffer state subsets: $ \mathcal{S}_1 = c_{b,1}(1 - \gamma) \mathcal{S} $, $ \mathcal{S}_2 = c_{b,2}(1 - \gamma) \mathcal{S} $
$\widehat{\mathcal{S}}_s$	booster state sets w.r.t. state s : $ \widehat{\mathcal{S}}_s = c_m(1 - \gamma) \mathcal{S} $
$\mathcal{S}_{\text{primary}}, \mathcal{S}_{\text{adj}}$	set of key primary states, set of key adjoint states
\bar{s}	adjoint state associated with primary state s
$\mathcal{M}, \mathcal{M}_{\text{modified}}$	MDPs constructed to prove Theorem 2 and Theorem 1
$t_s(\tau)$	crossing time: $\arg \min\{t \mid V^{(t)}(s) \geq \tau\}$

Table 2: Summary of notation and parameters.

A Preliminary facts

A.1 Basic properties of the constructed MDP

In this section, we provide more basic properties about the MDP we have constructed (see Section 3). Specifically, we present a miscellaneous collection of basic relations regarding more general policies, postponing the proof to Appendix A.4.

Lemma 8. *Consider any policy π , and recall the quantities defined in (18). Suppose that $\gamma^{2H} \geq 1/2$ and $0 < c_p \leq 1/6$.*

(i) *For any state $s \in \{3, \dots, H\}$, one has*

$$\gamma^{\frac{3}{2}}\tau_{s-1} \leq Q^\pi(s, a_0) = r_s + \gamma^2 p \tau_{s-2} \leq \gamma^{\frac{1}{2}}\tau_s, \quad (59a)$$

$$Q^\pi(s, a_1) = \gamma V^\pi(s-1), \quad (59b)$$

$$Q^\pi(s, a_2) = r_s + \gamma p V^\pi(s-2) \leq \gamma^{\frac{1}{2}}\tau_s. \quad (59c)$$

If one further has $V^\pi(s-2) \geq 0$, then $Q^\pi(s, a_2) \geq \gamma^{\frac{3}{2}}\tau_{s-1}$.

(ii) *If $V^\pi(s) \geq \tau_s$ for some $s \in \{3, \dots, H\}$, then we necessarily have*

$$\pi(a_1 | s) \geq \frac{1-\gamma}{2}. \quad (60)$$

(iii) *For any $\bar{s} \in \{\bar{1}, \dots, \bar{H}\}$, one has*

$$Q^\pi(\bar{s}, a_0) = \gamma \tau_s \quad \text{and} \quad Q^\pi(\bar{s}, a_1) = \gamma V^\pi(s), \quad (61)$$

where we recall the definition of $V^\pi(1)$ and $V^\pi(2)$ in (25). In addition, if $\pi(a_1 | \bar{s}) > 0$, then

$$V^\pi(\bar{s}) \geq \gamma \tau_s \quad \text{holds if and only if} \quad V^\pi(s) \geq \tau_s. \quad (62)$$

This means that: if $\pi^{(t)}(a_1 | \bar{s}) > 0$ holds for all $t \geq 0$, then one necessarily has

$$t_{\bar{s}}(\gamma \tau_s) = t_s(\tau_s). \quad (63)$$

(iv) For any policy π , we have

$$Q^\pi(1, a_0) = -\gamma^2, \quad Q^\pi(1, a_1) = \gamma^2, \quad V^\pi(1) = -\gamma^2\pi(a_0 | 1) + \gamma^2\pi(a_1 | 1), \quad (64a)$$

$$Q^\pi(2, a_0) = -\gamma^4, \quad Q^\pi(2, a_1) = \gamma^4, \quad V^\pi(2) = -\gamma^4\pi(a_0 | 2) + \gamma^4\pi(a_1 | 2). \quad (64b)$$

(v) Consider any policy π obeying $\min_{a,s} \pi(a | s) > 0$. For every $s \in \{3, \dots, H\}$, if $V^\pi(s) \geq \gamma^{\frac{1}{2}}\tau_s$ occurs, then one necessarily has $V^\pi(s-1) \geq \tau_{s-1}$.

(vi) If $V^\pi(s-2) < \tau_{s-2}$ and $\pi(a_1 | \overline{s-2}) > 0$, then

$$Q^\pi(s, a_0) - Q^\pi(s, a_2) = \gamma p(\gamma\tau_{s-2} - V^\pi(\overline{s-2})) > 0.$$

If $V^\pi(s-1) \leq \tau_{s-1}$ and $V^\pi(\overline{s-2}) \geq 0$, then

$$\min\{Q^\pi(s, a_0), Q^\pi(s, a_2)\} - Q^\pi(s, a_1) \geq (1 - \gamma)/8.$$

(vii) Consider the softmax PG update rule (12). One has for any $s \in \mathcal{S}$ and any θ ,

$$\sum_a \frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta(s, a)} = 0 \quad \text{and} \quad \sum_a \theta^{(t)}(s, a) = 0 \quad (65)$$

Remark 8. As it turns out, invoking Part (v) of Lemma 8 recursively reveals that: for any $2 \leq s \leq H$ and any $t < t_s(\tau_s)$, we have

$$V^{(t)}(s') < \gamma^{1/2}\tau_{s'} < \tau_{s'} \quad \text{for all } s' \text{ obeying } s \leq s' \leq H. \quad (66)$$

This in turn implies that $t_2(\tau_2) \leq t_3(\tau_3) \leq \dots \leq t_H(\tau_H)$ according to the definition (30).

Let us point out some implications of Lemma 8 that help guide our lower bound analysis. Once again, it is helpful to look at the results of this lemma when $\gamma \approx 1$ and $\gamma^H \approx 1$. In this case, the quantities defined in (18) obey $\tau_s \approx r_s \approx 1/2$, allowing us to obtain the following messages:

- Lemma 8(i) implies that, under mild conditions,

$$Q^\pi(s, a_0) \approx Q^\pi(s, a_2) \approx 1/2$$

holds any $s \in \{3, \dots, H\}$ and any policy π . In comparison to the optimal values (27), this result uncovers the strict sub-optimality of actions a_0 and a_2 , and indicates that one cannot possibly approach the optimal values unless $\pi(a_1 | s) \approx 1$.

- As further revealed by Lemma 8(ii), one needs to ensure a sufficiently large $\pi(a_1 | s)$ — i.e., $\pi(a_1 | s) \geq (1 - \gamma)/2$ — in order to achieve $V^\pi(s) \gtrapprox 1/2$.
- Lemma 8(iii) establishes an intimate connection between $V^\pi(s)$ and $V^\pi(\bar{s})$: if we hope to attain $V^\pi(\bar{s}) \gtrapprox 1/2$ for an adjoint state \bar{s} , then one needs to first ensure that its associated primary state achieves $V^\pi(s) \gtrapprox 1/2$. The equivalence property (63) allows one to propagate the crossing time of state s to that of state \bar{s} .
- In Lemma 8(iv), we make clear that the Q-functions w.r.t. the buffer states are independent of the policy in use.
- Lemma 8(v) further establishes an intriguing connection between the crossing time of state s and that of the preceding state $s-1$.
- Lemma 8(vi) uncovers that: (a) if $V^\pi(s-2)$ is not sufficiently large, then the Q-value associated with (s, a_0) dominates the one associated with (s, a_2) ; (b) if $V^\pi(s-1)$ is not large enough, then the Q-value associated with (s, a_1) is dominated by that of the other two.

- As indicated by Lemma 8(vii), the sum of the iterate $\theta^{(t)}(s, a)$ over a remains unchanged throughout the execution of the algorithm.

Another key feature that permeates our analysis is a certain monotonicity property of value function estimates as the iteration count t increases, which we discuss in the sequel. To begin with, akin to the monotonicity properties of gradient descent (Beck, 2017), the softmax PG update is known to achieve monotonic performance improvement in a pointwise manner, as summarized in the following lemma. The interested reader is referred to Agarwal et al. (2021, Lemma C.2) for details.

Lemma 9. *Consider the softmax PG method (12). One has*

$$V^{(t+1)}(s) \geq V^{(t)}(s) \quad \text{and} \quad Q^{(t+1)}(s, a) \geq Q^{(t)}(s, a)$$

for any state-action pair (s, a) and any $t \geq 0$, provided that $0 < \eta < (1 - \gamma)^2/5$.

The preceding monotonicity feature, in conjunction with the uniform initialization scheme, ensures non-negativity of value function estimates throughout the execution of the algorithm.

Lemma 10. *Consider the softmax PG method (12), and suppose the initial policy $\pi^{(0)}(\cdot | s)$ for any $s \in \mathcal{S}$ is given by a uniform distribution over the action space \mathcal{A}_s and $0 < \eta < (1 - \gamma)^2/5$. Then one has*

$$\forall t \geq 0, \forall s \in \mathcal{S} : \quad V^{(t)}(s) \geq 0.$$

Proof. The only negative rewards in our constructed MDP are $r(s_1, a_0)$ for $s_1 \in \mathcal{S}_1$ and $r(s_2, a_0)$ for $s_2 \in \mathcal{S}_2$. When $\pi^{(0)}(\cdot | s_1)$ is uniformly distributed, the MDP specification (22) gives

$$\forall s_1 \in \mathcal{S}_1 : \quad V^{(0)}(s_1) = 0.5r(s_1, a_0) + 0.5r(s_1, a_1) = 0.$$

Similarly, one has $V^{(0)}(s_2) = 0$ for all $s_2 \in \mathcal{S}_2$. Applying Lemma 9, we can demonstrate that $V^{(t)}(s) \geq V^{(0)}(s) \geq 0$ for any $s \in \mathcal{S}_1 \cup \mathcal{S}_2$ and any $t \geq 0$. From the Bellman equation, it is easily seen that the value function $V^{(t)}$ of any other state is a linear combination of $\{r(s, a) | s \notin \mathcal{S}_1, s \notin \mathcal{S}_2\}$, $\{V^{(t)}(s_1) | s_1 \in \mathcal{S}_1\}$ and $\{V^{(t)}(s_2) | s_2 \in \mathcal{S}_2\}$, which are all non-negative. It thus follows that $V^{(t)}(s) \geq 0$ for any $s \in \mathcal{S}$ and any $t \geq 0$. \square

A.2 A type of recursive relations

In addition, we make note of a sort of recursive relations that appear commonly when studying the dynamics of gradient descent (Beck, 2017). The proof of the following lemma can be found in Appendix A.5.

Lemma 11. *Consider a positive sequence $\{x_t\}_{t \geq 0}$.*

- (i) *Suppose that $x_t \leq x_{t-1}$ for all $t > 0$. If there exists some quantity $c_l > 0$ obeying $c_l x_0 \leq 1/2$ and*

$$x_t \geq x_{t-1} - c_l x_{t-1}^2 \quad \text{for all } t > 0, \tag{67a}$$

then one has

$$x_t \geq \frac{1}{2c_l t + \frac{1}{x_0}} \quad \text{for all } t \geq 0. \tag{67b}$$

- (ii) *If there exists some quantity $c_u > 0$ obeying*

$$x_t \leq x_{t-1} - c_u x_{t-1}^2 \quad \text{for all } t > 0, \tag{68a}$$

then it follows that

$$x_t \leq \frac{1}{c_u t + \frac{1}{x_0}} \quad \text{for all } t \geq 0. \tag{68b}$$

(iii) Suppose that $0 < x_t < c_x$ for all $t < t_0$ and $x_{t_0} \geq c_x$ for some quantity $c_x > 0$. Assume that

$$x_t \geq x_{t-1} + c_- x_{t-1}^2 \quad \text{for all } 0 < t \leq t_0 \quad (69a)$$

for some quantity $c_- > 0$. Then one necessarily has

$$t_0 \leq \frac{1 + c_- c_x}{c_- x_0}. \quad (69b)$$

(iv) Suppose that

$$0 \leq x_t \leq x_{t-1} + c_+ x_{t-1}^2 \quad \text{for all } 0 < t \leq t_0 \quad (70a)$$

for some quantity $c_+ > 0$. Then one necessarily has

$$t_0 \geq \frac{\frac{1}{x_0} - \frac{1}{x_{t_0}}}{c_+}. \quad (70b)$$

A.3 Proof of Lemma 1

(i) Let us start with state 0. Given that this is an absorbing state and that $r(0, a_0) = 0$, we have $V^*(0) = 0$.

(ii) Next, we turn to the buffer states in \mathcal{S}_1 and \mathcal{S}_2 . For any $s_1 \in \mathcal{S}_1$, the Bellman equation gives

$$Q^*(s_1, a_0) = r(s_1, a_0) + \gamma V^*(0) = -\gamma^2; \quad (71a)$$

$$Q^*(s_1, a_1) = r(s_1, a_1) + \gamma V^*(0) = \gamma^2. \quad (71b)$$

This in turn implies that $V^*(s_1) = Q^*(s_1, a_1) = \gamma^2$. Repeating the same argument, we arrive at $V^*(s_2) = Q^*(s_2, a_1) = r(s_2, a_1) = \gamma^4$ for any $s_2 \in \mathcal{S}_2$.

(iii) We then move on to the adjoint states $\bar{1}$ and $\bar{2}$. From the construction (23), the Bellman equation yields

$$\begin{aligned} Q^*(\bar{1}, a_0) &= r(\bar{1}, a_0) + \gamma V^*(0) = \gamma \tau_1 < \gamma/2, \\ Q^*(\bar{1}, a_1) &= r(\bar{1}, a_1) + \frac{\gamma}{|\mathcal{S}_1|} \sum_{s_1 \in \mathcal{S}_1} V^*(s_1) = \frac{\gamma}{|\mathcal{S}_1|} \sum_{s_1 \in \mathcal{S}_1} V^*(s_1) = \gamma^3, \end{aligned}$$

where the last identity follows since $V^*(s_1) = \gamma^2$. This in turn indicates that $V^*(\bar{1}) = \max\{Q^*(\bar{1}, a_0), Q^*(\bar{1}, a_1)\} = \gamma^3$, provided that $\gamma^2 \geq 1/2$. Similarly, repeating this argument shows that $V^*(\bar{2}) = \gamma^5$, as long as $\gamma^4 \geq 1/2$. As before, the optimal action in state $\bar{1}$ (resp. $\bar{2}$) is a_1 .

(iv) The next step is to determine $V^*(s)$ for any $s \in \{3, \dots, H\}$. Suppose that $V^*(\overline{s-2}) = \gamma^{2s-3}$ and $V^*(\overline{s-1}) = \gamma^{2s-1}$. Then the construction (20) together with the Bellman equation yields

$$\begin{aligned} Q^*(s, a_0) &= r(s, a_0) + \gamma V^*(0) = r_s + \gamma^2 p \tau_{s-2} < 2/3; \\ Q^*(s, a_1) &= r(s, a_1) + \gamma V^*(\overline{s-1}) = \gamma \gamma^{2s-1} = \gamma^{2s}; \\ Q^*(s, a_2) &= r(s, a_2) + \gamma(1-p)V^*(0) + \gamma p V^*(\overline{s-2}) = r_s + p \gamma^{2s-2} < 2/3. \end{aligned}$$

Consequently, one has $V^*(s) = Q^*(s, a_1) = \gamma^{2s}$ — namely, a_1 is the optimal action — as long as $\gamma^{2s} \geq 2/3$.

(v) We then turn attention to $V^*(\bar{s})$ for any $\bar{s} \in \{\bar{3}, \dots, \bar{H}\}$. Suppose that $V^*(s) = \gamma^{2s}$. In view of the construction (21) and the Bellman equation, one has

$$\begin{aligned} Q^*(\bar{s}, a_0) &= r(\bar{s}, a_0) + \gamma V^*(0) = \gamma \tau_s < 1/2; \\ Q^*(\bar{s}, a_1) &= r(\bar{s}, a_1) + \gamma V^*(s) = \gamma^{2s+1}. \end{aligned}$$

Hence, we have $V^*(\bar{s}) = Q^*(\bar{s}, a_1) = \gamma^{2s+1}$ — with the optimal action being a_1 — provided that $\gamma^{2s+1} \geq 1/2$.

(vi) Applying an induction argument based on Steps (iii), (iv) and (v), we conclude that

$$V^*(s) = \gamma^{2s} \quad \text{and} \quad V^*(\bar{s}) = \gamma^{2s+1} \quad (72)$$

for all $3 \leq s \leq H$, with the proviso that $\gamma^{2H} \geq 2/3$ and $\gamma^{2H+1} \geq 1/2$.

(vii) In view of our MDP construction, a negative immediate reward (which is either $-\gamma^2$ or $-\gamma^4$) is accrued only when the current state lies in the buffer sets \mathcal{S}_1 and \mathcal{S}_2 and when action a_0 is executed. However, once a_0 is taken, the MDP will transition to the absorbing state 0, with all subsequent rewards frozen to 0. In conclusion, the entire MDP trajectory cannot receive negative immediate rewards more than once, thus indicating that $Q^\pi(s, a) \geq \min\{-\gamma^2, -\gamma^4\} = -\gamma^2$ irrespective of π and (s, a) .

A.4 Proof of Lemma 8

Proof of Part (i). Before proceeding, we make note of a straightforward fact

$$V^\pi(0) = 0, \quad (73)$$

given that state 0 is an absorbing state and $r(0, a_0) = 0$.

For any $s \in \{3, \dots, H\}$, the construction (20) together with (73) and the Bellman equation yields

$$Q^\pi(s, a_0) = r(s, a_0) + \gamma V^\pi(0) = r_s + \gamma^2 p \tau_{s-2}; \quad (74a)$$

$$Q^\pi(s, a_1) = r(s, a_1) + \gamma V^\pi(s-1) = \gamma V^\pi(s-1); \quad (74b)$$

$$Q^\pi(s, a_2) = r(s, a_2) + \gamma(1-p)V^\pi(0) + \gamma p V^\pi(s-2) = r_s + \gamma p V^\pi(s-2). \quad (74c)$$

Recalling the choices of τ_s , r_s and p in (18), we can continue the derivation in (74a) to reach

$$\begin{aligned} Q^\pi(s, a_0) &= 0.5\gamma^{\frac{2s}{3}+\frac{5}{6}} + c_p(1-\gamma)\gamma^{\frac{2s}{3}+\frac{2}{3}} \\ \implies \gamma^{\frac{3}{2}}\tau_{s-1} &= 0.5\gamma^{\frac{2s}{3}+\frac{5}{6}} \leq Q^\pi(s, a_0) \leq 0.5\gamma^{\frac{2s}{3}+\frac{1}{2}} = \gamma^{\frac{1}{2}}\tau_s. \end{aligned}$$

Here, the last inequality is valid when $c_p \leq 1/6$, given that $\gamma^{\frac{1}{3}} + \frac{1-\gamma}{6}\gamma^{\frac{1}{6}} \leq 1$ holds for any $\gamma < 1$.

In addition, combining (74c) with (72), we arrive at

$$Q^\pi(s, a_2) \leq r_s + \gamma p V^\pi(s-2) = 0.5\gamma^{\frac{2s}{3}+\frac{5}{6}} + c_p(1-\gamma)\gamma^{2s-2} \leq 0.5\gamma^{\frac{2s}{3}+\frac{1}{2}} = \gamma^{\frac{1}{2}}\tau_s.$$

This is guaranteed to hold when $c_p \leq 1/6$, given that $\gamma^{\frac{1}{3}} + \frac{1-\gamma}{3}\gamma^{\frac{4s}{3}-\frac{5}{2}} \leq \gamma^{\frac{1}{3}} + \frac{1-\gamma}{3}\gamma^{\frac{3}{2}} \leq 1$ is valid for all $\gamma < 1$ and $s \geq 3$. Moreover, if one further has $V^\pi(s-2) \geq 0$, then it is seen from (74c) that

$$Q^\pi(s, a_2) \geq r_s = 0.5\gamma^{\frac{2s}{3}+\frac{5}{6}} = \gamma^{\frac{3}{2}}\tau_{s-1}. \quad (75)$$

Proof of Part (ii). By virtue of the construction (20), we can invoke the Bellman equation to show that

$$\begin{aligned} V^\pi(s) &= \pi(a_0 | s)Q^\pi(s, a_0) + \pi(a_1 | s)Q^\pi(s, a_1) + \pi(a_2 | s)Q^\pi(s, a_2) \\ &= \pi(a_1 | s) \cdot \gamma V^\pi(s-1) + \pi(a_0 | s)Q^\pi(s, a_0) + \pi(a_2 | s)Q^\pi(s, a_2) \\ &\leq \pi(a_1 | s)\gamma^{2s} + \{\pi(a_0 | s) + \pi(a_2 | s)\}\gamma^{\frac{1}{2}}\tau_s \\ &= \gamma^{2s}\pi(a_1 | s) + \gamma^{\frac{1}{2}}\tau_s(1 - \pi(a_1 | s)). \end{aligned} \quad (76)$$

Here, the second identity comes from (74b), the penultimate line follows from (59), (72), as well as the facts $V^\pi(s-1) \leq V^\pi(s-1)$, while the last inequality exploits the fact $\pi(a_0 | s) + \pi(a_2 | s) = 1 - \pi(a_1 | s)$.

If $V^\pi(s) \geq \tau_s$, then this together with the upper bound (76) necessarily requires that

$$\tau_s \leq \gamma^{2s}\pi(a_1 | s) + \gamma^{\frac{1}{2}}\tau_s(1 - \pi(a_1 | s)),$$

which is equivalent to saying that

$$\pi(a_1 | s) \geq \frac{\tau_s - \gamma^{\frac{1}{2}}\tau_s}{\gamma^{2s} - \gamma^{\frac{1}{2}}\tau_s} = \frac{1 - \gamma^{\frac{1}{2}}}{2\gamma^{\frac{4s}{3}} - \gamma^{\frac{1}{2}}} \geq \frac{1 - \gamma^{\frac{1}{2}}}{\gamma^{\frac{4s}{3}}} = \frac{1 - \gamma}{\gamma^{\frac{4s}{3}}(1 + \gamma^{\frac{1}{2}})} \geq \frac{1 - \gamma}{2}. \quad (77)$$

Putting these arguments together establishes the advertised result (60).

Proof of Part (iii). For any $\bar{s} \in \{\bar{3}, \dots, \bar{H}\}$, in view of the construction (21) and the Bellman equation, one has

$$\begin{aligned} Q^\pi(\bar{s}, a_0) &= r(\bar{s}, a_0) + \gamma V^\pi(0) = \gamma \tau_s; \\ Q^\pi(\bar{s}, a_1) &= r(\bar{s}, a_1) + \gamma V^\pi(s) = \gamma V^\pi(s). \end{aligned}$$

Regarding state $\bar{1}$, we have

$$\begin{aligned} Q^\pi(\bar{1}, a_0) &= r(\bar{1}, a_0) + \gamma V^\pi(0) = \gamma \tau_1; \\ Q^\pi(\bar{1}, a_1) &= r(\bar{1}, a_1) + \gamma \frac{1}{|\mathcal{S}_1|} \sum_{s' \in \mathcal{S}_1} V^\pi(s') = \gamma V^\pi(1). \end{aligned}$$

Similarly, one obtains $Q^\pi(\bar{2}, a_0) = \gamma \tau_2$ and $Q^\pi(\bar{2}, a_1) = \gamma V^\pi(2)$.

Next, let us decompose $V^\pi(\bar{s})$ as follows:

$$\begin{aligned} V^\pi(\bar{s}) &= \pi(a_0 | \bar{s}) Q^\pi(\bar{s}, a_0) + \pi(a_1 | \bar{s}) Q^\pi(\bar{s}, a_1) \\ &= \gamma \tau_s \pi(a_0 | \bar{s}) + \gamma \pi(a_1 | \bar{s}) V^\pi(s) = \gamma \tau_s + \gamma \pi(a_1 | \bar{s}) (V^\pi(s) - \tau_s), \end{aligned}$$

where we have used $\pi(a_0 | \bar{s}) + \pi(a_1 | \bar{s}) = 1$. From this relation and the assumption $\pi(a_1 | \bar{s}) > 0$, it is straightforward to see that $V^\pi(\bar{s}) \geq \gamma \tau_s$ if and only if $V^\pi(s) \geq \tau_s$. The claim (63) regarding $t_s(\tau_s)$ and $t_{\bar{s}}(\gamma \tau_s)$ then follows directly from the definition of t_s (see (30) and (31)).

Proof of Part (iv). For any $s_1 \in \mathcal{S}_1$, the Bellman equation yields

$$\begin{aligned} Q^\pi(s_1, a_0) &= r(s_1, a_0) + \gamma V^\pi(0) = -\gamma^2 + 0 = -\gamma^2, \\ Q^\pi(s_1, a_1) &= r(s_1, a_1) + \gamma V^\pi(0) = \gamma^2 + 0 = \gamma^2, \end{aligned}$$

and hence

$$V^\pi(s_1) = \pi(a_0 | s_1) Q^\pi(s_1, a_0) + \pi(a_1 | s_1) Q^\pi(s_1, a_1) = -\gamma^2 \pi(a_0 | s_1) + \gamma^2 \pi(a_1 | s_1).$$

A similar argument immediately yields that for any $s_2 \in \mathcal{S}_2$,

$$Q^\pi(s_2, a_0) = -\gamma^4, \quad Q^\pi(s_2, a_1) = \gamma^4, \quad \text{and} \quad V^\pi(s_2) = -\gamma^4 \pi(a_0 | s_2) + \gamma^4 \pi(a_1 | s_2).$$

These together with our notation convention (25) establish (64).

Proof of Part (v). Suppose instead that $V^\pi(s-1) < \tau_{s-1}$. In view of the basic property (62) in Lemma 8, this necessarily requires that

$$V^\pi(\overline{s-1}) < \gamma \tau_{s-1}. \tag{78}$$

Taking (78) together with the relation (59b) allows us to reach

$$Q^\pi(s, a_1) = \gamma V^\pi(\overline{s-1}) < \gamma^2 \tau_{s-1} = \gamma^{\frac{4}{3}} \tau_s. \tag{79}$$

In addition, the properties (59a) and (59c) imply that

$$Q^\pi(s, a_0) < \gamma^{\frac{1}{2}} \tau_s \quad \text{and} \quad Q^\pi(s, a_2) < \gamma^{\frac{1}{2}} \tau_s.$$

Putting everything together implies that

$$V^\pi(s) \leq \max \left\{ Q^\pi(s, a_0), Q^\pi(s, a_1), Q^\pi(s, a_2) \right\} < \gamma^{\frac{1}{2}} \tau_s,$$

which contradicts the assumption $V^\pi(s) \geq \gamma^{\frac{1}{2}} \tau_s$. This establishes the claimed result for any $s \in \{3, \dots, H\}$.

Proof of Part (vi). First, due to explicit expressions of the Q functions (74a) and (74c), one has

$$Q^\pi(s, a_0) - Q^\pi(s, a_2) = \gamma^2 p \tau_{s-2} - \gamma p V^\pi(\overline{s-2}) = \gamma p (\gamma \tau_{s-2} - V^\pi(\overline{s-2})) > 0,$$

where the last relation holds since $V^\pi(\overline{s-2}) < \gamma \tau_{s-2}$ when $V^\pi(s-2) < \tau_{s-2}$ (see (62)).

In addition, following the same derivation as for (79), we see that the condition $V^\pi(s-1) \leq \tau_{s-1}$ implies

$$Q^\pi(s, a_1) \leq \gamma^2 \tau_{s-1}.$$

It is also seen from Part (i) of this lemma that

$$Q^\pi(s, a_0) \geq \gamma^{3/2} \tau_{s-1} \quad \text{and} \quad Q^\pi(s, a_2) \geq \gamma^{3/2} \tau_{s-1},$$

provided that $V^\pi(\overline{s-2}) \geq 0$. Combining these two inequalities, we arrive at the claimed bound

$$\min \{Q^\pi(s, a_0), Q^\pi(s, a_2)\} - Q^\pi(s, a_1) \geq \gamma^{3/2} \tau_{s-1} - \gamma^2 \tau_{s-1} = \frac{\gamma^{3/2} (1 - \gamma) \tau_{s-1}}{1 + \gamma^{1/2}} \geq (1 - \gamma)/8,$$

where the last inequality holds if $\gamma^{2s/3+5/6} \geq \gamma^s \geq 1/2$.

Proof of Part (vii). According to the update rule (12), we have — for any policy π — that

$$\begin{aligned} \sum_a \frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta(s, a)} &= \sum_a \frac{1}{1 - \gamma} d_\mu^{\pi_\theta}(s) \pi_\theta(a | s) (Q^{\pi_\theta}(s, a) - V^{\pi_\theta}(s)) \\ &= \frac{1}{1 - \gamma} d_\mu^{\pi_\theta}(s) \left(\sum_a \pi_\theta(a | s) Q^{\pi_\theta}(s, a) - V^{\pi_\theta}(s) \sum_a \pi_\theta(a | s) \right) = 0, \end{aligned}$$

where we have used the identities $\sum_a \pi_\theta(a | s) = 1$ and $V^\pi(s) = \sum_a \pi(a | s) Q^\pi(s, a)$. As a result, if $\sum_a \theta^{(0)}(s, a) = 0$, then it follows from the PG update rule that $\sum_a \theta^{(t)}(s, a) = 0$.

A.5 Proof of Lemma 11

Proof of Part (i). Dividing both sides of (67a) by $x_t x_{t-1}$, we obtain

$$\frac{1}{x_{t-1}} \geq \frac{1}{x_t} - \frac{c_1 x_{t-1}}{x_t}.$$

If $c_1 x_0 \leq 1/2$, then the monotonicity assumption gives $c_1 x_t \leq 1/2$ for all $t \geq 0$. It then follows that

$$\frac{x_t}{x_{t-1}} \geq 1 - c_1 x_{t-1} \geq \frac{1}{2} \quad \implies \quad \frac{1}{x_{t-1}} \geq \frac{1}{x_t} - \frac{c_1 x_{t-1}}{x_t} \geq \frac{1}{x_t} - 2c_1.$$

Apply this relation recursively to deduce that

$$\frac{1}{x_t} \leq \frac{1}{x_{t-1}} + 2c_1 \leq \dots \leq \frac{1}{x_0} + 2c_1 t.$$

This readily concludes the proof of (67b).

Proof of Part (ii). Similarly, divide both sides of (68a) by $x_t x_{t-1}$ to derive

$$\frac{1}{x_{t-1}} \leq \frac{1}{x_t} - \frac{c_u x_{t-1}}{x_t} \leq \frac{1}{x_t} - c_u,$$

given the monotonicity and positivity assumption $0 < x_t \leq x_{t-1}$. Invoking this inequality recursively gives

$$\frac{1}{x_t} \geq \frac{1}{x_{t-1}} + c_u \geq \dots \geq \frac{1}{x_0} + c_u t,$$

thus establishing the advertised bound (68b).

Proof of Part (iii). We now turn attention to (69b). As is clearly seen, the non-negative sequence $\{x_t\}$ majorizes another sequence $\{y_t\}$ generated as follows (in the sense that $x_t \geq y_t$ for all $0 < t \leq t_0$)

$$y_0 = x_0 \quad \text{and} \quad y_t = y_{t-1} + c_- y_{t-1}^2 \quad \text{for all } 0 < t \leq t_0 \quad (80)$$

Dividing both sides of the second equation of (80) by $y_{t-1}y_t$, we reach

$$\frac{1}{y_{t-1}} = \frac{1}{y_t} + c_- \frac{y_{t-1}}{y_t} \geq \frac{1}{y_t} + \frac{c_-}{1 + c_- c_x}.$$

To see why the last inequality holds, note that, according to the first equation of (80) and the assumption $x_{t-1} < c_x$ (and hence $y_{t-1} \leq x_{t-1} < c_x$), we have

$$\frac{y_t}{y_{t-1}} = 1 + c_- y_{t-1} \leq 1 + c_- c_x.$$

As a result, we can apply the preceding inequalities recursively to derive

$$\frac{1}{y_0} \geq \frac{1}{y_1} + \frac{c_-}{1 + c_- c_x} \geq \dots \geq \frac{1}{y_{t_0}} + \frac{c_-}{1 + c_- c_x} t_0 \geq \frac{c_-}{1 + c_- c_x} t_0,$$

and hence we arrive at (69b),

$$t_0 \leq \frac{1 + c_- c_x}{c_- y_0} = \frac{1 + c_- c_x}{c_- x_0}.$$

Proof of Part (iv). The proof of (70b) is quite similar to that of (69b). Let us construct another non-negative sequence $\{z_t\}$ as follows

$$z_0 = x_0 \quad \text{and} \quad z_t = z_{t-1} + c_+ z_{t-1}^2 \quad \text{for all } 0 < t \leq t_0. \quad (81)$$

Comparing this with (70a) clearly reveals that $z_t \geq x_t$. Divide both sides of (81) by $z_t z_{t-1}$ to reach

$$\frac{1}{z_{t-1}} = \frac{1}{z_t} + c_+ \frac{z_{t-1}}{z_t} \leq \frac{1}{z_t} + c_+,$$

where the last inequality is valid since, by construction, $z_t \geq z_{t-1}$. Applying this relation recursively yields

$$\frac{1}{z_0} \leq \frac{1}{z_{t_0}} + c_+ t_0,$$

which taken together with the fact $z_0 = x_0$ and $z_{t_0} \geq x_{t_0}$ leads to

$$t_0 \geq \frac{\frac{1}{z_0} - \frac{1}{z_{t_0}}}{c_+} \geq \frac{\frac{1}{x_0} - \frac{1}{x_{t_0}}}{c_+}.$$

B Discounted state visitation probability (Lemmas 2-3)

In this section, we establish our bounds concerning the discounted state visitation probability, as claimed in Lemma 2 and Lemma 3. Throughout this section, we denote by $\mathbb{P}(\cdot | \pi)$ the probability distribution when policy π is adopted. Also, we recall that μ is taken to be a uniform distribution over all states.

B.1 Lower bounds: proof of Lemma 2

Consider an arbitrary policy π , and let $\{s^k\}_{k \geq 0}$ represent an MDP trajectory. For any $s \in \{3, \dots, H\}$, it follows from the definition (13) of d_μ^π that

$$d_\mu^\pi(s) = (1 - \gamma) \sum_{k=0}^{\infty} \gamma^k \mathbb{P}(s^k = s | s^0 \sim \mu, \pi) \quad (82)$$

$$\begin{aligned}
&\geq (1 - \gamma)\gamma\mathbb{P}(s^1 = s \mid s^0 \sim \mu, \pi) \geq (1 - \gamma)\gamma \sum_{s' \in \widehat{\mathcal{S}}_s} \mathbb{P}(s^1 = s \mid s^0 = s', \pi) \mathbb{P}(s_0 = s' \mid s^0 \sim \mu) \\
&= (1 - \gamma)\gamma \cdot \frac{|\widehat{\mathcal{S}}_s|}{|\mathcal{S}|} = c_m \gamma (1 - \gamma)^2.
\end{aligned}$$

Here, the penultimate identity is valid due to the construction (24) and the assumption that μ is uniformly distributed, whereas the last identity results from the assumption (17). This establishes (36a). Repeating the same argument also reveals that

$$d_\mu^\pi(\bar{s}) \geq c_m \gamma (1 - \gamma)^2$$

for any $\bar{s} \in \{\bar{1}, \dots, \bar{H}\}$, thus validating the lower bound (36b).

In addition, for any $s \in \mathcal{S}_1$, the MDP construction (24) allows one to derive

$$\begin{aligned}
d_\mu^\pi(s) &= (1 - \gamma) \sum_{k=0}^{\infty} \gamma^k \mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi) \geq \gamma(1 - \gamma) \mathbb{P}(s^1 = s \mid s^0 \sim \mu, \pi) \\
&\geq \gamma(1 - \gamma) \mathbb{P}(s^1 = s \mid s^0 \in \widehat{\mathcal{S}}_1) \mathbb{P}(s_0 \in \widehat{\mathcal{S}}_1 \mid s^0 \sim \mu) \\
&= \gamma(1 - \gamma) \cdot \frac{1}{|\mathcal{S}_1|} \cdot \frac{|\widehat{\mathcal{S}}_1|}{|\mathcal{S}|} = \gamma(1 - \gamma) \frac{c_m}{c_{b,1}} \cdot \frac{1}{|\mathcal{S}|}.
\end{aligned}$$

Here, the last line holds due to the fact that μ is uniformly distributed and the assumptions (16) and (17). We have thus concluded the proof for (36c). The proof for (36d) follows from an identical argument and is hence omitted.

B.2 Upper bounds: proof of Lemma 3

B.2.1 Preliminary facts

Before embarking on the proof, we collect several basic yet useful properties that happen when $t < t_s(\tau_s)$. The first-time readers can proceed directly to Appendix B.2.2.

Properties about $Q^{(t)}(\bar{s}, a)$. Combine the property (61) in Lemma 8 with (32) to yield that: for any $1 \leq s \leq H$ and any $t < t_s(\tau_s)$, one has

$$Q^{(t)}(\bar{s}, a_1) = \gamma V^{(t)}(s) < \gamma \tau_s = Q^{(t)}(\bar{s}, a_0) \quad (83)$$

In addition, combining the property (61) in Lemma 8 with (66) yields: for any $2 \leq s \leq H$,

$$V^{(t)}(\bar{s}') \leq \max \left\{ Q^{(t)}(\bar{s}', a_0), Q^{(t)}(\bar{s}', a_1) \right\} = \max \left\{ \gamma \tau_{s'}, \gamma V^{(t)}(s') \right\} = \gamma \tau_{s'} \quad (84)$$

holds for all s' obeying $s \leq s' \leq H$ and all $t < t_s(\tau_s)$. As a remark, (83) indicates that a_1 remains unfavored (according to the current estimate $Q^{(t)}$) before the iteration number hits $t_s(\tau_s)$.

Properties about $Q^{(t)}(s+1, a)$ and $Q^{(t)}(s+2, a)$. First, combining (84) with the relation (74) reveals that: for any $2 \leq s \leq H-1$ and any $t < t_s(\tau_s)$,

$$Q^{(t)}(s+1, a_0) = r_{s+1} + \gamma^2 p \tau_{s-1} \geq r_{s+1}, \quad (85a)$$

$$Q^{(t)}(s+1, a_1) = \gamma V^{(t)}(\bar{s}) \leq \gamma^2 \tau_s = \gamma^{1/2} r_{s+1}, \quad (85b)$$

$$Q^{(t)}(s+1, a_2) = r_{s+1} + \gamma p V^{(t)}(\overline{s-1}) \geq r_{s+1} \quad (85c)$$

hold as long as $V^{(t)}(\overline{s-1}) \geq 0$ (which is guaranteed by Lemma 10). Similarly, (84) and (74) also give

$$Q^{(t)}(s+2, a_0) = r_{s+2} + \gamma^2 p \tau_s$$

$$\begin{aligned}
Q^{(t)}(s+2, a_1) &= \gamma V^{(t)}(\overline{s+1}) = \gamma^2 \tau_{s+1} = \gamma^{1/2} r_{s+2}, \\
Q^{(t)}(s+2, a_2) &= r_{s+2} + \gamma p V^{(t)}(\bar{s}) \leq r_{s+2} + \gamma^2 p \tau_s
\end{aligned}$$

for any $1 \leq s \leq H-2$ and any $t < t_s(\tau_s)$. Consequently, we have

$$Q^{(t)}(s+1, a_1) \leq \min \{Q^{(t)}(s+1, a_0), Q^{(t)}(s+1, a_2)\}, \quad \text{if } 2 \leq s \leq H-1 \quad (86a)$$

$$Q^{(t)}(s+2, a_2) \leq Q^{(t)}(s+2, a_0), \quad \text{if } 1 \leq s \leq H-2 \quad (86b)$$

for all $t < t_s(\tau_s)$. In other words, the above two inequalities reveal that actions a_1 and a_2 are perceived as suboptimal (based on the current Q-function estimates) before the iteration count surpasses $t_s(\tau_s)$.

Next, consider any $2 \leq s \leq H-1$ and any $t < t_s(\tau_s)$. It has already been shown above that

$$Q^{(t)}(s+1, a) \geq Q^{(t)}(s+1, a_1), \quad a \in \{a_0, a_2\}. \quad (87a)$$

A similar argument also implies that, for any $t < t_s(\tau_s)$,

$$Q^{(t)}(s+2, a_0) \geq Q^{(t)}(s+2, a_2), \quad (87b)$$

which forms another property useful for our subsequent analysis.

B.2.2 Proof of the upper bounds (37a) and (37b)

We now turn attention to upper bounding $d_\mu^{(t)}(s)$ for any $s \in \{3, \dots, H\}$. By virtue of the expansion (82), upper bounding $d_\mu^{(t)}(s)$ requires controlling $\mathbb{P}(s^k = s | s^0 \sim \mu, \pi^{(t)})$ for all $k \geq 0$. In light of this, our analysis consists of (i) developing upper bounds on the inter-related quantities $\mathbb{P}(s^k = s | s^0 \sim \mu, \pi^{(t)})$ and $\mathbb{P}(s^k = \bar{s} | s^0 \sim \mu, \pi^{(t)})$ for any $k \geq 0$, and (ii) combining these upper bounds to control $d_\mu^{(t)}(s)$. At the core of our analysis is the following upper bounds on the t -th policy iterate, which will be established in Appendix B.2.6.

Lemma 12. *Under the assumption (35), for any $2 \leq s \leq H$ and any $t < t_s(\tau_s)$, one has*

$$\pi^{(t)}(a_1 | \bar{s}) \leq \pi^{(t)}(a_0 | \bar{s}) \quad \text{and} \quad \pi^{(t)}(a_1 | \bar{s}) \leq 1/2. \quad (88a)$$

Furthermore,

$$\pi^{(t)}(a_1 | s+1) \leq \min \{\pi^{(t)}(a_0 | s+1), \pi^{(t)}(a_2 | s+1)\} \quad \text{and} \quad \pi^{(t)}(a_1 | s+1) \leq 1/3 \quad (88b)$$

hold if $2 \leq s \leq H-1$, and

$$\pi^{(t)}(a_2 | s+2) \leq \pi^{(t)}(a_0 | s+2) \quad \text{and} \quad \pi^{(t)}(a_2 | s+2) \leq 1/2 \quad (88c)$$

hold if $1 \leq s \leq H-2$.

In words, Lemma 12 posits that, at the beginning, the policy iterate $\pi^{(t)}$ does not assign too much probability mass on actions that are currently perceived as suboptimal (see the remarks in Appendix B.2.1). With this lemma in place, we are positioned to establish the advertised upper bound.

Step 1: bounding $\mathbb{P}(s^k = s | s^0 \sim \mu, \pi^{(t)})$. For any $t < t_s(\tau_s)$ and any $s \in \{3, \dots, H\}$, making use of the upper bound (88a) and the MDP construction in Section 3 yields

$$\begin{aligned}
\mathbb{P}(s^0 = s | s^0 \sim \mu, \pi^{(t)}) &= 1/|\mathcal{S}|, \\
\mathbb{P}(s^1 = s | s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^0 \in \widehat{\mathcal{S}}_s) + \pi^{(t)}(a_1 | \bar{s}) \mathbb{P}(s^0 = \bar{s}) \leq \frac{|\widehat{\mathcal{S}}_s|}{|\mathcal{S}|} + \frac{1}{2} \cdot \frac{1}{|\mathcal{S}|} \leq \frac{2|\widehat{\mathcal{S}}_s|}{|\mathcal{S}|} = 2c_m(1-\gamma), \\
\mathbb{P}(s^k = s | s^0 \sim \mu, \pi^{(t)}) &= \pi^{(t)}(a_1 | \bar{s}) \mathbb{P}(s^{k-1} = \bar{s} | s^0 \sim \mu, \pi^{(t)}) \leq \frac{1}{2} \mathbb{P}(s^{k-1} = \bar{s} | s^0 \sim \mu, \pi^{(t)})
\end{aligned}$$

for all $k \geq 2$. Note that the above calculation exploits the fact that μ is a uniform distribution.

Step 2: bounding $\mathbb{P}(s^k = \bar{s} \mid s^0 \sim \mu, \pi^{(t)})$. Given that μ is a uniform distribution, one has

$$\mathbb{P}(s^0 = \bar{s} \mid s^0 \sim \mu, \pi^{(t)}) = 1/|\mathcal{S}| \quad (90a)$$

for any $s \in \mathcal{S}$. With (88b) and (88c) in mind, the MDP construction in Section 3 allows one to show that

$$\begin{aligned} \mathbb{P}(s^1 = \bar{s} \mid s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^0 \in \widehat{\mathcal{S}}_{\bar{s}}) + \pi^{(t)}(a_1 \mid s+1) \mathbb{P}(s^0 = s+1) + \pi^{(t)}(a_2 \mid s+2) \mathbb{P}(s^0 = s+2) \\ &\leq \frac{|\widehat{\mathcal{S}}_{\bar{s}}|}{|\mathcal{S}|} + \frac{1}{3|\mathcal{S}|} + \frac{1}{2|\mathcal{S}|} \leq \frac{2|\widehat{\mathcal{S}}_{\bar{s}}|}{|\mathcal{S}|} = 2c_m(1-\gamma) \end{aligned} \quad (90b)$$

holds for any $2 \leq s \leq H-2$ and any $t < t_s(\tau_s)$, and in addition,

$$\begin{aligned} \mathbb{P}(s^k = \bar{s} \mid s_0 \sim \mu, \pi^{(t)}) &\leq \pi^{(t)}(a_1 \mid s+1) \mathbb{P}(s^{k-1} = s+1 \mid s^0 \sim \mu, \pi^{(t)}) + \pi^{(t)}(a_2 \mid s+2) \mathbb{P}(s^{k-1} = s+2 \mid s^0 \sim \mu, \pi^{(t)}) \\ &\leq \frac{1}{3} \mathbb{P}(s^{k-1} = s+1 \mid s_0 \sim \mu, \pi^{(t)}) + \frac{1}{2} \mathbb{P}(s^{k-1} = s+2 \mid s^0 \sim \mu, \pi^{(t)}) \end{aligned} \quad (90c)$$

hold for any $k \geq 2$, $2 \leq s \leq H-2$, and any $t < t_s(\tau_s)$. Moreover, invoking (88b) and the MDP construction once again reveals that

$$\begin{aligned} \mathbb{P}(s^k = \overline{H-1} \mid s^0 \sim \mu, \pi^{(t)}) &\leq \pi^{(t)}(a_1 \mid H) \mathbb{P}(s^{k-1} = H \mid s^0 \sim \mu, \pi^{(t)}) \leq \frac{1}{3} \mathbb{P}(s^{k-1} = H \mid s^0 \sim \mu, \pi^{(t)}) \\ \mathbb{P}(s^k = \overline{H} \mid s^0 \sim \mu, \pi^{(t)}) &= 0 \end{aligned}$$

hold for any $k \geq 2$ and any $t < t_s(\tau_s)$. In addition, it is seen that

$$\begin{aligned} \mathbb{P}(s^1 = \overline{H-1} \mid s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^0 \in \widehat{\mathcal{S}}_{\overline{H-1}}) + \mathbb{P}(s^0 = H) = \frac{|\widehat{\mathcal{S}}_{\overline{H-1}}|}{|\mathcal{S}|} + \frac{1}{|\mathcal{S}|} \leq \frac{2|\widehat{\mathcal{S}}_{\overline{H-1}}|}{|\mathcal{S}|} = 2c_m(1-\gamma), \\ \mathbb{P}(s^1 = \overline{H} \mid s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^0 \in \widehat{\mathcal{S}}_{\overline{H}}) = \frac{|\widehat{\mathcal{S}}_{\overline{H}}|}{|\mathcal{S}|} = c_m(1-\gamma). \end{aligned}$$

Step 3: putting all this together. Combining the preceding upper bounds on both $\mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)})$ and $\mathbb{P}(s^k = \bar{s} \mid s^0 \sim \mu, \pi^{(t)})$ ($k \geq 1$) and recognizing the monotonicity property (33), we immediately arrive at the following crude bounds

$$\begin{aligned} \max_{3 \leq s \leq H, t < t_s(\tau_s)} \left\{ \mathbb{P}(s^0 = s \mid s^0 \sim \mu, \pi^{(t)}), \mathbb{P}(s^1 = s \mid s^0 \sim \mu, \pi^{(t)}) \right\} &\leq 1/|\mathcal{S}| \leq 2c_m(1-\gamma) \\ \max_{2 \leq s \leq H, t < t_s(\tau_s)} \left\{ \mathbb{P}(s^0 = \bar{s} \mid s^0 \sim \mu, \pi^{(t)}), \mathbb{P}(s^1 = \bar{s} \mid s^0 \sim \mu, \pi^{(t)}) \right\} &\leq 2c_m(1-\gamma) \\ \max_{3 \leq s \leq H, t < t_s(\tau_s)} \mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)}) &\leq \frac{5}{6} \max_{2 \leq s \leq H, t < t_s(\tau_s)} \mathbb{P}(s^{k-1} = \bar{s} \mid s^0 \sim \mu, \pi^{(t)}) \\ \max_{2 \leq s \leq H, t < t_s(\tau_s)} \mathbb{P}(s^k = \bar{s} \mid s^0 \sim \mu, \pi^{(t)}) &\leq \frac{5}{6} \max_{3 \leq s \leq H, t < t_s(\tau_s)} \mathbb{P}(s^{k-1} = s \mid s^0 \sim \mu, \pi^{(t)}) \end{aligned}$$

for any $k \geq 2$. It is then straightforward to deduce that

$$\max_{3 \leq s \leq H, t < t_s(\tau_s)} \mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)}) \leq \left(\frac{5}{6}\right)^{k-1} 2c_m(1-\gamma) \quad (91a)$$

$$\max_{2 \leq s \leq H, t < t_s(\tau_s)} \mathbb{P}(s^k = \bar{s} \mid s^0 \sim \mu, \pi^{(t)}) \leq \left(\frac{5}{6}\right)^{k-1} 2c_m(1-\gamma) \quad (91b)$$

for any $k \geq 1$. In turn, these bounds give rise to

$$d_\mu^{(t)}(s) = (1-\gamma) \sum_{k=0}^{\infty} \gamma^k \mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)}) \leq (1-\gamma) \left\{ 2c_m(1-\gamma) + \sum_{k=1}^{\infty} \left(\frac{5}{6}\right)^{k-1} 2c_m(1-\gamma) \right\}$$

$$\leq 2c_m(1-\gamma)^2 + \frac{1}{1-5/6} \cdot 2c_m(1-\gamma)^2 = 14c_m(1-\gamma)^2 \quad (92a)$$

for any $3 \leq s \leq H$ and any $t < t_s(\tau_s)$. This establishes the claimed upper bound (37a) as long as Lemma 12 is valid. Further, replacing s with \bar{s} in (92) also reveals that

$$d_\mu^{(t)}(\bar{s}) \leq 14c_m(1-\gamma)^2 \quad (92b)$$

for any $2 \leq s \leq H$ and any $t < t_s(\tau_s)$, thus concluding the proof of (37b).

B.2.3 Proof of the upper bound (37c)

We now consider any $s \in \mathcal{S}_2$. From our MDP construction, we have

$$\begin{aligned} \mathbb{P}(s^0 = s \mid s^0 \sim \mu, \pi^{(t)}) &= 1/|\mathcal{S}|, \\ \mathbb{P}(s^1 = s \mid s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^1 = s \mid s^0 \in \widehat{\mathcal{S}}_2, \pi^{(t)}) \mathbb{P}(s^0 \in \widehat{\mathcal{S}}_2) + \mathbb{P}(s^1 = s \mid s^0 = \bar{2}, \pi^{(t)}) \mathbb{P}(s^0 = \bar{2}) \\ &\leq \frac{1}{|\mathcal{S}_2|} \frac{|\widehat{\mathcal{S}}_2|}{|\mathcal{S}|} + \frac{1}{|\mathcal{S}_2|} \frac{1}{|\mathcal{S}|} \leq \frac{2|\widehat{\mathcal{S}}_2|}{|\mathcal{S}_2||\mathcal{S}|} = \frac{2c_m}{c_{b,2}|\mathcal{S}|}, \\ \mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^k = s \mid s^{k-1} = \bar{2}, \pi^{(t)}) \mathbb{P}(s^{k-1} = \bar{2} \mid s^0 \sim \mu, \pi^{(t)}) \\ &\leq \frac{1}{2|\mathcal{S}_2|} \mathbb{P}(s^{k-1} = \bar{2} \mid s^0 \sim \mu, \pi^{(t)}) \end{aligned}$$

for any $k \geq 2$ and any $s \in \mathcal{S}_2$. In addition, our bound in (91b) gives

$$\mathbb{P}(s^{k-1} = \bar{2} \mid s^0 \sim \mu, \pi^{(t)}) \leq \left(\frac{5}{6}\right)^{k-2} 2c_m(1-\gamma)$$

for any $k \geq 2$ and any $t < t_2(\tau_2)$. Consequently, we arrive at

$$\mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)}) \leq \frac{1}{|\mathcal{S}_2|} \mathbb{P}(s^{k-1} = \bar{2} \mid s^0 \sim \mu, \pi^{(t)}) \leq \frac{c_m(1-\gamma)}{|\mathcal{S}_2|} \left(\frac{5}{6}\right)^{k-2} = \frac{c_m}{c_{b,2}|\mathcal{S}|} \left(\frac{5}{6}\right)^{k-2}. \quad (94)$$

Armed with the preceding inequalities, we can derive

$$\begin{aligned} d_\mu^{(t)}(s) &= (1-\gamma) \sum_{k=0}^{\infty} \gamma^k \mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)}) \\ &\leq (1-\gamma) \left\{ \frac{1}{|\mathcal{S}|} + \gamma \cdot \frac{2c_m}{c_{b,2}|\mathcal{S}|} + \sum_{k=2}^{\infty} \gamma^k \frac{c_m}{c_{b,2}|\mathcal{S}|} \left(\frac{5}{6}\right)^{k-2} \right\} \\ &\leq \frac{1-\gamma}{|\mathcal{S}|} \left(1 + \frac{2c_m}{c_{b,2}}\right) + \frac{c_m(1-\gamma)}{(1-5/6)c_{b,2}|\mathcal{S}|} = \frac{1-\gamma}{|\mathcal{S}|} \left(1 + \frac{8c_m}{c_{b,2}}\right) \end{aligned}$$

for any $s \in \mathcal{S}_2$ and any $t < t_2(\tau_2)$, thus concluding the advertised upper bound for $s \in \mathcal{S}_2$.

B.2.4 Proof of the upper bound (37d)

It follows from our MDP construction that

$$\begin{aligned} \mathbb{P}(s^0 = \bar{1} \mid s^0 \sim \mu, \pi^{(t)}) &= 1/|\mathcal{S}|, \\ \mathbb{P}(s^1 = \bar{1} \mid s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^0 \in \widehat{\mathcal{S}}_1) + \mathbb{P}(s^0 = 3) = \frac{|\widehat{\mathcal{S}}_1|}{|\mathcal{S}|} + \frac{1}{|\mathcal{S}|}. \end{aligned}$$

Moreover, for any $k \geq 2$ and any $t < t_3(\tau_3)$, one can derive

$$\mathbb{P}(s^k = \bar{1} \mid s^0 \sim \mu, \pi^{(t)}) = \pi^{(t)}(a_2 \mid 3) \mathbb{P}(s^{k-1} = 3 \mid s^0 \sim \mu, \pi^{(t)}) \leq \left(\frac{5}{6}\right)^{k-2} 2c_m(1-\gamma), \quad (95)$$

where the last inequality arises from (91a). Putting these bounds together leads to

$$\begin{aligned} d_\mu^{(t)}(\bar{1}) &= (1 - \gamma) \sum_{k=0}^{\infty} \gamma^k \mathbb{P}(s^k = \bar{1} \mid s^0 \sim \mu, \pi^{(t)}) \leq (1 - \gamma) \left\{ \frac{1}{|\mathcal{S}|} + \gamma \left(\frac{|\widehat{\mathcal{S}}_1|}{|\mathcal{S}|} + \frac{1}{|\mathcal{S}|} \right) + \sum_{k=2}^{\infty} \left(\frac{5}{6} \right)^{k-2} 2c_m(1 - \gamma) \right\} \\ &\leq (1 - \gamma) \left\{ \frac{2|\widehat{\mathcal{S}}_1|}{|\mathcal{S}|} + \frac{1}{1 - 5/6} 2c_m(1 - \gamma) \right\} = 14c_m(1 - \gamma)^2, \end{aligned}$$

where we have used the assumption that $|\widehat{\mathcal{S}}_1| = c_m(1 - \gamma)|\mathcal{S}|$. When $t < t_2(\tau_2)$, the monotonicity property (33) indicates that $t < t_3(\tau_3)$, thus concluding the proof of (37d).

B.2.5 Proof of the upper bound (37e)

In view of our MDP construction, for any $s \in \mathcal{S}_1$ and any $t < \min\{t_1(\tau_1), t_2(\tau_2)\}$ we have

$$\begin{aligned} \mathbb{P}(s^0 = s \mid s^0 \sim \mu, \pi^{(t)}) &= 1/|\mathcal{S}|, \\ \mathbb{P}(s^1 = s \mid s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^1 = s \mid s^0 \in \widehat{\mathcal{S}}_1, \pi^{(t)}) \mathbb{P}(s^0 \in \widehat{\mathcal{S}}_1) + \mathbb{P}(s^1 = s \mid s^0 = \bar{1}, \pi^{(t)}) \mathbb{P}(s^0 = \bar{1}) \\ &\leq \frac{1}{|\mathcal{S}_1|} \frac{|\widehat{\mathcal{S}}_1|}{|\mathcal{S}|} + \frac{1}{|\mathcal{S}_1|} \frac{1}{|\mathcal{S}|} \leq \frac{1}{|\mathcal{S}_1|} \frac{2|\widehat{\mathcal{S}}_1|}{|\mathcal{S}|} = \frac{2c_m}{c_{b,1}|\mathcal{S}|}, \\ \mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)}) &\leq \mathbb{P}(s^k = s \mid s^{k-1} = \bar{1}, \pi^{(t)}) \mathbb{P}(s^{k-1} = \bar{1} \mid s^0 \sim \mu, \pi^{(t)}) \\ &\leq \frac{1}{|\mathcal{S}_1|} \mathbb{P}(s^{k-1} = \bar{1} \mid s^0 \sim \mu, \pi^{(t)}) \leq \frac{2c_m}{c_{b,1}|\mathcal{S}|} \left(\frac{5}{6} \right)^{k-3}, \end{aligned}$$

where k is any integer obeying $k \geq 2$. Here, the last inequality comes from (95). These bounds taken collectively demonstrate that

$$\begin{aligned} d_\mu^{(t)}(s) &= (1 - \gamma) \sum_{k=0}^{\infty} \gamma^k \mathbb{P}(s^k = s \mid s^0 \sim \mu, \pi^{(t)}) \\ &\leq (1 - \gamma) \left\{ \frac{1}{|\mathcal{S}|} + \gamma \cdot \frac{2c_m}{c_{b,1}|\mathcal{S}|} + \sum_{k=2}^{\infty} \gamma^k \frac{2c_m}{c_{b,1}|\mathcal{S}|} \left(\frac{5}{6} \right)^{k-3} \right\} \\ &\leq \frac{1 - \gamma}{|\mathcal{S}|} \left(1 + \frac{2c_m}{c_{b,1}} \right) + \frac{\frac{6}{5} \cdot 2c_m(1 - \gamma)}{(1 - 5/6)c_{b,1}|\mathcal{S}|} \leq \frac{1 - \gamma}{|\mathcal{S}|} \left(1 + \frac{17c_m}{c_{b,1}} \right) \end{aligned}$$

for any $s \in \mathcal{S}_1$ and any $t < \min\{t_1(\tau_1), t_2(\tau_2)\}$. This completes the proof.

B.2.6 Proof of Lemma 12

In order to prove this lemma, we are in need of the following auxiliary result, whose proof can be found in Appendix B.2.7.

Lemma 13. *Consider any state $1 \leq s \leq H$. Suppose that $0 < \eta \leq (1 - \gamma)/2$.*

(i) *If the following conditions*

$$\begin{aligned} Q^{(t)}(s, a_0) - Q^{(t)}(s, a_1) &\geq 0, & Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) &\geq 0 \\ \pi^{(t-1)}(a_1 \mid s) &\leq \min \{ \pi^{(t-1)}(a_0 \mid s), \pi^{(t-1)}(a_2 \mid s) \} \end{aligned}$$

hold, then one has $\pi^{(t)}(a_1 \mid s) \leq 1/3$ and $\pi^{(t)}(a_1 \mid s) \leq \min \{ \pi^{(t)}(a_0 \mid s), \pi^{(t)}(a_2 \mid s) \}$.

(ii) *If the following conditions*

$$Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) \geq 0 \quad \text{and} \quad \pi^{(t-1)}(a_2 \mid s) \leq \pi^{(t-1)}(a_0 \mid s)$$

hold, then one has $\pi^{(t)}(a_2 \mid s) \leq 1/2$ and $\pi^{(t)}(a_2 \mid s) \leq \pi^{(t)}(a_0 \mid s)$.

(iii) If the following conditions

$$Q^{(t)}(\bar{s}, a_0) - Q^{(t)}(\bar{s}, a_1) \geq 0 \quad \text{and} \quad \pi^{(t-1)}(a_1 | \bar{s}) \leq \pi^{(t-1)}(a_0 | \bar{s})$$

hold, then one has $\pi^{(t)}(a_1 | \bar{s}) \leq 1/2$ and $\pi^{(t)}(a_1 | \bar{s}) \leq \pi^{(t)}(a_0 | \bar{s})$.

Remark 9. In words, Lemma 13 develops nontrivial upper bounds on the policy associated with actions that are currently perceived as suboptimal. As we shall see, such upper bounds — which are strictly below 1 — translate to some contraction factors that enable the advertised result of this lemma.

With Lemma 13 in place, we proceed to prove Lemma 12 by induction. Let us start from the base case with $t = 0$. Given that the initial policy is chosen to be uniformly distributed, we have

$$\begin{aligned} \pi^{(0)}(a_1 | s) &= \pi^{(0)}(a_0 | s) = \pi^{(0)}(a_2 | s), & 3 \leq s \leq H; \\ \pi^{(0)}(a_1 | \bar{s}) &= \pi^{(0)}(a_0 | \bar{s}), & 1 \leq s \leq H. \end{aligned}$$

Therefore, the claim (88) trivially holds for $t = 0$.

Next, we move on to the induction step. Suppose that the induction hypothesis (88) holds for the t -th iteration, and we intend to establish it for the $(t + 1)$ -th iteration. Apply Lemma 13 with Conditions (83) and (88a) to yield

$$\pi^{(t+1)}(a_1 | \bar{s}) \leq \pi^{(t+1)}(a_0 | \bar{s})$$

with the proviso that $0 < \eta \leq (1 - \gamma)/2$. Clearly, this also implies that $\pi^{(t+1)}(a_1 | \bar{s}) \leq 1/2$. Further, invoke Lemma 13 once again with Condition (87) and the induction hypothesis (88) to arrive at

$$\begin{aligned} \pi^{(t+1)}(a_1 | s + 1) &\leq \min \{ \pi^{(t+1)}(a_0 | s + 1), \pi^{(t+1)}(a_2 | s + 1) \}, & \text{if } 2 \leq s \leq H - 1; \\ \pi^{(t+1)}(a_2 | s + 2) &\leq \pi^{(t+1)}(a_0 | s + 2), & \text{if } 1 \leq s \leq H - 2. \end{aligned}$$

A straightforward consequence is $\pi^{(t+1)}(a_1 | s + 1) \leq 1/3$ and $\pi^{(t+1)}(a_2 | s + 2) \leq 1/2$. The proof is thus complete by induction.

B.2.7 Proof of Lemma 13

First of all, suppose that $Q^{(t)}(s, a_0) - Q^{(t)}(s, a_1) \geq 0$ and $\pi^{(t-1)}(a_0 | s) \geq \pi^{(t-1)}(a_1 | s)$ hold true. Combining this result with the PG update rule (12) gives

$$\begin{aligned} \theta^{(t)}(s, a_1) &= \theta^{(t-1)}(s, a_1) + \frac{\eta}{1 - \gamma} d_\mu^{(t-1)}(s) \pi^{(t-1)}(a_1 | s) A^{(t-1)}(s, a_1) \\ &\leq \theta^{(t-1)}(s, a_1) + \frac{\eta}{1 - \gamma} d_\mu^{(t-1)}(s) \pi^{(t-1)}(a_1 | s) A^{(t-1)}(s, a_0). \end{aligned}$$

Consequently, applying this inequality and using the PG update rule (12) yield

$$\begin{aligned} &\theta^{(t)}(s, a_1) - \theta^{(t)}(s, a_0) \\ &\leq \theta^{(t-1)}(s, a_1) + \frac{\eta}{1 - \gamma} d_\mu^{(t-1)}(s) \pi^{(t-1)}(a_1 | s) A^{(t-1)}(s, a_0) \\ &\quad - \theta^{(t-1)}(s, a_0) - \frac{\eta}{1 - \gamma} d_\mu^{(t-1)}(s) \pi^{(t-1)}(a_0 | s) A^{(t-1)}(s, a_0) \\ &\leq \left\{ \theta^{(t-1)}(s, a_1) - \theta^{(t-1)}(s, a_0) \right\} + \left\{ \pi^{(t-1)}(a_0 | s) - \pi^{(t-1)}(a_1 | s) \right\} \left| \frac{\eta}{1 - \gamma} d_\mu^{(t-1)}(s) A^{(t-1)}(s, a_0) \right|, \end{aligned} \quad (97)$$

where the last line arises by combining terms and invoking the assumption $\pi^{(t-1)}(a_0 | s) \geq \pi^{(t-1)}(a_1 | s)$.

Additionally, it is seen from the definition of the advantage function that

$$|A^{(t-1)}(s, a_0)| \leq \max_{\pi, a} |Q^\pi(s, a)| + \max_{\pi} |V^\pi(s)| \leq 2, \quad (98)$$

where the last inequality follows from Lemma 1. Recognizing that $d_\mu^{(t-1)}(s) \leq 1$, one obtains

$$\left| \frac{\eta}{1 - \gamma} d_\mu^{(t-1)}(s) A^{(t-1)}(s, a_0) \right| \leq \frac{\eta}{1 - \gamma} \cdot 2 \leq 1, \quad (99)$$

with the proviso that $0 < \eta \leq (1 - \gamma)/2$.

Substituting (99) into (97) then yields

$$\begin{aligned} (97) &\leq \left\{ \theta^{(t-1)}(s, a_1) - \theta^{(t-1)}(s, a_0) \right\} + \left\{ \pi^{(t-1)}(a_0 | s) - \pi^{(t-1)}(a_1 | s) \right\} \\ &\leq \left\{ \theta^{(t-1)}(s, a_1) - \theta^{(t-1)}(s, a_0) \right\} - \left\{ \theta^{(t-1)}(s, a_0) - \theta^{(t-1)}(s, a_1) \right\} = 0, \end{aligned} \quad (100)$$

where both the first line and the last identity rely on the fact that $\theta^{(t-1)}(s, a_1) \leq \theta^{(t-1)}(s, a_0)$ — an immediate consequence of the assumption $\pi^{(t-1)}(a_1 | s) \leq \pi^{(t-1)}(a_0 | s)$. To see why the inequality (100) holds, it suffices to make note of the following consequence of softmax parameterization:

$$\begin{aligned} \pi^{(t-1)}(a_0 | s) - \pi^{(t-1)}(a_1 | s) &= \pi^{(t-1)}(a_1 | s) \left\{ \exp[\theta^{(t-1)}(s, a_0) - \theta^{(t-1)}(s, a_1)] - 1 \right\} \\ &\stackrel{(a)}{\leq} \frac{\exp[\theta^{(t-1)}(s, a_0) - \theta^{(t-1)}(s, a_1)] - 1}{\exp[\theta^{(t-1)}(s, a_0) - \theta^{(t-1)}(s, a_1)] + 1} \\ &\stackrel{(b)}{\leq} \theta^{(t-1)}(s, a_0) - \theta^{(t-1)}(s, a_1), \end{aligned}$$

where (b) follows since $\frac{e^x - 1}{e^x + 1} \leq x$ for all $x \geq 0$, and the validity of (a) is guaranteed since

$$\begin{aligned} \pi^{(t-1)}(a_1 | s) &= \frac{\exp(\theta^{(t-1)}(s, a_1))}{\sum_a \exp(\theta^{(t-1)}(s, a))} \leq \frac{\exp(\theta^{(t-1)}(s, a_1))}{\exp(\theta^{(t-1)}(s, a_0)) + \exp(\theta^{(t-1)}(s, a_1))} \\ &= \frac{1}{\exp[\theta^{(t-1)}(s, a_0) - \theta^{(t-1)}(s, a_1)] + 1}. \end{aligned}$$

To conclude, the above result (100) implies that

$$\pi^{(t)}(a_0 | s) \geq \pi^{(t)}(a_1 | s). \quad (101)$$

Repeating the above argument immediately reveals that: if

$$Q^{(t-1)}(s, a_2) \geq Q^{(t-1)}(s, a_1) \quad \text{and} \quad \pi^{(t-1)}(a_2 | s) \geq \pi^{(t-1)}(a_1 | s),$$

then one has $\pi^{(t)}(a_2 | s) \geq \pi^{(t)}(a_1 | s)$, which together with (101) indicates that

$$\begin{aligned} \pi^{(t)}(a_1 | s) &\leq \min \{ \pi^{(t)}(a_0 | s), \pi^{(t)}(a_2 | s) \} \\ \implies \pi^{(t)}(a_1 | s) &\leq \frac{\pi^{(t)}(a_0 | s) + \pi^{(t)}(a_1 | s) + \pi^{(t)}(a_2 | s)}{3} = \frac{1}{3}. \end{aligned}$$

This establishes Part (i) of Lemma 13.

The proofs of Parts (ii) and (iii) follow from exactly the same argument as for Part (i), and are hence omitted for the sake of brevity.

C Crossing times of the first few states (Lemma 4)

This section presents the proof of Lemma 4 regarding the crossing times w.r.t. \mathcal{S}_1 , \mathcal{S}_2 , and state $\bar{1}$.

C.1 Crossing times for the buffer states in \mathcal{S}_1 and \mathcal{S}_2

We first present the proof of the relation (38a) regarding several quantities about t_1 and t_2 .

Step 1: characterize the policy gradients. Our analysis largely relies on understanding the policy gradient dynamics, towards which we need to first characterize the gradient. Recalling that the gradient of $V^{(t)}$ w.r.t. $\theta_t(1, a_1)$ (cf. (12b)) is given by

$$\begin{aligned}\frac{\partial V^{(t)}(\mu)}{\partial \theta(1, a_1)} &= \frac{1}{1-\gamma} d_\mu^{(t)}(1) \pi^{(t)}(a_1 | 1) \left\{ Q^{(t)}(1, a_1) - V^{(t)}(1) \right\} \\ &= \frac{1}{1-\gamma} d_\mu^{(t)}(1) \pi^{(t)}(a_1 | 1) \left\{ Q^{(t)}(1, a_1) - \pi^{(t)}(a_0 | 1) Q^{(t)}(1, a_0) - \pi^{(t)}(a_1 | 1) Q^{(t)}(1, a_1) \right\} \\ &= \frac{1}{1-\gamma} d_\mu^{(t)}(1) \pi^{(t)}(a_1 | 1) \pi^{(t)}(a_0 | 1) \left\{ Q^{(t)}(1, a_1) - Q^{(t)}(1, a_0) \right\} \\ &= \frac{2\gamma^2}{1-\gamma} d_\mu^{(t)}(1) \pi^{(t)}(a_1 | 1) \pi^{(t)}(a_0 | 1) > 0,\end{aligned}\tag{102}$$

where in the last step we use $Q^{(t)}(1, a_1) - Q^{(t)}(1, a_0) = 2\gamma^2$ (see (64)). The same calculation also yields

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(2, a_1)} = \frac{2\gamma^4}{1-\gamma} d_\mu^{(t)}(2) \pi^{(t)}(a_1 | 2) \pi^{(t)}(a_0 | 2) > 0.\tag{103}$$

As an immediate consequence, the PG update rule (12a) reveals that both $\theta^{(t)}(1, a_1)$ (resp. $\pi^{(t)}(1, a_1)$) and $\theta^{(t)}(2, a_1)$ (resp. $\pi^{(t)}(2, a_1)$) are monotonically increasing with t throughout the execution of the algorithm, which together with the initial condition $\pi^{(0)}(a_0 | 1) = \pi^{(0)}(a_1 | 1) = \pi^{(0)}(a_0 | 2) = \pi^{(0)}(a_1 | 2)$ as well as the identities $\theta^{(t)}(1, a_1) = -\theta^{(t)}(1, a_0)$ and $\theta^{(t)}(2, a_1) = -\theta^{(t)}(2, a_0)$ (due to (65)) gives

$$\pi^{(t)}(a_0 | 1) \leq \pi^{(t)}(a_1 | 1) \quad \text{and} \quad \pi^{(t)}(a_0 | 2) \leq \pi^{(t)}(a_1 | 2) \quad \text{for all } t \geq 0.\tag{104}$$

Step 2: determine the range of $\pi^{(t)}(\cdot | 1)$ and $\pi^{(t)}(\cdot | 2)$. From the basic property (64), the value function of the buffer states in \mathcal{S}_1 — abbreviated by $V^{(t)}(1)$ as in the notation convention (25) — satisfies

$$V^{(t)}(1) = -\gamma^2 \pi^{(t)}(a_0 | 1) + \gamma^2 \pi^{(t)}(a_1 | 1) = -\gamma^2 + 2\gamma^2 \pi^{(t)}(a_1 | 1),\tag{105}$$

given that $\pi^{(t)}(a_0 | 1) + \pi^{(t)}(a_1 | 1) = 1$. Therefore, for any $t < t_1(\gamma^2 - 1/4)$ — which means $V^{(t)}(1) < \gamma^2 - 1/4$ according to the definition (31) — one has the following upper bound:

$$V^{(t)}(1) = -\gamma^2 + 2\gamma^2 \pi^{(t)}(a_1 | 1) < \gamma^2 - 1/4.$$

This is equivalent to requiring that

$$\pi^{(t)}(a_1 | 1) < 1 - (8\gamma^2)^{-1} \leq 7/8\tag{106}$$

and, consequently, $\pi^{(t)}(a_0 | 1) = 1 - \pi^{(t)}(a_1 | 1) \geq 1/8$ for any $t < t_1(\gamma^2 - 1/4)$. Putting this and (104) together further implies — for every $t < t_1(\gamma^2 - 1/4)$ — that:

$$1/8 \leq \pi^{(t)}(a_0 | 1) \leq \pi^{(t)}(a_1 | 1) \leq 7/8.\tag{107}$$

Step 3: determine the range of policy gradients. In addition to showing the non-negativity of $\frac{\partial V^{(t)}(\mu)}{\partial \theta(1, a_1)}$ and $\frac{\partial V^{(t)}(\mu)}{\partial \theta(2, a_1)}$ for all $t \geq 0$, we are also in need of bounding their magnitudes. Towards this, invoke the property (107) to bound the derivative (102) by

$$\frac{7\gamma^2}{32(1-\gamma)} d_\mu^{(t)}(1) \leq \frac{\partial V^{(t)}(\mu)}{\partial \theta(1, a_1)} \leq \frac{\gamma^2}{2(1-\gamma)} d_\mu^{(t)}(1)\tag{108}$$

for any $t < t_1(\gamma^2 - 1/4)$, where we have used the elementary facts

$$\min_{1/8 \leq x \leq 7/8} x(1-x) = 7/64 \quad \text{and} \quad \max_{0 \leq x \leq 1} x(1-x) = 1/4.$$

Similarly, repeating the above argument with the gradient expression (103) leads to

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(2, a_1)} \leq \frac{\gamma^4}{2(1-\gamma)} d_\mu^{(t)}(2) \quad \text{for all } t \geq 0; \quad (109a)$$

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(2, a_1)} \geq \frac{7\gamma^4}{32(1-\gamma)} d_\mu^{(t)}(2) \quad \text{for all } 0 \leq t < t_2(\gamma^4 - 1/4). \quad (109b)$$

Further, note that Lemma 2 and Lemma 3 deliver upper and lower bounds on the quantities $d_\mu^{(t)}(1)$ and $d_\mu^{(t)}(2)$, which allow us to deduce that

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(1, a_1)} \geq \frac{7\gamma^3 c_m}{32c_{b,1}|\mathcal{S}|} \quad \text{for all } t < t_1(\gamma^2 - 1/4); \quad (110a)$$

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(2, a_1)} \geq \frac{7\gamma^5 c_m}{32c_{b,2}|\mathcal{S}|} \quad \text{for all } t < t_2(\gamma^4 - 1/4); \quad (110b)$$

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(1, a_1)} \leq \frac{\gamma^2(1 + 17c_m/c_{b,1})}{2|\mathcal{S}|} \quad \text{for all } t < \min\{t_1(\tau_1), t_2(\tau_2)\}; \quad (110c)$$

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(2, a_1)} \leq \frac{\gamma^4(1 + 8c_m/c_{b,2})}{2|\mathcal{S}|} \quad \text{for all } t < t_2(\tau_2). \quad (110d)$$

Step 4: develop an upper bound on $t_1(\gamma^2 - 1/4)$. The preceding bounds allow us to develop an upper bound on $t_1(\gamma^2 - 1/4)$. To do so, it is first observed from the fact $\theta^{(t)}(1, a_0) = -\theta^{(t)}(1, a_1)$ (due to (65)) that

$$\pi^{(t)}(a_1 | 1) = \frac{\exp(\theta^{(t)}(1, a_1))}{\exp(\theta^{(t)}(1, a_0)) + \exp(\theta^{(t)}(1, a_1))} = 1 - \frac{1}{1 + \exp(2\theta^{(t)}(1, a_1))}.$$

Recognizing that $V^{(t)}(1) < \gamma^2 - 1/4$ occurs if and only if $\pi^{(t)}(a_1 | 1) < 1 - (8\gamma^2)^{-1}$ (see (106)), we can easily demonstrate that

$$\theta^{(t)}(1, a_1) \leq \frac{1}{2} \log(8\gamma^2 - 1) \leq \frac{1}{2} \log 7 \quad \text{for all } t < t_1(\gamma^2 - 1/4). \quad (111)$$

If $t_1(\gamma^2 - 1/4) \geq \lceil \frac{32 \log(7)c_{b,1}|\mathcal{S}|}{7\gamma^3 c_m \eta} \rceil$, then taking $t = \lceil \frac{32 \log(7)c_{b,1}|\mathcal{S}|}{7\gamma^3 c_m \eta} \rceil$ together with (110) and (12a) yields

$$\theta^{(t)}(1, a_1) \geq \theta^{(0)}(1, a_1) + \eta \frac{7\gamma^3 c_m}{32c_{b,1}|\mathcal{S}|} t = \eta \frac{7\gamma^3 c_m}{32c_{b,1}|\mathcal{S}|} t \geq \log 7,$$

thus leading to contradiction with (111). As a result, one arrives at the following upper bound:

$$t_1(\tau_1) \leq t_1(\gamma^2 - 1/4) \leq \frac{32 \log(7)c_{b,1}|\mathcal{S}|}{7\gamma^3 c_m \eta} \leq \frac{15c_{b,1}|\mathcal{S}|}{c_m \eta}, \quad (112)$$

with the proviso that $\gamma \geq 0.85$ (so that $\tau_1 \leq \gamma^2 - 1/4$).

An upper bound on $t_2(\gamma^4 - 1/4)$ (and hence $t_2(\tau_2)$) can be obtained in a completely analogous manner

$$t_2(\tau_2) \leq t_2(\gamma^4 - 1/4) \leq \frac{15c_{b,2}|\mathcal{S}|}{c_m \eta},$$

provided that $\gamma \geq 0.95$ (so that $\tau_2 \leq \gamma^4 - 1/4$). We omit the proof of this part for the sake of brevity.

Step 5: develop a lower bound on $t_2(\tau_2)$. Repeating the argument in (106) and (111), we see that $V^{(t)}(2) \geq \tau_2$ if and only if $\pi^{(t)}(a_1 | 2) \geq \frac{1}{2} + \frac{\tau_2}{2\gamma^4}$, which is also equivalent to

$$\theta^{(t)}(2, a_1) \geq \frac{1}{2} \log \left(\frac{1}{\frac{1}{2} - \frac{\tau_2}{2\gamma^4}} - 1 \right) > \frac{1}{2} \log 3,$$

as long as $2\tau_2 > \gamma^4$. Of necessity, this implies that $\theta^{(t)}(2, a_1) > \frac{1}{2} \log 3$ when $t = t_2(\tau_2)$. If $t_2(\tau_2) \leq \frac{|\mathcal{S}| \log 3}{2\eta\gamma^4(1+8c_m/c_{b,2})}$, then invoking (110) and (12a) and taking $t = t_2(\tau_2)$ yield

$$\theta^{(t)}(2, a_1) \leq \theta^{(0)}(2, a_1) + \eta \frac{\gamma^4}{2|\mathcal{S}|} \left(1 + \frac{8c_m}{c_{b,2}}\right) t = \frac{\eta\gamma^4 t}{2|\mathcal{S}|} \left(1 + \frac{8c_m}{c_{b,2}}\right) \leq \frac{1}{2} \log 3,$$

thus resulting in contradiction. We can thus conclude that

$$t_2(\tau_2) > \frac{|\mathcal{S}| \log 3}{\eta\gamma^4(1+8c_m/c_{b,2})} > \frac{|\mathcal{S}| \log 3}{\eta(1+8c_m/c_{b,2})}. \quad (113)$$

As an important byproduct, comparing (113) with (112) immediately reveals that

$$t_2(\tau_2) \geq t_1(\gamma^2 - 1/4) \geq t_1(\tau_1), \quad (114)$$

with the proviso that $\frac{\log 3}{1+8c_m/c_{b,2}} \geq \frac{15c_{b,1}}{c_m}$ and $\gamma \geq 0.87$ (so that $\gamma^2 - 1/4 > \tau_1$).

Step 6: develop a lower bound on $t_1(\tau_1)$. Repeat the analysis in (106) and (111) to show that: $V^{(t)}(1) \geq \tau_1$ if and only if

$$\theta^{(t)}(1, a_1) \geq \frac{1}{2} \log \left(\frac{1}{\frac{1}{2} - \frac{\tau_1}{2\gamma^2}} - 1 \right) > \frac{1}{2} \log 3.$$

Clearly, this lower bound should hold if $t = t_1(\tau_1)$. In addition, in view of (114), one has $\min\{t_1(\tau_1), t_2(\tau_2)\} = t_1(\tau_1)$. If $t_1(\tau_1) \leq \frac{|\mathcal{S}| \log 3}{\eta\gamma^2(1+17c_m/c_{b,1})}$, then setting $t = t_1(\tau_1) = \min\{t_1(\tau_1), t_2(\tau_2)\}$ and applying (110) and (12a) lead to

$$\theta^{(t)}(1, a_1) \leq \theta^{(0)}(1, a_1) + \eta \frac{\gamma^2(1+17c_m/c_{b,1})}{2|\mathcal{S}|} t = \frac{\eta\gamma^2 t(1+17c_m/c_{b,1})}{2|\mathcal{S}|} \leq \frac{1}{2} \log 3,$$

which is contradictory to the preceding lower bound. This in turn implies that

$$t_1(\tau_1) \geq \frac{|\mathcal{S}| \log 3}{\eta\gamma^2(1+17c_m/c_{b,1})} > \frac{|\mathcal{S}| \log 3}{\eta(1+17c_m/c_{b,1})}. \quad (115)$$

C.2 Crossing times for the adjoint state $\bar{1}$

We now move on to the proof of (38b). Note that we have developed a lower bound on $t_2(\tau_2)$ in (113). In order to justify the advertised result $t_2(\tau_2) > t_{\bar{1}}(\gamma^3 - 1/4)$, it thus suffices to demonstrate that

$$t_{\bar{1}}(\gamma^3 - 1/4) \leq \frac{|\mathcal{S}| \log 3}{\eta(1+8c_m/c_{b,2})}, \quad (116)$$

a goal we aim to accomplish in this subsection.

To do so, we divide into two cases. In the scenario where $t_1(\tau_1) \geq t_{\bar{1}}(\gamma^3 - 1/4)$, the bound (112) derived previously immediately leads to the desired bound:

$$t_{\bar{1}}(\gamma^3 - 1/4) \leq t_1(\tau_1) \leq \frac{15c_{b,1}|\mathcal{S}|}{c_m\eta} \leq \frac{|\mathcal{S}| \log 3}{\eta(1+8c_m/c_{b,2})},$$

with the proviso that $\frac{15c_{b,1}}{c_m} \leq \frac{\log 3}{1+8c_m/c_{b,2}}$. Consequently, the subsequent analysis concentrates on establishing (116) for the case where

$$t_1(\tau_1) < t_{\bar{1}}(\gamma^3 - 1/4).$$

In what follows, we divide into three stages and investigate each one separately, after presenting some basic gradient calculations that shall be invoked frequently.

Gradient characterizations. To begin with, observe from (12) that

$$\begin{aligned}\frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_1)} &= \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{1}) \pi^{(t)}(a_1 | \bar{1}) \left(Q^{(t)}(\bar{1}, a_1) - V^{(t)}(\bar{1}) \right) \\ &= \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{1}) \pi^{(t)}(a_1 | \bar{1}) \left(Q^{(t)}(\bar{1}, a_1) - \sum_{a \in \{a_0, a_1\}} \pi^{(t)}(a | \bar{1}) Q^{(t)}(\bar{1}, a) \right) \\ &= \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{1}) \pi^{(t)}(a_1 | \bar{1}) \pi^{(t)}(a_0 | \bar{1}) \left(Q^{(t)}(\bar{1}, a_1) - Q^{(t)}(\bar{1}, a_0) \right),\end{aligned}\tag{117a}$$

which makes use of the fact $\pi^{(t)}(a_0 | \bar{1}) + \pi^{(t)}(a_1 | \bar{1}) = 1$. Analogously, we have

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_0)} = -\frac{1}{1-\gamma} d_\mu^{(t)}(\bar{s}) \pi^{(t)}(a_1 | \bar{1}) \pi^{(t)}(a_0 | \bar{1}) \left(Q^{(t)}(\bar{1}, a_1) - Q^{(t)}(\bar{1}, a_0) \right).\tag{117b}$$

Stage 1: any t obeying $t < t_1(\tau_1)$. We start by looking at each term in the gradient expression (117a) separately. First, note that when $t < t_1(\tau_1)$, one has $V^{(t)}(1) < \tau_1$, which combined with (61) in Lemma 8 indicates that $Q^{(t)}(\bar{1}, a_1) = \gamma V^{(t)}(1) < \gamma \tau_1 = Q^{(t)}(\bar{1}, a_0)$. In fact, from the definition (18a) of τ_1 , the property (61) and Lemma 10, we have

$$1/2 \geq Q^{(t)}(\bar{1}, a_0) > Q^{(t)}(\bar{1}, a_1) = \gamma V^{(t)}(1) \geq 0.$$

Additionally, recall that $t_1(\tau_1) < t_2(\tau_2)$ (see (114)). Lemma 3 then tells us that $d_\mu^{(t)}(\bar{1}) \leq 14c_m(1-\gamma)^2$ during this stage. Substituting these into (117a) and using $\pi^{(t)}(a_0 | \bar{1}) \leq 1$, we arrive at

$$0 \geq \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_1)} \geq -7c_m(1-\gamma)\pi^{(t)}(a_1 | \bar{1}),\tag{118}$$

which together with the PG update rule (12) also indicates that $\theta^{(t)}(\bar{1}, a_1)$ (and hence $\pi^{(t)}(a_1 | \bar{1})$) is monotonically non-increasing with t in this stage. Invoke the auxiliary fact in Lemma 14 to reach

$$\pi^{(t+1)}(a_1 | \bar{1}) - \pi^{(t)}(a_1 | \bar{1}) \geq 2\eta\pi^{(t)}(a_1 | \bar{1}) \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_1)} \geq -14\eta c_m(1-\gamma) \left[\pi^{(t)}(a_1 | \bar{1}) \right]^2.$$

Taking the preceding recursive relation together with Lemma 11 and recalling the initialization $\pi^{(0)}(a_1 | \bar{1}) = 1/2$, we can guarantee that

$$\pi^{(t)}(a_1 | \bar{1}) \geq \frac{1}{28\eta c_m(1-\gamma)t + 2} \quad \text{for all } t \leq t_1(\tau_1)\tag{119}$$

provided that $14\eta c_m(1-\gamma) \leq 1$. In conclusion, the above calculation precludes $\pi^{(t)}(a_1 | \bar{1})$ from decaying to zero too quickly, an observation that is particularly useful for our analysis in Stage 3.

Stage 2: any t obeying $t_1(\tau_1) \leq t < t_1(\gamma^2 - 1/4)$. The only step lies in extending the lower bound (119) to this stage. From the definition (31) of $t_1(\tau_1)$ as well as the monotonicity of $V^{(t)}(1)$ (see Lemma 9), we know that

$$V^{(t)}(1) \geq V^{(t_1(\tau_1))}(1) \geq \tau_1 \quad \text{for all } t \geq t_1(\tau_1),$$

provided that $\eta < (1-\gamma)^2/5$. This taken together with the property (61) in Lemma 8 reveals that

$$Q^{(t)}(\bar{1}, a_1) - Q^{(t)}(\bar{1}, a_0) \geq 0 \quad \text{for all } t \geq t_1(\tau_1),$$

and hence $\pi^{(t)}(a_1 | \bar{1})$ is non-decreasing in t during this stage. Therefore, we have

$$\pi^{(t)}(a_1 | \bar{1}) \geq \pi^{(t_1(\tau_1))}(a_1 | \bar{1}) \geq \frac{1}{28\eta c_m(1-\gamma)t_1(\tau_1) + 2}, \quad t \geq t_1(\tau_1),\tag{120}$$

where the first inequality follows from the non-decreasing property established above, and the second inequality follows from the lower bound (119). In fact, we have established a lower bound on $\pi^{(t)}(a_1 | \bar{1})$ that holds for the entire trajectory of the algorithm.

Stage 3: any t obeying $t_1(\gamma^2 - 1/4) \leq t \leq t_{\bar{1}}(\gamma^3 - 1/4)$. To facilitate analysis, we single out a time threshold t' as follows:

$$t' := \min \left\{ t \mid \pi^{(t)}(a_0 \mid \bar{1}) < 1/2 \right\}. \quad (121)$$

We begin by developing an upper bound on $\pi^{(t)}(a_0 \mid \bar{1})$ for any $t \geq \max\{t', t_1(\gamma^2 - 1/4)\}$. Towards this, with the help of (61) in Lemma 8 we make the observation that: for any $t \geq t_1(\gamma^2 - 1/4)$, one has

$$Q^{(t)}(\bar{1}, a_1) - Q^{(t)}(\bar{1}, a_0) = \gamma V^{(t)}(1) - \gamma \tau_1 \geq \gamma(\gamma^2 - 1/4) - \gamma \tau_1 \geq 0.1 \quad (122)$$

as long as $\gamma \geq 0.92$, which combined with (117b) indicates that

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_0)} < 0 \quad \text{for all } t \geq t_1(\gamma^2 - 1/4). \quad (123)$$

Recognizing that $d_\mu^\pi(\bar{1}) \geq c_m \gamma(1 - \gamma)^2$ (see Lemma 2), we can continue the derivation (117b) to derive

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_0)} \leq -\frac{1}{1 - \gamma} \gamma c_m (1 - \gamma)^2 \cdot \frac{1}{2} \cdot \pi^{(t)}(a_0 \mid \bar{1}) \cdot 0.1 = -0.05 c_m \gamma (1 - \gamma) \pi^{(t)}(a_0 \mid \bar{1})$$

for any $t \geq \max\{t', t_1(\gamma^2 - 1/4)\}$, which implies

$$\pi^{(t)}(a_1 \mid \bar{1}) \geq \pi^{(t')}(a_1 \mid \bar{1}) = 1 - \pi^{(t')}(a_0 \mid \bar{1}) \geq 1/2 \quad \text{for any } t \geq t'.$$

Invoke Lemma 14 to arrive at

$$\pi^{(t+1)}(a_0 \mid \bar{1}) - \pi^{(t)}(a_0 \mid \bar{1}) \leq \frac{\eta}{2} \pi^{(t)}(a_0 \mid \bar{1}) \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_0)} \leq -\frac{\eta}{40} c_m \gamma (1 - \gamma) \left[\pi^{(t)}(a_0 \mid \bar{1}) \right]^2,$$

provided that $2\eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_0)} \geq -1$, which is guaranteed by $\eta < (1 - \gamma)/2$. Recalling that $\pi^{(t)}(a_0 \mid \bar{1}) \leq 1/2$ for this entire stage, one can apply Lemma 11 to obtain

$$\pi^{(t)}(a_0 \mid \bar{1}) \leq \frac{1}{\frac{\eta}{40} c_m \gamma (1 - \gamma) \left(t - \max\{t', t_1(\gamma^2 - 1/4)\} \right) + 2} \quad (124)$$

for any $t \geq \max\{t', t_1(\gamma^2 - 1/4)\}$.

With the above upper bound (124) in place, we are capable of showing that the target quantity $t_{\bar{1}}(\gamma^3 - 1/4)$ is not much larger than $\max\{t', t_1(\gamma^2 - 1/4)\}$. To show this, we first note that the value function of the adjoint state $\bar{1}$ obeys (see Part (iii) in Lemma 8)

$$\begin{aligned} V^{(t)}(\bar{1}) &= \pi^{(t)}(a_0 \mid \bar{1}) Q^{(t)}(\bar{1}, a_0) + \pi^{(t)}(a_1 \mid \bar{1}) Q^{(t)}(\bar{1}, a_1) = \gamma \tau_1 \pi^{(t)}(a_0 \mid \bar{1}) + \gamma \pi^{(t)}(a_1 \mid \bar{1}) V^{(t)}(1) \\ &= \gamma \tau_1 \pi^{(t)}(a_0 \mid \bar{1}) + \gamma V^{(t)}(1) \left\{ 1 - \pi^{(t)}(a_0 \mid \bar{1}) \right\} \geq \gamma \tau_1 \pi^{(t)}(a_0 \mid \bar{1}) + \gamma (\gamma^2 - 1/4) \left\{ 1 - \pi^{(t)}(a_0 \mid \bar{1}) \right\} \\ &= \gamma \left\{ \tau_1 - \gamma^2 + 1/4 \right\} \pi^{(t)}(a_0 \mid \bar{1}) + \gamma^3 - \gamma/4, \end{aligned}$$

where the inequality holds since $V^{(t)}(1) \geq \gamma^2 - 1/4$ in this stage (given that $t \geq t_1(\gamma^2 - 1/4)$). Recognizing that $0.5\gamma^{2/3} - \gamma^2 + 1/4 < 0$ for any $\gamma \geq 0.85$, we can rearrange terms to demonstrate that $V^{(t)}(\bar{1}) \geq \gamma^3 - 1/4$ holds once

$$\pi^{(t)}(a_0 \mid \bar{1}) \leq \frac{1 - \gamma}{4\gamma(\gamma^2 - 1/4 - 0.5\gamma^{2/3})}.$$

In fact, for any $\gamma \geq 0.85$, the above inequality is guaranteed to hold as long as $\pi^{(t)}(a_0 \mid \bar{1}) \leq 1 - \gamma$ since $4\gamma(\gamma^2 - 1/4 - 0.5\gamma^{2/3}) < 1$. In view of (124), we can achieve $\pi^{(t)}(a_0 \mid \bar{1}) \leq 1 - \gamma$ as soon as $t - \max\{t', t_1(\gamma^2 - 1/4)\}$ surpasses $\frac{40}{c_m \gamma \eta (1 - \gamma)^2}$. As a consequence, we reach

$$t_{\bar{1}}(\gamma^3 - 1/4) \leq \max\{t', t_1(\gamma^2 - 1/4)\} + \frac{40}{c_m \gamma \eta (1 - \gamma)^2}. \quad (125)$$

Armed with the relation (125), the goal of upper bounding $t_{\bar{1}}(\gamma^3 - 1/4)$ can be accomplished by controlling t' . To this end, we claim for the moment that

$$t' \leq \frac{1121t_1(\gamma^2 - 1/4)}{\gamma}. \quad (126)$$

If this claim holds, then combining it with (125) and (112) would result in the advertised bound (116):

$$t_{\bar{1}}(\gamma^3 - 1/4) \leq \frac{9972c_{b,1}|\mathcal{S}|}{\gamma^4 c_m \eta} + \frac{40}{c_m \gamma \eta (1 - \gamma)^2} \leq \frac{|\mathcal{S}|}{4\gamma^4 \eta} \leq \frac{|\mathcal{S}| \log 3}{\eta(1 + 8c_m/c_{b,2})},$$

where the penultimate inequality relies on the assumptions $\frac{c_{b,1}}{c_m} \leq \frac{1}{79776}$ and $|\mathcal{S}| \geq \frac{320\gamma^3}{c_m(1-\gamma)^2}$, and the last one holds as long as $\frac{1}{4\gamma^4} \leq \frac{\log 3}{1+8c_m/c_{b,2}}$. To finish up, it suffices to establish the claim (126).

Proof of the claim (126). It is sufficient to consider the case where $t' > t_1(\gamma^2 - 1/4)$; otherwise the inequality (126) is trivially satisfied. Since Lemma 2 tells us that $d_\mu^\pi(\bar{1}) \geq c_m \gamma (1 - \gamma)^2$, we can see from (117a) that, for any t with $t_1(\gamma^2 - 1/4) \leq t < t'$,

$$\begin{aligned} \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_1)} &= \frac{1}{1 - \gamma} d_\mu^{(t)}(\bar{1}) \pi^{(t)}(a_1 | \bar{1}) \pi^{(t)}(a_0 | \bar{1}) \left\{ Q^{(t)}(\bar{1}, a_1) - Q^{(t)}(\bar{1}, a_0) \right\} \\ &\geq 0.05c_m \gamma (1 - \gamma) \pi^{(t)}(a_1 | \bar{1}) > 0, \end{aligned}$$

where the last line follows by combining (122) and the fact that $\pi^{(t)}(a_0 | \bar{1}) \geq 1/2$ for any $t < t'$ (see the definition (121) of t'). According to Lemma 14, we can demonstrate that

$$\pi^{(t+1)}(a_1 | \bar{1}) - \pi^{(t)}(a_1 | \bar{1}) \geq \eta \pi^{(t)}(a_1 | \bar{1}) \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{1}, a_1)} \geq 0.05\eta c_m \gamma (1 - \gamma) \left[\pi^{(t)}(a_1 | \bar{1}) \right]^2$$

for any t obeying $t_1(\gamma^2 - 1/4) \leq t < t'$. Invoking Lemma 11, we then have

$$\begin{aligned} t' &\leq \frac{1 + 0.025\eta c_m \gamma (1 - \gamma)}{0.05\eta c_m \gamma (1 - \gamma) \pi^{(t_1(\gamma^2 - 1/4))}(a_1 | \bar{1})} + t_1(\gamma^2 - 1/4) < \frac{40}{\eta c_m \gamma (1 - \gamma) \pi^{(t_1(\gamma^2 - 1/4))}(a_1 | \bar{1})} + t_1(\gamma^2 - 1/4) \\ &\leq \frac{40(28t_1(\tau_1) + 2)}{\gamma} + t_1(\gamma^2 - 1/4) \\ &\leq \frac{1121t_1(\gamma^2 - 1/4)}{\gamma} \end{aligned}$$

as claimed, where the second line follows from (120).

C.3 Auxiliary facts

In this subsection, we collect some elementary facts that have been used multiple times in the proof of Lemma 4. Specifically, the lemma below makes clear an explicit link between the gradient $\nabla_\theta V^{(t)}(\mu)$ and the difference between two consecutive policy iterates.

Lemma 14. *Consider any s whose associated action space is $\{a_0, a_1\}$.*

- If $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq 0$, then one has

$$\pi^{(t+1)}(a_1 | s) - \pi^{(t)}(a_1 | s) \geq 2\eta \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)}. \quad (127)$$

- If $\pi^{(t+1)}(a_0 | s) \geq 1/2$ and $-1 \leq 2\eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq 0$, then we have

$$\pi^{(t+1)}(a_1 | s) - \pi^{(t)}(a_1 | s) \leq \frac{\eta}{2} \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)}. \quad (128)$$

- If $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \geq 0$ and if $\pi^{(t+1)}(a_0 | s) \geq 1/2$, then one has

$$\pi^{(t+1)}(a_1 | s) - \pi^{(t)}(a_1 | s) \geq \eta \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)}. \quad (129)$$

Proof of Lemma 14. We make note of the following elementary identity

$$\begin{aligned} \frac{e^{\theta_1}}{e^{\theta_1} + e^{-\theta_1}} - \frac{e^{\theta_2}}{e^{\theta_2} + e^{-\theta_2}} &= \frac{e^{\theta_1 - \theta_2} - e^{-\theta_1 + \theta_2}}{(e^{\theta_1} + e^{-\theta_1})(e^{\theta_2} + e^{-\theta_2})} = \frac{e^{-\theta_1}}{e^{\theta_1} + e^{-\theta_1}} \frac{e^{\theta_2}}{e^{\theta_2} + e^{-\theta_2}} (e^{2(\theta_1 - \theta_2)} - 1) \\ &= \left(1 - \frac{e^{\theta_1}}{e^{\theta_1} + e^{-\theta_1}}\right) \frac{e^{\theta_2}}{e^{\theta_2} + e^{-\theta_2}} (e^{2(\theta_1 - \theta_2)} - 1), \end{aligned}$$

which allows us to write

$$\begin{aligned} \pi^{(t+1)}(a_1 | s) - \pi^{(t)}(a_1 | s) &= \pi^{(t+1)}(a_0 | s) \pi^{(t)}(a_1 | s) \left\{ \exp \left[2\theta_{t+1}(s, a_1) - 2\theta_t(s, a_1) \right] - 1 \right\} \\ &= \pi^{(t+1)}(a_0 | s) \pi^{(t)}(a_1 | s) \left\{ \exp \left[2\eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \right] - 1 \right\}. \end{aligned} \quad (130)$$

- If $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq 0$, then one can deduce that

$$(130) \geq 2\eta \pi^{(t+1)}(a_0 | s) \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \geq 2\eta \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)},$$

where the first inequality relies on the elementary fact $e^x - 1 \geq x$ for all $x \in \mathbb{R}$, and the second one holds since $\pi^{(t+1)}(a_0 | s) \leq 1$ and $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq 0$.

- If $-1 \leq 2\eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq 0$ and $\pi^{(t+1)}(a_0 | s) \geq 1/2$, then one has

$$(130) \leq \eta \pi^{(t+1)}(a_0 | s) \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq \frac{\eta}{2} \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)},$$

where the first inequality comes from the elementary inequality $e^x - 1 \leq 0.5x$ for any $-1 \leq x \leq 0$, and the last inequality is valid since $\pi^{(t+1)}(a_0 | s) \geq 1/2$ and $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq 0$.

- If $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \geq 0$ and if $\pi^{(t+1)}(a_0 | s) \geq 1/2$, then it follows that

$$(130) \geq 2\eta \pi^{(t+1)}(a_0 | s) \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \geq \eta \pi^{(t)}(a_1 | s) \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)},$$

as claimed in (129). □

D Analysis for the initial stage (Lemma 5)

This section establishes Lemma 5, which investigates the dynamics of $\theta^{(t)}(s, a)$ prior to the threshold $t_{s-2}(\tau_{s-2})$. Before proceeding, let us introduce a rescaled version of $\pi^{(t)}(s, a)$ that is sometimes convenient to work with:

$$\hat{\pi}^{(t)}(s, a) := \exp \left(\theta^{(t)}(s, a) - \max_{a' \in \mathcal{A}_s} \theta^{(t)}(s, a') \right) \quad (131)$$

for any state-action pair (s, a) . This is orderwise equivalent to $\pi^{(t)}(s, a)$ since

$$\hat{\pi}^{(t)}(s, a) = \frac{\exp(\theta^{(t)}(s, a))}{\max_{a' \in \mathcal{A}_s} \exp(\theta^{(t)}(s, a'))} \geq \frac{\exp(\theta^{(t)}(s, a))}{\sum_{a' \in \mathcal{A}_s} \exp(\theta^{(t)}(s, a'))} = \pi^{(t)}(s, a); \quad (132a)$$

$$\hat{\pi}^{(t)}(s, a) = \frac{\exp(\theta^{(t)}(s, a))}{\max_{a' \in \mathcal{A}_s} \exp(\theta^{(t)}(s, a'))} \leq \frac{\exp(\theta^{(t)}(s, a))}{\frac{1}{3} \sum_{a' \in \mathcal{A}_s} \exp(\theta^{(t)}(s, a'))} = 3\pi^{(t)}(s, a). \quad (132b)$$

D.1 Two key properties

Our proof is based on the following claim: in order to establish the advertised results of Lemma 5, it suffices to justify the following two properties

$$\frac{1}{1 + 56c_m\eta(1 - \gamma)t} \leq \hat{\pi}^{(t)}(a_1 | s) \leq \frac{1}{1 + \frac{c_m\gamma}{35}\eta(1 - \gamma)^2t} \quad (133)$$

$$\text{and} \quad Q^{(t)}(s, a_2) - V^{(t)}(s) \geq 0 \quad (134)$$

hold for any $t \leq t_{s-2}(\tau_{s-2})$. In light of this claim, our subsequent analysis consists of validating these two inequalities separately, which forms the main content of Section D.2.

We now move on to justify the above claim, namely, Lemma 5 is valid as long as the two key properties (133) and (134) hold true. First, recall that Lemma 12 together with (33) and Lemma 4 tells us that

$$\theta^{(t)}(s, a_0) \geq \theta^{(t)}(s, a_2) \geq \theta^{(t)}(s, a_1) \quad \text{for all } t < t_{s-2}(\tau_{s-2}) \leq t_{s-1}(\tau_{s-1}) \quad (135)$$

for any $3 \leq s \leq H$. Next, note that the gradient takes the following form (cf. (12))

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a)} = \frac{1}{1 - \gamma} d_\mu^{(t)}(s) \pi^{(t)}(a | s) (Q^{(t)}(s, a) - V^{(t)}(s)), \quad a \in \{a_0, a_1, a_2\} \quad (136)$$

which together with the assumption $Q^{(t)}(s, a_2) - V^{(t)}(s) \geq 0$ (cf. (134)) implies that

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} \geq 0 \quad \text{for all } t < t_{s-2}(\tau_{s-2}).$$

Consequently, $\theta^{(t)}(s, a_2)$ keeps increasing before t exceeds $t_{s-2}(\tau_{s-2})$. This combined with the relation (135), the initialization $\theta^{(0)}(s, a_0) = \theta^{(0)}(s, a_2) = 0$ and the constraint $\sum_a \theta^{(t)}(s, a) = 0$ (see Part (vii) of Lemma 8) reveals that

$$\theta^{(t)}(s, a_0) \geq \theta^{(t)}(s, a_2) \geq 0 \geq \theta^{(t)}(s, a_1) \quad \text{for all } t < t_{s-2}(\tau_{s-2}), \quad (137)$$

thereby confirming the desired property (41).

Further, given the non-negativity of $\theta^{(t)}(s, a_2)$ stated in (137), one can readily derive

$$\begin{aligned} \hat{\pi}^{(t)}(a_1 | s) &= \exp \left(\theta^{(t)}(s, a_1) - \max_{a'} \theta^{(t)}(s, a') \right) = \exp \left(\theta^{(t)}(s, a_1) - \theta^{(t)}(s, a_0) \right) \\ &= \exp \left(2\theta^{(t)}(s, a_1) + \theta^{(t)}(s, a_2) \right) \geq \exp \left(2\theta^{(t)}(s, a_1) \right), \end{aligned}$$

where the last line also makes use of the identity $\theta^{(t)}(s, a_0) = -\theta^{(t)}(s, a_1) - \theta^{(t)}(s, a_2)$ (see Part (vii) of Lemma 8). With this observation in mind, the assumed property (133) directly leads to the advertised result (40).

D.2 Proof of the properties (133) and (134)

This subsection presents the proofs of the two key properties, which are somewhat intertwined and require a bit of induction. Before proceeding, we make note of the initialization $\hat{\pi}^{(0)}(a_1 | s) = 1$, which clearly satisfies the property (133) for this base case. Our proof consists of two steps to be detailed below. As can be easily seen, combining these two steps in an inductive manner immediately establishes both properties (133) and (134) for any $t \leq t_{s-2}(\tau_{s-2})$.

Step 1: justifying (134) for the t -th iteration if (133) holds for the t -th iteration

We first turn to the proof of the inequality (134), assuming that (133) holds for the t -th iteration. According to (132) and (133), we have

$$\pi^{(t)}(a_1 | s) \geq \frac{1}{3} \hat{\pi}^{(t)}(a_1 | s) \geq \frac{1}{3 + 168c_m\eta(1 - \gamma)t}. \quad (138)$$

By virtue of the auxiliary fact (146c) in Lemma 15 (see Section D.3), one has

$$Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) \leq \frac{\gamma p}{\frac{c_m \gamma}{2} \eta (1 - \gamma) t + \frac{1}{\gamma^{\tau_s - 2}}}. \quad (139)$$

Given that $p := c_p(1 - \gamma)$ for some small constant $0 < c_p < \frac{1}{2016}$, the above two inequalities allow one to ensure that

$$Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) < (\gamma^{3/2} - \gamma^2) \tau_{s-1} \pi^{(t)}(a_1 | s). \quad (140)$$

With the above relation in mind, we are ready to control $Q^{(t)}(s, a_2) - V^{(t)}(s)$ as follows

$$\begin{aligned} Q^{(t)}(s, a_2) - V^{(t)}(s) &= \left(\sum_a \pi^{(t)}(a | s) \right) Q^{(t)}(s, a_2) - \sum_a \pi^{(t)}(a | s) Q^{(t)}(s, a) \\ &= \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \right) - \pi^{(t)}(a_0 | s) \left(Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) \right) \\ &\geq \pi^{(t)}(a_1 | s) (\gamma^{3/2} - \gamma^2) \tau_{s-1} - \left(Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) \right) \\ &> 0. \end{aligned}$$

Here, the second lines arise from the auxiliary facts in Lemma 15, while the last inequality is a consequence of (140). Then we complete the proof of the inequality (134).

Step 2: justifying (133) for the $(t + 1)$ -th iteration if (134) holds up to the t -th iteration

Suppose that the inequality (134) holds up to the t -th iteration. To validate (133) for the $(t + 1)$ -th iteration, we claim for the moment that

$$-14c_m(1 - \gamma)\widehat{\pi}^{(t)}(a_1 | s) \leq \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq -\frac{c_m \gamma}{24}(1 - \gamma)^2 \widehat{\pi}^{(t)}(a_1 | s) < 0 \quad (141)$$

as long as $t \leq t_s(\tau_s)$. Let us take this claim as given, and return to prove it shortly.

Recall from (135) that

$$\theta^{(t)}(s, a_0) \geq \theta^{(t)}(s, a_2) \geq \theta^{(t)}(s, a_1)$$

and hence $\theta_{\max}^{(t)}(s) = \theta^{(t)}(s, a_0)$ is increasing with t during this stage, and as a result,

$$\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) + \theta_{\max}^{(t)}(s) - \theta_{\max}^{(t+1)}(s) \leq \theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) \leq 0.$$

The gradient expression (136) combined with the satisfaction of (134) up to the t -th iteration implies that $\theta^{(t)}(s, a_2)$ is increasing up to the t -th iteration. Given that $\sum_a \theta^{(t)}(s, a) = 0$ (see Part (vii) of Lemma 8), we can derive

$$\begin{aligned} \theta_{\max}^{(t)}(s) - \theta_{\max}^{(t+1)}(s) &= \theta^{(t)}(s, a_0) - \theta^{(t+1)}(s, a_0) = \theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) + \theta^{(t+1)}(s, a_2) - \theta^{(t)}(s, a_2) \\ &\geq \theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1), \end{aligned}$$

thus indicating that

$$\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) + \theta_{\max}^{(t)}(s) - \theta_{\max}^{(t+1)}(s) \geq 2 \left(\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) \right).$$

These combined with Lemma 16 in Section D.3 guarantee that

$$\begin{aligned} \widehat{\pi}^{(t+1)}(a_1 | s) - \widehat{\pi}^{(t)}(a_1 | s) &\geq 2\widehat{\pi}^{(t)}(a_1 | s) \left(\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) \right), \\ \widehat{\pi}^{(t+1)}(a_1 | s) - \widehat{\pi}^{(t)}(a_1 | s) &\leq 0.7\widehat{\pi}^{(t)}(a_1 | s) \left(\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) \right), \end{aligned}$$

and as a consequence,

$$2\hat{\pi}^{(t)}(a_1 | s) \cdot \eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq \hat{\pi}^{(t+1)}(a_1 | s) - \hat{\pi}^{(t)}(a_1 | s) \leq 0.7\hat{\pi}^{(t)}(a_1 | s) \cdot \eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)}.$$

Taking this collectively with (141), we reach

$$-28c_m\eta(1-\gamma)\left[\hat{\pi}^{(t)}(a_1 | s)\right]^2 \leq \hat{\pi}^{(t+1)}(a_1 | s) - \hat{\pi}^{(t)}(a_1 | s) \leq -\frac{c_m\gamma}{35}\eta(1-\gamma)^2\left[\hat{\pi}^{(t)}(a_1 | s)\right]^2. \quad (142)$$

Apply Lemma 11 together with the initialization $\hat{\pi}^{(0)}(a_1 | s) = 1$ to arrive at

$$\frac{1}{1 + 56c_m\eta(1-\gamma)(t+1)} \leq \hat{\pi}^{(t+1)}(a_1 | s) \leq \frac{1}{1 + \frac{c_m\gamma}{35}\eta(1-\gamma)^2(t+1)}. \quad (143)$$

Proof of the inequality (141). Recall the gradient expression (136):

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} = \frac{1}{1-\gamma} d_\mu^{(t)}(s) \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_1) - V^{(t)}(s) \right), \quad (144)$$

each term of which will be bounded separately.

The first step is to control $Q^{(t)}(s, a_1) - V^{(t)}(s)$, towards which we start with the following decomposition

$$\begin{aligned} Q^{(t)}(s, a_1) - V^{(t)}(s) &= Q^{(t)}(s, a_1) - \sum_{a \in \{a_0, a_1, a_2\}} \pi^{(t)}(a | s) Q^{(t)}(s, a) \\ &= -\pi^{(t)}(a_0 | s) \left(Q^{(t)}(s, a_0) - Q^{(t)}(s, a_1) \right) - \pi^{(t)}(a_2 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \right). \end{aligned} \quad (145)$$

The auxiliary facts stated in Lemma 15 (see Appendix D.3) imply that

$$Q^{(t)}(s, a_0) - Q^{(t)}(s, a_1) \geq Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \geq (\gamma^{3/2} - \gamma^2)\tau_{s-1},$$

while Lemma 1 and Lemma 10 tell us that

$$Q^{(t)}(s, a_0) - Q^{(t)}(s, a_1) \leq V^*(s) - 0 = \gamma^{2s}.$$

At the same time, the auxiliary fact (146a) in Lemma 15 (see Appendix D.3) taken together with the gradient expression (12b) guarantees that

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_0)} \geq \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} \geq \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)}$$

and hence $\theta^{(t)}(s, a_1) \leq \theta^{(t)}(s, a_2) \leq \theta^{(t)}(s, a_0)$ (or equivalently $\pi^{(t)}(a_1 | s) \leq \pi^{(t)}(a_2 | s) \leq \pi^{(t)}(a_0 | s)$) during this stage. As a result,

$$\pi^{(t)}(a_1 | s) \leq 1/3 \quad \text{and} \quad 1 \geq \pi^{(t)}(a_0 | s) + \pi^{(t)}(a_2 | s) \geq 2/3.$$

Substituting the preceding bounds into the decomposition (145), we arrive at

$$\begin{aligned} Q^{(t)}(s, a_1) - V^{(t)}(s) &\leq -\left(\pi^{(t)}(a_0 | s) + \pi^{(t)}(a_2 | s)\right) \min \left\{ Q^{(t)}(s, a_0) - Q^{(t)}(s, a_1), Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \right\} \\ &\leq -\frac{2}{3}(\gamma^{\frac{3}{2}} - \gamma^2)\tau_{s-1} = -\frac{2}{3} \frac{\gamma^{\frac{3}{2}}\tau_{s-1}}{1 + \sqrt{\gamma}}(1-\gamma) \leq -\frac{1-\gamma}{8}, \end{aligned}$$

provided that $\gamma \geq 0.85$. Meanwhile, it follows from Lemma 1 and Lemma 10 that

$$Q^{(t)}(s, a_1) - V^{(t)}(s) \geq 0 - V^*(s) \geq -1.$$

Further, from Lemma 3, we have learned that $c_m\gamma(1-\gamma)^2 \leq d_\mu^{(t)}(s) \leq 14c_m(1-\gamma)^2$ for any $t \leq t_s(\tau_s)$. Substituting the above bounds into (144) and invoking (132), we establish the desired inequality (141).

D.3 Auxiliary facts

We now gather a few basic facts that are useful throughout this section. The first lemma presents some preliminary facts regarding the difference of Q-function estimates across different actions in the current setting; the proof is deferred to Appendix D.3.1.

Lemma 15. *Consider any $t < t_{s-2}(\tau_{s-2})$. Under the assumption (35), the following are satisfied*

$$Q^{(t)}(s, a_0) > Q^{(t)}(s, a_2) > Q^{(t)}(s, a_1), \quad (146a)$$

$$Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \geq (\gamma^{3/2} - \gamma^2)\tau_{s-1}, \quad (146b)$$

$$Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) \leq \frac{\gamma p}{\frac{c_m \gamma}{2} \eta (1 - \gamma) t + \frac{1}{\gamma \tau_{s-2}}}. \quad (146c)$$

Remark 10. Lemma 15 makes clear that — before t exceeds $t_{s-2}(\tau_{s-2})$ — action a_0 is perceived as the best choice, with a_1 being the least favorable one. In the meantime, it also reveals that (i) $Q^{(t)}(s, a_2)$ is considerably larger than $Q^{(t)}(s, a_1)$, while (ii) the gap between $Q^{(t)}(s, a_0)$ and $Q^{(t)}(s, a_2)$ decays at least as rapidly as $O(1/t)$ in this stage.

The second lemma is concerned with the consecutive difference between two rescaled policy iterates. The proof can be found in Appendix D.3.2.

Lemma 16. *Suppose that $0 < \eta \leq (1 - \gamma)/6$. For any $t \geq 0$ and any $3 \leq s \leq H$, define $\theta_{\max}^{(t)}(s) := \max_a \theta^{(t)}(s, a)$. If we write*

$$\hat{\pi}^{(t+1)}(a_1 | s) - \hat{\pi}^{(t)}(a_1 | s) = c \hat{\pi}^{(t)}(a_1 | s) \left(\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) + \theta_{\max}^{(t)}(s) - \theta_{\max}^{(t+1)}(s) \right) \quad (147)$$

for some $c \in \mathbb{R}$, then we necessarily have

$$\begin{aligned} c \in [1, 1.5] & \quad \text{if } \theta^{(t+1)}(s, a_1) \geq \theta^{(t)}(s, a_1) \quad \text{and} \quad \theta_{\max}^{(t)}(s) \geq \theta_{\max}^{(t+1)}(s); \\ c \in (0.72, 1] & \quad \text{if } \theta^{(t+1)}(s, a_1) \leq \theta^{(t)}(s, a_1) \quad \text{and} \quad \theta_{\max}^{(t)}(s) \leq \theta_{\max}^{(t+1)}(s). \end{aligned}$$

D.3.1 Proof of Lemma 15

In view of Lemma 10, one has $V^{(t)}(\overline{s-2}) \geq 0$ for all $t \geq 0$. Therefore, the relation (59) yields

$$Q^{(t)}(s, a_2) = r_s + \gamma p V^{(t)}(\overline{s-2}) \geq r_s = \gamma^{3/2} \tau_{s-1}.$$

In addition, for any $t < t_{s-2}(\tau_{s-2}) \leq t_{s-1}(\tau_{s-1}) \leq t_{s-1}(\gamma \tau_{s-1})$ (see Lemma 8 and Lemma 4), we have $V^{(t)}(\overline{s-1}) < \gamma \tau_{s-1}$, and hence it is seen from the relation (59) that

$$Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) = Q^{(t)}(s, a_2) - \gamma V^{(t)}(\overline{s-1}) \geq (\gamma^{3/2} - \gamma^2)\tau_{s-1} > 0,$$

as claimed in (146b). Also, Part (i) of Lemma 8 tells us that

$$Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) = r_s + \gamma^2 p \tau_{s-2} - r_s - \gamma p V^{(t)}(\overline{s-2}) = \gamma p \left(\gamma \tau_{s-2} - V^{(t)}(\overline{s-2}) \right) \geq 0,$$

where the last inequality holds for any $t < t_{s-2}(\tau_{s-2})$ (see Part (iii) of Lemma 8). These taken together validate (146a).

It remains to justify (146c), which is the content of the rest of this proof. The main step lies in demonstrating that, for any $t < t_s(\tau_s)$ and any $1 \leq s \leq H$,

$$\gamma \tau_s - V^{(t)}(\overline{s}) \leq \frac{1}{\frac{c_m \gamma}{2} \eta (1 - \gamma) t + \frac{1}{\gamma \tau_s}}. \quad (148)$$

If this were true, then taking it together with the following property (which is a consequence of (59))

$$Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) = \gamma p \left(\gamma \tau_{s-2} - V^{(t)}(\overline{s-2}) \right), \quad (149)$$

would establish the inequality (146c). It then boils down to justifying (148). Towards this, we first make the observation that

$$\begin{aligned} V^{(t)}(\bar{s}) - \gamma\tau_s &= \pi^{(t)}(a_0 | \bar{s})Q^{(t)}(\bar{s}, a_0) + \pi^{(t)}(a_1 | \bar{s})Q^{(t)}(\bar{s}, a_1) - \gamma\tau_s \\ &= \pi^{(t)}(a_0 | \bar{s})\gamma\tau_s + \pi^{(t)}(a_1 | \bar{s})Q^{(t)}(\bar{s}, a_1) - \gamma\tau_s \\ &= \pi^{(t)}(a_1 | \bar{s})\left(Q^{(t)}(\bar{s}, a_1) - \gamma\tau_s\right), \end{aligned} \quad (150)$$

where the second line holds since $Q^{(t)}(\bar{s}, a_0) = \gamma\tau_s$ (see (61)). Additionally, recall from the definition that for any $t < t_s(\tau_s)$, one has $V^{(t)}(s) < \tau_s$ and hence

$$\begin{aligned} \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{s}, a_1)} &= \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{s}) \pi^{(t)}(a_1 | \bar{s}) \left(Q^{(t)}(\bar{s}, a_1) - \pi^{(t)}(a_0 | \bar{s})Q^{(t)}(\bar{s}, a_0) - \pi^{(t)}(a_1 | \bar{s})Q^{(t)}(\bar{s}, a_1) \right) \\ &= \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{s}) \pi^{(t)}(a_0 | \bar{s}) \pi^{(t)}(a_1 | \bar{s}) \left(Q^{(t)}(\bar{s}, a_1) - Q^{(t)}(\bar{s}, a_0) \right) \\ &= \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{s}) \pi^{(t)}(a_0 | \bar{s}) \pi^{(t)}(a_1 | \bar{s}) \left(\gamma V^{(t)}(s) - \gamma\tau_s \right) < 0, \end{aligned} \quad (151)$$

where the last line makes use of the identities in (61). This means that $\theta^{(t)}(\bar{s}, a_1)$ keeps decreasing, and hence $\theta^{(t)}(\bar{s}, a_1) \leq 0$ given the initialization $\theta^{(0)}(\bar{s}, a_1) = 0$. As an immediate consequence, one has $\theta^{(t)}(\bar{s}, a_0) = -\theta^{(t)}(\bar{s}, a_1) \geq 0$ and $\pi^{(t)}(a_0 | \bar{s}) \geq 1/2$. Taking this observation together with (151) and Lemma 2 gives

$$\begin{aligned} \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{s}, a_1)} &= \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{s}) \pi^{(t)}(a_0 | \bar{s}) \pi^{(t)}(a_1 | \bar{s}) \left(Q^{(t)}(\bar{s}, a_1) - Q^{(t)}(\bar{s}, a_0) \right) \\ &= \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{s}) \pi^{(t)}(a_0 | \bar{s}) \pi^{(t)}(a_1 | \bar{s}) \left(Q^{(t)}(\bar{s}, a_1) - \gamma\tau_s \right) \\ &\leq \frac{c_m \gamma}{2} (1-\gamma) \pi^{(t)}(a_1 | \bar{s}) \left(Q^{(t)}(\bar{s}, a_1) - \gamma\tau_s \right) < 0. \end{aligned}$$

Moreover, combine (151) with Lemma 3 and Lemma 1 to yield

$$\left| \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{s}, a_1)} \right| \leq \frac{1}{1-\gamma} d_\mu^{(t)}(\bar{s}) \left| Q^{(t)}(\bar{s}, a_1) - \gamma\tau_s \right| \leq \frac{1}{1-\gamma} 14c_m (1-\gamma)^2 \left(|Q^{(t)}(\bar{s}, a_1)| + \gamma\tau_s \right) \leq 28c_m (1-\gamma),$$

If $28c_m \eta (1-\gamma) < 1/2$, then the above inequalities taken together with Lemma 14 give

$$\pi^{(t+1)}(a_1 | \bar{s}) - \pi^{(t)}(a_1 | \bar{s}) \leq \frac{c_m \gamma}{2} \eta (1-\gamma) \left[\pi^{(t)}(a_1 | \bar{s}) \right]^2 \left(Q^{(t)}(\bar{s}, a_1) - \gamma\tau_s \right) \quad (152)$$

for all $t < t_s(\tau_s)$. This combined with (150) and the monotonicity of $Q^{(t)}(\bar{s}, a_1)$ (see Lemma 9) gives

$$\begin{aligned} \gamma\tau_s - V^{(t+1)}(\bar{s}) &= \pi^{(t+1)}(a_1 | \bar{s}) \left(\gamma\tau_s - Q^{(t+1)}(\bar{s}, a_1) \right) \leq \pi^{(t+1)}(a_1 | \bar{s}) \left(\gamma\tau_s - Q^{(t)}(\bar{s}, a_1) \right) \\ &\leq \left\{ \pi^{(t)}(a_1 | \bar{s}) - \frac{c_m \gamma}{2} \eta (1-\gamma) \left[\pi^{(t)}(a_1 | \bar{s}) \right]^2 \left(\gamma\tau_s - Q^{(t)}(\bar{s}, a_1) \right) \right\} \left(\gamma\tau_s - Q^{(t)}(\bar{s}, a_1) \right) \\ &= \left\{ 1 - \frac{\eta c_m \gamma (1-\gamma)}{2} \left(\gamma\tau_s - V^{(t)}(\bar{s}) \right) \right\} \left(\gamma\tau_s - V^{(t)}(\bar{s}) \right), \end{aligned}$$

where the penultimate line follows from the inequality (152) for the iteration $t-1$, and the last identity makes use of (150). In conclusion, we have arrived at the following inductive relation

$$\gamma\tau_s - V^{(t+1)}(\bar{s}) \leq \gamma\tau_s - V^{(t)}(\bar{s}) - \frac{\eta c_m \gamma (1-\gamma)}{2} \left(\gamma\tau_s - V^{(t)}(\bar{s}) \right)^2,$$

which bears resemblance to the recursive relations studied in Lemma 11. Recognizing that $\gamma\tau_s - V^{(0)}(\bar{s}) \leq \gamma\tau_{s-2}$ (since $V^{(0)}(\bar{s}) \geq 0$ according to Lemma 10), we can invoke Lemma 11 to derive

$$\gamma\tau_s - V^{(t)}(\bar{s}) \leq \frac{1}{\frac{c_m \gamma}{2} \eta (1-\gamma) t + \frac{1}{\gamma\tau_s}}.$$

Putting the above pieces together concludes the proof of (146c).

D.3.2 Proof of Lemma 16

From the definition (131), direct calculations lead to

$$\begin{aligned}\widehat{\pi}^{(t+1)}(a_1 | s) - \widehat{\pi}^{(t)}(a_1 | s) &= \exp\left(\theta^{(t+1)}(s, a_1) - \theta_{\max}^{(t+1)}(s)\right) - \exp\left(\theta^{(t)}(s, a_1) - \theta_{\max}^{(t)}(s)\right) \\ &= \widehat{\pi}^{(t)}(a_1 | s) \left\{ \exp\left(\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) + \theta_{\max}^{(t)}(s) - \theta_{\max}^{(t+1)}(s)\right) - 1 \right\}.\end{aligned}$$

According to Lemma 1, we have $|Q^{(t)}(s, a)| \leq 1$ and $|V^{(t)}(s)| \leq 1$, which indicates — for any action $a \in \{a_0, a_1, a_2\}$ — that

$$\left| \theta^{(t+1)}(s, a) - \theta^{(t)}(s, a) \right| = \left| \eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a)} \right| = \frac{\eta}{1 - \gamma} d_{\mu}^{(t)}(s) \pi^{(t)}(a | s) \left| Q^{(t)}(s, a) - V^{(t)}(s) \right| \leq \frac{1}{3},$$

provided that $\eta \leq (1 - \gamma)/6$. An immediate consequence is that $|\theta_{\max}^{(t+1)}(s) - \theta_{\max}^{(t)}(s)| \leq 1/3$ and hence

$$\left| \theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) + \theta_{\max}^{(t)}(s) - \theta_{\max}^{(t+1)}(s) \right| \leq 2/3.$$

This taken together with the following elementary facts

$$(e^x - 1)/x \in [1, 1.5) \quad \text{for } 0 \leq x < 2/3, \quad \text{and} \quad (e^x - 1)/x \in (0.72, 1] \quad \text{for } -2/3 < x \leq 0$$

establishes the claim (147).

E Analysis for the intermediate stage (Lemma 6)

We now turn attention to Lemma 6, which studies the dynamics during an intermediate stage between $t_{s-2}(\tau_{s-2})$ and $t_{s-1}(\tau_s)$.

E.1 Main steps

Key facts regarding crossing times. Our proof for Lemma 6 relies on several crucial properties regarding the crossing times for both the key primary states and the adjoint states, as stated in the following two lemmas.

Lemma 17. Suppose that (35) holds. There exists some constant $0 < c_0 \leq \frac{1222}{c_m \gamma}$ such that:

$$t_s(\gamma^{2s} - 1/4) - \max \left\{ t_{s-1}(\gamma^{2s-1} - 1/4), t_s(\tau_s) \right\} \leq \frac{c_0}{\eta(1 - \gamma)^2} \quad (153)$$

holds for every $3 \leq s \leq H$, and

$$t_{\bar{s}}(\gamma^{2s+1} - 1/4) - \max \left\{ t_s(\gamma^{2s} - 1/4), t_{\bar{s}}(\tau_{s+1}) \right\} \leq \frac{c_0}{\eta(1 - \gamma)^2} \quad (154)$$

holds for every $1 \leq s \leq H$.

Lemma 18. Suppose that (35) holds and

$$t_3(\tau_3) > t_2(\gamma^4 - 1/4). \quad (155)$$

Then for every $3 \leq s \leq H$, we have

$$t_s(\gamma^{2s} - 1/4) - t_s(\tau_s) \leq \frac{2sc_0}{\eta(1 - \gamma)^2}, \quad (156a)$$

$$t_{s-1}(\gamma^{2s-1} - 1/4) - t_{s-1}(\tau_s) \leq \frac{2sc_0}{\eta(1 - \gamma)^2}. \quad (156b)$$

In addition, if we further have $t_{s-1}(\tau_{s-1}) > t_{s-2}(\tau_{s-1}) + \frac{2sc_0}{\eta(1 - \gamma)^2}$, then

$$t_{s-2}(\gamma^{2s-3} - 1/4) \leq t_{s-1}(\tau_s). \quad (156c)$$

Furthermore, (156c) still holds for $s = 3$ without requiring the assumption (155).

The proofs of the above two lemmas are postponed to Appendix E.2 and Appendix E.3, respectively. Let us take a moment to explain these two lemmas; to provide some intuitions, let us treat $\gamma^{2s} \approx 1$. Lemma 17 makes clear that: once the value function estimates for states $\overline{s-1}$ and s are both sufficiently large (i.e., $V^{(t)}(\overline{s-1}) \gtrsim 0.75$ and $V^{(t)}(s) \gtrsim 0.5$), then it does not take long for $V^{(t)}(s)$ to (approximately) exceed 0.75. A similar message holds true if we replace s (resp. $\overline{s-1}$) with \overline{s} (resp. s). Built upon this observation, Lemma 18 further reveals that: the time taken for $V^{(t)}(s)$ (resp. $V^{(t)}(\overline{s-1})$) to rise from 0.5 to 0.75 is fairly short.

Proof of Lemma 6. We are now in a position to present the proof of Lemma 6. To begin with, recall from Lemma 8 that: for any $t \leq t_{\overline{s-1}}(\tau_s) \leq t_{\overline{s-1}}(\tau_{s-1})$, one has

$$Q^{(t)}(s, a_1) = \gamma V^{(t)}(\overline{s-1}) \leq \gamma \tau_s \leq \min \left\{ Q^{(t)}(s, a_0), Q^{(t)}(s, a_2) \right\}. \quad (157)$$

Given that $V^{(t)}(s)$ is a convex combination of $\{Q^{(t)}(s, a)\}_{a \in \{a_0, a_1, a_2\}}$, one has $V^{(t)}(s) - Q^{(t)}(s, a_1) \geq 0$, which together with the gradient expression (136) indicates that

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \leq 0 \quad (158)$$

and hence $\theta^{(t)}(s, a_1)$ is non-increasing with t for any $t < t_{\overline{s-1}}(\tau_s)$. Additionally, we have learned from Lemma 18 that

$$t_{\overline{s-1}}(\tau_s) \geq t_{\overline{s-2}}(\gamma^{2s-3} - 1/4) \geq t_{\overline{s-2}}(\gamma \tau_{s-2}) = t_{s-2}(\tau_{s-2}),$$

where the second inequality holds since $\gamma^{2s-3} - 1/4 \geq \gamma \tau_{s-2}$, and the last identity results from Part (iii) of Lemma 8. This combined with the non-increasing nature of $\theta^{(t)}(s, a_1)$ readily establishes the advertised inequality $\theta^{(t_{\overline{s-1}}(\tau_s))}(s, a_1) \leq \theta^{(t_{s-2}(\tau_{s-2}))}(s, a_1)$.

The next step is to justify $\theta^{(t_{\overline{s-1}}(\tau_s))}(s, a_2) \geq 0$. Notice that for $t > t_{s-2}(\tau_{s-2})$, we have $V^{(t)}(s-2) > \tau_{s-2}$, and then $V^{(t)}(\overline{s-2}) > \gamma \tau_{s-2}$ by (62), which leads to $Q^{(t)}(s, a_2) > Q^{(t)}(s, a_0)$ by (59) in Lemma 8. Recall (157) that $Q^{(t)}(s, a_1) \leq \gamma \tau_s \leq \min \left\{ Q^{(t)}(s, a_0), Q^{(t)}(s, a_2) \right\}$. Then, one has $Q^{(t)}(s, a_2) - V^{(t)}(s) \geq 0$, which together with the gradient expression (136) indicates that

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} \geq 0 \quad (159)$$

and hence $\theta^{(t)}(s, a_2)$ is non-decreasing with t for any $t < t_{\overline{s-1}}(\tau_s)$. This establishes $\theta^{(t_{\overline{s-1}}(\tau_s))}(s, a_2) \geq 0$.

E.2 Proof of Lemma 17

For every $t \geq \max \{t_{\overline{s-1}}(\gamma^{2s-1} - 1/4), t_s(\tau_s)\}$, we isolate the following properties that will prove useful.

- The definition (30) of $t_{\overline{s-1}}(\cdot)$ together with the monotonicity property in Lemma 9 requires that $V^{(t)}(\overline{s-1}) \geq \gamma^{2s-1} - 1/4$, and hence it is seen from (59) that

$$Q^{(t)}(s, a_1) = \gamma V^{(t)}(\overline{s-1}) \geq \gamma^{2s} - \gamma/4. \quad (160)$$

- In the meantime, since $t \geq t_s(\tau_s)$, Lemma 8 (cf. (60)) guarantees that

$$\pi^{(t)}(a_1 | s) \geq (1 - \gamma)/2. \quad (161)$$

- Given that $t_s(\tau_s) \geq t_{s-2}(\tau_{s-2})$ (see (33)) and the monotonicity property in Lemma 9, one has $V^{(t)}(\overline{s-2}) \geq \tau_{s-2}$, and thus we can see from (59) that

$$Q^{(t)}(s, a_2) - Q^{(t)}(s, a_0) = \gamma p(V^{(t)}(\overline{s-2}) - \tau_{s-2}) \geq 0. \quad (162)$$

- In addition, Lemma 8 ensures that both $Q^{(t)}(s, a_2)$ and $Q^{(t)}(s, a_0)$ are bounded above by $\gamma^{1/2}\tau_s$. Therefore, it is easily seen that

$$Q^{(t)}(s, a_1) \geq \gamma^{2s} - \gamma/4 > \gamma^{1/2}\tau_s \geq Q^{(t)}(s, a_2) \geq Q^{(t)}(s, a_0), \quad (163)$$

where the first inequality comes from (160), the second one holds when $\gamma^{2s} > 0.75$, and the last inequality has been justified in (162).

- Moreover, given that $V^{(t)}(s) \geq \tau_s$ (since $t \geq t_s(\tau_s)$), one further has

$$\begin{aligned} Q^{(t)}(s, a_1) - \max\{Q^{(t)}(s, a_2), Q^{(t)}(s, a_0)\} &> V^{(t)}(s) - \max\{Q^{(t)}(s, a_2), Q^{(t)}(s, a_0)\} \\ &> \tau_s - \gamma^{1/2}\tau_s > 0. \end{aligned} \quad (164)$$

Here, the first inequality comes from (163), while the penultimate inequality is a consequence of (163).

- We have seen from the above bullet points that

$$Q^{(t)}(s, a_1) > V^{(t)}(s) > \max\{Q^{(t)}(s, a_2), Q^{(t)}(s, a_0)\}, \quad (165)$$

which combined with the gradient expression (136) reveals that

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} > 0 > \max\left\{\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_0)}, \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)}\right\}. \quad (166)$$

With the above properties in place, we are now ready to prove our lemma, for which we shall look at the key primary states $3 \leq s \leq H$ and the adjoint states separately.

Analysis for the key primary states. Let us start with any state $3 \leq s \leq H$, and control $t_s(\gamma^{2s} - 1/4)$ as claimed in (153). As before, define

$$\theta_{\max}^{(t)}(s) := \max_a \theta^{(t)}(s, a).$$

From the above fact (166), we know that $\theta^{(t)}(s, a_1)$ keeps increasing with t while $\theta^{(t)}(s, a_0), \theta^{(t)}(s, a_2)$ are both decreasing with t . As a result, once $\theta^{(t)}(s, a_1) = \theta_{\max}^{(t)}(s)$, then $\theta^{(t)}(s, a_1)$ will remain equal to $\theta_{\max}^{(t)}(s)$ for the subsequent iterations. This allows us to divide into two stages as follows.

- **Stage 1: the duration when $\theta^{(t)}(s, a_1) < \theta_{\max}^{(t)}(s)$.** Our aim is to show that this stage contains at most $O(\frac{1}{\eta(1-\gamma)^2})$ iterations. In order to prove this, the starting point is again the gradient expression (136):

$$\begin{aligned} \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} &= \frac{1}{1-\gamma} d_{\mu}^{(t)}(s) \pi^{(t)}(a_1 | s) (Q^{(t)}(s, a_1) - V^{(t)}(s)) \\ &\geq c_m \gamma (1-\gamma) \pi^{(t)}(a_1 | s) (Q^{(t)}(s, a_1) - V^{(t)}(s)), \end{aligned} \quad (167)$$

where the last line relies on Lemma 2 and the fact $Q^{(t)}(s, a_1) > V^{(t)}(s)$ (cf. (165)). Regarding the size of $Q^{(t)}(s, a_1) - V^{(t)}(s)$, we make the observation that

$$\begin{aligned} Q^{(t)}(s, a_1) - V^{(t)}(s) &= \pi^{(t)}(a_0 | s) (Q^{(t)}(s, a_1) - Q^{(t)}(s, a_0)) + \pi^{(t)}(a_2 | s) (Q^{(t)}(s, a_1) - Q^{(t)}(s, a_2)) \\ &\geq (\pi^{(t)}(a_0 | s) + \pi^{(t)}(a_2 | s)) \left(Q^{(t)}(s, a_1) - \max_{a \in \{a_0, a_2\}} Q^{(t)}(s, a) \right) \\ &\stackrel{(i)}{\geq} \frac{1}{2} \left(Q^{(t)}(s, a_1) - \max_{a \in \{a_0, a_2\}} Q^{(t)}(s, a) \right) \stackrel{(ii)}{\geq} \frac{1}{2} (\gamma^{2s} - \gamma/4 - \gamma^{1/2}\tau_s) \stackrel{(iii)}{\geq} \frac{1}{16}. \end{aligned}$$

Here, (i) follows since $\theta^{(t)}(s, a_1) < \theta_{\max}^{(t)}(s)$ during this stage and, therefore, $\pi^{(t)}(a_1 | s) \leq 1/2$; (ii) arises from the relation (163); and (iii) holds whenever $\gamma^{2s} - \gamma/4 > 5/8$. Substitution into (167) yields

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \geq \frac{1}{16} c_m \gamma (1-\gamma) \pi^{(t)}(a_1 | s) \geq \frac{1}{48} c_m \gamma (1-\gamma) \hat{\pi}^{(t)}(a_1 | s), \quad (168)$$

where the last inequality comes from (132). In addition, recall that $\theta^{(t)}(s, a_1)$ is increasing with t , while $\theta^{(t)}(s, a_0)$ and $\theta^{(t)}(s, a_2)$ are both decreasing (and hence $\theta_{\max}^{(t)}(s)$ is also decreasing). Invoking Lemma 16 then yields

$$\begin{aligned}\widehat{\pi}^{(t+1)}(a_1 | s) - \widehat{\pi}^{(t)}(a_1 | s) &\geq \widehat{\pi}^{(t)}(a_1 | s) \left(\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) + \theta_{\max}^{(t)}(s) - \theta_{\max}^{(t+1)}(s) \right) \\ &\geq \widehat{\pi}^{(t)}(a_1 | s) \left(\theta^{(t+1)}(s, a_1) - \theta^{(t)}(s, a_1) \right) \\ &= \widehat{\pi}^{(t)}(a_1 | s) \cdot \eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \geq \frac{1}{48} c_m \eta \gamma (1 - \gamma) \left[\widehat{\pi}^{(t)}(a_1 | s) \right]^2,\end{aligned}$$

where the last line arises from (168). Given this recursive relation, Lemma 11 implies that: if $\widehat{\pi}^{(t)}(a_1 | s) < 1$ (or equivalently, $\theta^{(t)}(s, a_1) < \theta_{\max}^{(t)}(s)$), then one necessarily has

$$t - t_{0,1} \leq \frac{1 + \frac{1}{48} c_m \eta \gamma (1 - \gamma)}{\frac{1}{48} c_m \eta \gamma (1 - \gamma) \pi^{(t_0)}(a_1 | s)} \leq \frac{2}{\frac{1}{48} c_m \eta \gamma (1 - \gamma) \pi^{(t_{0,1})}(a_1 | s)} \leq \frac{240}{c_m \eta \gamma (1 - \gamma)^2},$$

with $t_{0,1} := \max \{t_{s-1}(\gamma^{2s-1} - 1/4), t_s(\tau_s)\}$. Here, the last inequality relies on the property (161).

- **Stage 2: the duration when $\theta^{(t)}(s, a_1) = \theta_{\max}^{(t)}(s)$.** For this stage, we intend to demonstrate that it takes at most $O(\frac{1}{\eta(1-\gamma)^2})$ iterations to achieve $\max \{\pi^{(t)}(a_0 | s), \pi^{(t)}(a_2 | s)\} \leq (1 - \gamma)/8$. To this end, we again begin by studying the gradient as follows:

$$\begin{aligned}\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} &= \frac{1}{1 - \gamma} d_{\mu}^{(t)}(s) \pi^{(t)}(a_2 | s) \left(Q^{(t)}(s, a_2) - V^{(t)}(s) \right) \leq c_m \gamma (1 - \gamma) \pi^{(t)}(a_2 | s) \left(Q^{(t)}(s, a_2) - V^{(t)}(s) \right) \\ &\leq \frac{1}{3} c_m \gamma (1 - \gamma) \widehat{\pi}^{(t)}(a_2 | s) \left(Q^{(t)}(s, a_2) - V^{(t)}(s) \right).\end{aligned}$$

Here, the first inequality comes from Lemma 2 and the fact $Q^{(t)}(s, a_2) < V^{(t)}(s)$ (see (165)), whereas the last inequality is a consequence of (132). In order to control $Q^{(t)}(s, a_2) - V^{(t)}(s)$, we observe that

$$\begin{aligned}Q^{(t)}(s, a_2) - V^{(t)}(s) &= \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \right) + \pi^{(t)}(a_0 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_0) \right) \\ &\leq \pi^{(t)}(a_1 | s) \left(\gamma^{1/2} \tau_s - \gamma^{2s} + \gamma/4 \right) + \pi^{(t)}(a_2 | s) \gamma p \left(V^{(t)}(\overline{s-2}) - \tau_{s-2} \right) \\ &\leq \frac{1}{3} \left(\gamma^{1/2} \tau_s - \gamma^{2s} + \gamma/4 \right) + \gamma p \leq -\frac{1}{24},\end{aligned}$$

where the second line arises from (163) and (162), and the last line holds since $V^{(t)}(\overline{s-2}) \leq 1$ as well as the fact $\pi^{(t)}(a_1 | s) \geq 1/3$ during this stage (since $\theta^{(t)}(s, a_1) = \theta_{\max}^{(t)}(s)$). Putting the above two bounds together leads to

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} \leq -\frac{1}{72} c_m \gamma (1 - \gamma) \widehat{\pi}^{(t)}(a_2 | s). \quad (169)$$

Next, Lemma 16 tells us that

$$\begin{aligned}\widehat{\pi}^{(t+1)}(a_2 | s) - \widehat{\pi}^{(t)}(a_2 | s) &\leq 0.72 \widehat{\pi}^{(t)}(a_2 | s) \left(\theta^{(t+1)}(s, a_2) - \theta^{(t)}(s, a_2) + \theta_{\max}^{(t)}(s) - \theta_{\max}^{(t+1)}(s) \right) \\ &\leq 0.72 \widehat{\pi}^{(t)}(a_2 | s) \left(\theta^{(t+1)}(s, a_2) - \theta^{(t)}(s, a_2) \right) = 0.72 \widehat{\pi}^{(t)}(a_2 | s) \cdot \eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} \\ &\leq -0.01 \eta c_m \gamma (1 - \gamma) \left[\widehat{\pi}^{(t)}(a_2 | s) \right]^2,\end{aligned}$$

where the first inequality makes use of the facts $\theta^{(t+1)}(s, a_2) \leq \theta^{(t)}(s, a_2)$ and $\theta_{\max}^{(t)}(s) = \theta^{(t)}(s, a_1) \leq \theta^{(t+1)}(s, a_1) = \theta_{\max}^{(t+1)}(s)$ (see (166)). Denoting by $t_{0,2}$ the first iteration in this stage, we can invoke Lemma 11 to reach

$$\widehat{\pi}^{(t-t_{0,2})}(a_2 | s) \leq \frac{1}{0.01 \eta c_m \gamma (1 - \gamma) (t - t_{0,2}) + 1}. \quad (170)$$

As a consequence, once $t - t_{0,2}$ exceeds

$$\frac{800}{\eta c_m \gamma (1 - \gamma)^2},$$

then one has $\pi^{(t)}(a_2 | s) \leq (1 - \gamma)/8$. The same conclusion holds for a_0 as well.

Combining the above analysis for the two stages, we see that: if

$$t - t_{0,1} \geq \frac{240}{\eta c_m \gamma (1 - \gamma)^2} + \frac{800}{\eta c_m \gamma (1 - \gamma)^2} = \frac{1040}{\eta c_m \gamma (1 - \gamma)^2}$$

with $t_{0,1} := \max \{t_{s-1}(\gamma^{2s-1} - 1/4), t_s(\tau_s)\}$, then one has

$$\pi^{(t)}(a_1 | s) = 1 - \pi^{(t)}(a_0 | s) - \pi^{(t)}(a_2 | s) \geq 1 - (1 - \gamma)/4,$$

which combined with (163) leads to

$$V^{(t)}(s) \geq \pi^{(t)}(a_1 | s) Q^{(t)}(s, a_1) \geq (1 - (1 - \gamma)/4)(\gamma^{2s} - \gamma/4) \geq \gamma^{2s} - 1/4.$$

This means that one necessarily has $t \geq t_s(\gamma^{2s} - 1/4)$. It then follows that

$$t_s(\gamma^{2s} - 1/4) - \max \{t_{s-1}(\gamma^{2s-1} - 1/4), t_s(\tau_s)\} = t_s(\gamma^{2s} - 1/4) - t_{0,1} \leq \frac{1040}{\eta c_m \gamma (1 - \gamma)^2},$$

thus concluding the proof of (153).

Analysis for the adjoint states. We then move forward to the adjoint states $\{\bar{1}, \dots, \bar{H}\}$ and control $t_{\bar{s}}(\gamma^{2s+1} - 1/4)$ as desired in (154). The proof consists of studying the dynamic for any t obeying

$$\max \{t_s(\gamma^{2s} - 1/4), t_{\bar{s}}(\tau_{s+1})\} \leq t \leq t_{\bar{s}}(\gamma^{2s+1} - 1/4).$$

Once again, we divide into two stages and analyze each of them separately.

- **Stage 1: the duration where $\theta^{(t)}(\bar{s}, a_1) < \theta^{(t)}(\bar{s}, a_0)$.** We aim to demonstrate that it takes no more than $O(\frac{1}{\eta(1-\gamma)^2})$ iterations for $\theta^{(t)}(\bar{s}, a_1)$ to surpass $\theta^{(t)}(\bar{s}, a_0)$. In order to do so, note that

$$\begin{aligned} \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{s}, a_1)} &= \frac{1}{1 - \gamma} d_{\mu}^{(t)}(\bar{s}) \pi^{(t)}(a_1 | \bar{s}) \pi^{(t)}(a_0 | \bar{s}) \left(Q^{(t)}(\bar{s}, a_1) - Q^{(t)}(\bar{s}, a_0) \right) \\ &\geq \frac{1}{16} c_m \gamma (1 - \gamma) \pi^{(t)}(a_1 | \bar{s}) > 0. \end{aligned} \tag{171}$$

Here, the last line applies Lemma 2 and makes use of the fact

$$Q^{(t)}(\bar{s}, a_1) - Q^{(t)}(\bar{s}, a_0) = \gamma V^{(t)}(s) - \gamma \tau_s \geq \gamma(\gamma^{2s} - 1/4 - \tau_s) \geq 1/8. \tag{172}$$

where the inequality comes from the assumption $t \geq t_s(\gamma^{2s} - 1/4)$ as well as the monotonicity property in Lemma 9. As a result, the PG update rule (12a) implies that $\theta^{(t)}(\bar{s}, a_1)$ is increasing in t , and hence $\theta^{(t)}(\bar{s}, a_0)$ is decreasing in t (since $\sum_a \theta^{(t)}(s, a) = 0$); these taken collectively mean that

$$\theta^{(t+1)}(\bar{s}, a_1) - \theta^{(t)}(\bar{s}, a_1) + \theta^{(t)}(\bar{s}, a_0) - \theta^{(t+1)}(\bar{s}, a_0) \geq \theta^{(t+1)}(\bar{s}, a_1) - \theta^{(t)}(\bar{s}, a_1) \geq 0.$$

Invoking Lemma 16 then reveals that

$$\begin{aligned} \hat{\pi}^{(t+1)}(a_1 | \bar{s}) - \hat{\pi}^{(t)}(a_1 | \bar{s}) &\geq \hat{\pi}^{(t)}(a_1 | \bar{s}) \left(\theta^{(t+1)}(\bar{s}, a_1) - \theta^{(t)}(\bar{s}, a_1) + \theta^{(t)}(\bar{s}, a_0) - \theta^{(t+1)}(\bar{s}, a_0) \right) \\ &\geq \hat{\pi}^{(t)}(a_1 | \bar{s}) \left(\theta^{(t+1)}(\bar{s}, a_1) - \theta^{(t)}(\bar{s}, a_1) \right) = \eta \hat{\pi}^{(t)}(a_1 | \bar{s}) \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{s}, a_1)} \end{aligned}$$

$$\geq \frac{1}{48} \eta c_m \gamma (1 - \gamma) \left[\widehat{\pi}^{(t)}(a_1 | \bar{s}) \right]^2,$$

where the last inequality relies on (171) and (132). Given this recursive relation, Lemma 11 tells us that: one has $\widehat{\pi}^{(t)}(a_1 | \bar{s}) \geq 1$ (which means a_1 becomes the favored action by (131)) as soon as $t - t_{0,3}$ exceeds

$$\frac{2}{\frac{1}{48} \eta c_m \gamma (1 - \gamma) \widehat{\pi}^{(t_{0,3})}(a_1 | \bar{s})} \leq \frac{96}{\eta c_m \gamma (1 - \gamma) \pi^{(t_{0,3})}(a_1 | \bar{s})} \leq \frac{1152}{\eta c_m \gamma (1 - \gamma)^2},$$

where $t_{0,3} := \max \left\{ t_s (\gamma^{2s} - 1/4), t_{\bar{s}}(\tau_{s+1}) \right\}$. Here, the last inequality is valid as long as

$$\pi^{(t_{0,3})}(a_1 | \bar{s}) \geq (1 - \gamma)/12 \quad (173)$$

holds. It thus remains to justify (173). Towards this end, observe that for any $t \geq t_{\bar{s}}(\tau_{s+1})$,

$$\begin{aligned} \tau_{s+1} &\leq V^{(t)}(\bar{s}) = \pi^{(t)}(a_0 | \bar{s}) Q^{(t)}(\bar{s}, a_0) + \pi^{(t)}(a_1 | \bar{s}) Q^{(t)}(\bar{s}, a_1) = \pi^{(t)}(a_0 | \bar{s}) \gamma \tau_s + \pi^{(t)}(a_1 | \bar{s}) \gamma V^{(t)}(s) \\ &= \gamma \tau_s + \pi^{(t)}(a_1 | \bar{s}) \gamma \left(V^{(t)}(s) - \tau_s \right) \leq \gamma \tau_s + \pi^{(t)}(a_1 | \bar{s}) \gamma, \end{aligned}$$

and, as a result,

$$\pi^{(t)}(a_1 | \bar{s}) \geq \frac{\tau_{s+1}}{\gamma} - \tau_s = \frac{1}{2} \frac{\gamma^{\frac{2}{3}s + \frac{2}{3}} - \gamma^{\frac{2}{3}s+1}}{\gamma} = \frac{\gamma^{\frac{2}{3}s-1}}{2} \left(\gamma^{\frac{2}{3}} - \gamma \right) \geq \frac{1 - \gamma}{12},$$

provided that $\gamma \geq 0.9$ (so that $\gamma^{\frac{2}{3}} - \gamma \geq 0.3(1 - \gamma)$) and $\gamma^{\frac{2}{3}H} \geq 0.7$. This concludes the analysis of this stage.

- **Stage 2: the duration where $\theta^{(t)}(\bar{s}, a_1) \geq \theta^{(t)}(\bar{s}, a_0)$.** Similar to the above argument, we intend to show that it takes at most $O\left(\frac{1}{\eta(1-\gamma)^2}\right)$ iterations for $\pi^{(t)}(a_0 | \bar{s}) \leq 1 - \gamma$ to occur. From the gradient expression and the property (172), we obtain

$$\begin{aligned} \frac{\partial V^{(t)}(\mu)}{\partial \theta(\bar{s}, a_0)} &= \frac{1}{1 - \gamma} d_{\mu}^{(t)}(\bar{s}) \pi^{(t)}(a_0 | \bar{s}) \pi^{(t)}(a_1 | \bar{s}) \left(Q^{(t)}(\bar{s}, a_0) - Q^{(t)}(\bar{s}, a_1) \right) \\ &\leq -\frac{1}{16} c_m \gamma (1 - \gamma) \pi^{(t)}(a_0 | \bar{s}) \leq -\frac{1}{48} c_m \gamma (1 - \gamma) \widehat{\pi}^{(t)}(a_0 | \bar{s}), \end{aligned}$$

where the first inequality uses Lemma 2 and the property $\pi^{(t)}(a_1 | \bar{s}) \geq 1/2$ (since $\theta^{(t)}(\bar{s}, a_1) \geq \theta^{(t)}(\bar{s}, a_0)$), and the last inequality relies on (132). Repeating a similar argument as above, we can demonstrate that

$$\widehat{\pi}^{(t+1)}(a_0 | \bar{s}) - \widehat{\pi}^{(t)}(a_0 | \bar{s}) \leq -\frac{1}{70} \eta c_m \gamma (1 - \gamma) \left[\widehat{\pi}^{(t)}(a_0 | \bar{s}) \right]^2. \quad (174)$$

This combined with Lemma 11 implies that

$$\widehat{\pi}^{(t)}(a_0 | \bar{s}) \leq \frac{1}{\frac{1}{70} \eta c_m \gamma (1 - \gamma) (t - t_{0,4}) + 1}, \quad (175)$$

with $t_{0,4}$ denoting the first iteration of this stage. Consequently, one has $\widehat{\pi}^{(t)}(a_0 | \bar{s}) \leq 1 - \gamma$ — and therefore $\pi^{(t)}(a_0 | \bar{s}) \leq 1 - \gamma$ according to (132) — as soon as $t - t_{0,4}$ exceeds

$$\frac{70}{\eta c_m \gamma (1 - \gamma)^2}.$$

Finally, if $\pi^{(t)}(a_0 | \bar{s}) \leq 1 - \gamma$, then one has

$$V^{(t)}(\bar{s}) = \pi^{(t)}(a_0 | \bar{s}) Q^{(t)}(\bar{s}, a_0) + \pi^{(t)}(a_1 | \bar{s}) Q^{(t)}(\bar{s}, a_1) = \pi^{(t)}(a_0 | \bar{s}) \gamma \tau_s + \pi^{(t)}(a_1 | \bar{s}) \gamma V^{(t)}(s)$$

$$\begin{aligned}
&\geq \pi^{(t)}(a_0 | \bar{s})\gamma\tau_s + \left(1 - \pi^{(t)}(a_0 | \bar{s})\right)\gamma(\gamma^{2s} - 1/4) = \pi^{(t)}(a_0 | \bar{s})(\gamma\tau_s - \gamma^{2s+1} - \gamma/4) + \gamma^{2s+1} - \gamma/4 \\
&\geq (1 - \gamma)(\gamma\tau_s - \gamma^{2s+1} - \gamma/4) + \gamma^{2s+1} - \gamma/4 \geq \gamma^{2s+1} - 1/4,
\end{aligned}$$

where the first inequality holds by recalling that $t \geq t_s(\gamma^{2s} - 1/4)$. Consequently, putting the above pieces (regarding the duration of the two stages) together allows us to conclude that

$$t_{\bar{s}}(\gamma^{2s+1} - 1/4) - \max\left\{t_s(\gamma^{2s} - 1/4), t_{\bar{s}}(\tau_{s+1})\right\} \leq \frac{1152}{\eta c_m \gamma (1 - \gamma)^2} + \frac{70}{\eta c_m \gamma (1 - \gamma)^2} = \frac{1222}{\eta c_m \gamma (1 - \gamma)^2}$$

as claimed.

E.3 Proof of Lemma 18

Before proceeding, we first single out two properties that play a crucial role in the proof of Lemma 18.

Lemma 19. *The following basic properties hold true for any $2 \leq s \leq H$:*

$$t_s(\tau_s) \geq t_{s-1}(\tau_s); \quad (176a)$$

$$t_{s-1}(\tau_s) \geq t_{s-1}(\tau_{s-1}). \quad (176b)$$

The proof of this auxiliary lemma is deferred to the end of this subsection. Equipped with this result, we are positioned to present the proof of Lemma 18. To begin with, we seek to bound the quantity $t_s(\gamma^{2s} - 1/4) - t_s(\tau_s)$. Apply Lemma 17 with a little algebra to yield

$$\begin{aligned}
t_s(\gamma^{2s} - 1/4) - t_s(\tau_s) &\leq \max\left\{t_{s-1}(\gamma^{2s-1} - 1/4), t_s(\tau_s)\right\} + \frac{c_0}{\eta(1 - \gamma)^2} - t_s(\tau_s) \\
&= \max\left\{t_{s-1}(\gamma^{2s-1} - 1/4) - t_s(\tau_s), 0\right\} + \frac{c_0}{\eta(1 - \gamma)^2}. \quad (177)
\end{aligned}$$

With the assistance of the bound (176a) in Lemma 19, we can continue the bound in (177) to derive

$$\begin{aligned}
t_s(\gamma^{2s} - 1/4) - t_s(\tau_s) &\leq \max\left\{t_{s-1}(\gamma^{2s-1} - 1/4) - t_{s-1}(\tau_s), 0\right\} + \frac{c_0}{\eta(1 - \gamma)^2} \\
&= t_{s-1}(\gamma^{2s-1} - 1/4) - t_{s-1}(\tau_s) + \frac{c_0}{\eta(1 - \gamma)^2}. \quad (178)
\end{aligned}$$

To continue, we shall bound the quantity $t_{s-1}(\gamma^{2s-1} - 1/4) - t_{s-1}(\tau_s)$. Similar to the derivation of the inequality (177), we can apply Lemma 17 to show that

$$\begin{aligned}
t_{s-1}(\gamma^{2s-1} - 1/4) - t_{s-1}(\tau_s) &\leq \max\left\{t_{s-1}(\gamma^{2s-2} - 1/4), t_{s-1}(\tau_s)\right\} + \frac{c_0}{\eta(1 - \gamma)^2} - t_{s-1}(\tau_s) \\
&= \max\left\{t_{s-1}(\gamma^{2s-2} - 1/4) - t_{s-1}(\tau_s), 0\right\} + \frac{c_0}{\eta(1 - \gamma)^2} \\
&\leq \max\left\{t_{s-1}(\gamma^{2s-2} - 1/4) - t_{s-1}(\tau_{s-1}), 0\right\} + \frac{c_0}{\eta(1 - \gamma)^2} \\
&= t_{s-1}(\gamma^{2s-2} - 1/4) - t_{s-1}(\tau_{s-1}) + \frac{c_0}{\eta(1 - \gamma)^2}, \quad (179)
\end{aligned}$$

where the third line makes use of (176b) in Lemma 19.

Applying the inequalities (178) and (179) recursively, one arrives at

$$\begin{aligned}
t_s(\gamma^{2s} - 1/4) - t_s(\tau_s) &\leq t_{s-1}(\gamma^{2s-2} - 1/4) - t_{s-1}(\tau_{s-1}) + \frac{2c_0}{\eta(1 - \gamma)^2} \leq \dots \\
&\leq t_3(\gamma^6 - 1/4) - t_3(\tau_3) + \frac{2(s-3)c_0}{\eta(1 - \gamma)^2}. \quad (180)
\end{aligned}$$

To continue, note that Lemma 17 and the bound (176a) give

$$\begin{aligned} t_3(\gamma^6 - 1/4) &\leq \max \left\{ t_2(\gamma^5 - 1/4), t_3(\tau_3) \right\} + \frac{c_0}{\eta(1-\gamma)^2}, \\ t_2(\gamma^5 - 1/4) &\leq \max \left\{ t_2(\gamma^4 - 1/4), t_2(\tau_3) \right\} + \frac{c_0}{\eta(1-\gamma)^2} \leq \max \left\{ t_2(\gamma^4 - 1/4), t_3(\tau_3) \right\} + \frac{c_0}{\eta(1-\gamma)^2}, \end{aligned}$$

which together leads to

$$t_3(\gamma^6 - 1/4) \leq \max \left\{ t_2(\gamma^4 - 1/4), t_3(\tau_3) \right\} + \frac{2c_0}{\eta(1-\gamma)^2}. \quad (181)$$

Plugging back to (180) leads to

$$\begin{aligned} t_s(\gamma^{2s} - 1/4) - t_s(\tau_s) &\leq \max \left\{ t_2(\gamma^4 - 1/4), t_3(\tau_3) \right\} - t_3(\tau_3) + \frac{2c_0}{\eta(1-\gamma)^2} + \frac{2(s-3)c_0}{\eta(1-\gamma)^2} \\ &\leq \frac{(2s-4)c_0}{\eta(1-\gamma)^2}, \end{aligned} \quad (182)$$

where the last step arises from the assumption (155), that is, $t_2(\gamma^4 - 1/4) < t_3(\tau_3)$.

Further, the above inequality taken together with (179) yields

$$t_{s-1}(\gamma^{2s-1} - 1/4) - t_{s-1}(\tau_s) \leq \frac{(2s-4)c_0}{\eta(1-\gamma)^2} + \frac{c_0}{\eta(1-\gamma)^2} = \frac{(2s-3)c_0}{\eta(1-\gamma)^2}. \quad (183)$$

We have thus established (156a) and (156b).

Finally, we turn to the proof of (156c). In view of (156b), one has

$$t_{s-2}(\gamma^{2s-3} - 1/4) - t_{s-2}(\tau_{s-1}) \leq \frac{2sc_0}{\eta(1-\gamma)^2}.$$

In addition,

$$\begin{aligned} t_{s-1}(\tau_s) - t_{s-2}(\tau_{s-1}) &\geq t_{s-1}(\gamma\tau_{s-1}) - t_{s-2}(\tau_{s-1}) = t_{s-1}(\tau_{s-1}) - t_{s-2}(\tau_{s-1}) \\ &\geq t_{s-1}(\tau_{s-1}) - t_{s-2}(\tau_{s-1}) > \frac{2sc_0}{\eta(1-\gamma)^2}, \end{aligned}$$

where the identity in the first line comes from Part (iii) of Lemma 8, and the last inequality uses the assumption $t_{s-1}(\tau_{s-1}) > t_{s-2}(\tau_{s-1}) + \frac{2sc_0}{\eta(1-\gamma)^2}$. Combining the above two inequalities justifies the validity of the advertised inequality (156c). Then, we establish (156c) for $s = 3$ through Lemma 4, which gives

$$t_1(\gamma^3 - 1/4) \leq t_2(\tau_2) \leq t_2(\tau_3), \quad (184)$$

where the last inequality comes from (176b).

Proof of Lemma 19. To begin with, the claim (176a) holds when $s = 2$ as a result of the inequality (38b) in Lemma 4. We now turn to the case with $3 \leq s \leq H$. In view of the property (59) in Lemma 8, we have

$$\max \left\{ Q^{(t)}(s, a_0), Q^{(t)}(s, a_2) \right\} \leq \gamma^{\frac{1}{2}} \tau_s \quad \text{and} \quad Q^{(t)}(s, a_1) = \gamma V^{(t)}(\overline{s-1}).$$

Recognizing that $V^{(t)}(s)$ is a convex combination of $\{Q^{(t)}(s, a)\}_{a \in \{a_0, a_1, a_2\}}$, we know that if $V^{(t)}(s) \geq \tau_s$, then one necessarily has $Q^{(t)}(s, a_1) > \tau_s$, or equivalently, $V^{(t)}(\overline{s-1}) > \tau_s/\gamma \geq \tau_s$. This essentially means that $t_s(\tau_s) \geq t_{s-1}(\tau_s)$, thus establishing the claim (176a).

Similarly, Lemma 8 (cf. (61)) also tells us that

$$Q^{(t)}(\overline{s}, a_0) = \gamma \tau_s \quad \text{and} \quad Q^{(t)}(\overline{s}, a_1) = \gamma V^{(t)}(s).$$

This means that if $V^{(t)}(s-1) \leq \tau_{s-1}$, then

$$V^{(t)}(\overline{s-1}) \leq \max \left\{ Q^{(t)}(\overline{s-1}, a_0), Q^{(t)}(\overline{s-1}, a_1) \right\} \leq \gamma \tau_{s-1} \leq \tau_s.$$

Consequently, we conclude that $t_{s-1}(\tau_s) \geq t_{s-1}(\tau_{s-1})$, as claimed in (176b). \square

F Analysis for the blowing-up lemma (Lemma 7)

In this section, we establish the blowing-up phenomenon as asserted in Lemma 7.

F.1 Which reference point t_{ref} shall we choose?

Let us specify the time instance t_{ref} as required in Lemma 7 as follows

$$t_{\text{ref}} := \min \left\{ t \in [t_{s-1}(\tau_s), t_s(\tau_s)) \mid c_{\text{ref}}(1-\gamma)\pi^{(t)}(a_0 | s) \leq \pi^{(t)}(a_1 | s) \right\}, \quad (185)$$

where $c_{\text{ref}} \in (0, 1/3)$ is some constant to be specified shortly.

Existence. An important step is to justify that (185) is well-defined, namely, there exists at least one time instance within $[t_{s-1}(\tau_s), t_s(\tau_s))$ that satisfies $c_{\text{ref}}(1-\gamma)\pi^{(t)}(a_0 | s) \leq \pi^{(t)}(a_1 | s)$. Towards this, we note that if the time instance $t_{s-1}(\tau_s)$ obeys

$$c_{\text{ref}}(1-\gamma)\pi^{(t)}(a_0 | s) \leq \pi^{(t)}(a_1 | s) \quad \text{when } t = t_{s-1}(\tau_s),$$

then we simply have $t_{\text{ref}} = t_{s-1}(\tau_s)$. We then move on to the complement case where

$$\begin{aligned} c_{\text{ref}}(1-\gamma)\pi^{(t_{s-1}(\tau_s))}(a_0 | s) &> \pi^{(t_{s-1}(\tau_s))}(a_1 | s), \\ \text{or equivalently, } \theta^{(t_{s-1}(\tau_s))}(s, a_0) &> \theta^{(t_{s-1}(\tau_s))}(s, a_1) - \log(c_{\text{ref}}(1-\gamma)). \end{aligned} \quad (186)$$

To justify that the construction (185) makes sense, it suffices to show that the endpoint $t_s(\tau_s)$ obeys

$$c_{\text{ref}}(1-\gamma)\pi^{(t_s(\tau_s))}(a_0 | s) < \pi^{(t_s(\tau_s))}(a_1 | s). \quad (187)$$

In order to validate (187), recall that the inequality (60) in Lemma 8 ensures that

$$\pi^{(t_s(\tau_s))}(a_1 | s) \geq \frac{1-\gamma}{2},$$

given that $V^{(t_s(\tau_s))}(s) \geq \tau_s$. Therefore, the inequality (187) must be satisfied when $c_{\text{ref}} < 1/2$, given that the left-hand side of (187) obeys

$$c_{\text{ref}}(1-\gamma)\pi^{(t_s(\tau_s))}(a_0 | s) \leq c_{\text{ref}}(1-\gamma) < \frac{1-\gamma}{2}.$$

This in turn validates the existence of (187) for this case.

Several immediate properties about t_{ref} and $t_{s-1}(\tau_s)$. We pause to single out a couple of immediate properties about the t_{ref} constructed above as well as $t_{s-1}(\tau_s)$.

Consider the case where $t_{s-1}(\tau_s)$ obeys

$$\begin{aligned} c_{\text{ref}}(1-\gamma)\pi^{(t_{s-1}(\tau_s))}(a_0 | s) &\leq \pi^{(t_{s-1}(\tau_s))}(a_1 | s), \\ \text{or equivalently, } \theta^{(t_{s-1}(\tau_s))}(s, a_0) &\leq \theta^{(t_{s-1}(\tau_s))}(s, a_1) - \log(c_{\text{ref}}(1-\gamma)), \end{aligned}$$

then one has $t_{\text{ref}} = t_{s-1}(\tau_s)$ (as discussed above). As can be clearly seen, $t_{s-1}(\tau_s)$ satisfies the advertised inequality (45a) by taking $c_{\text{ref}} \geq c_p/8064$. Additionally, let us first recall from (156c) in Lemma 18 that

$$t_{s-1}(\tau_s) = \max \left\{ t_{s-2}(\gamma^{2s-3} - 1/4), t_{s-1}(\tau_s) \right\}.$$

This combined with Lemma 6 (see (43)) tells us that

$$\theta^{(t_{s-1}(\tau_s))}(s, a_1) = \theta^{(t_{s-1}(\tau_s))}(s, a_1) \leq \theta^{(t_{s-2}(\tau_{s-2}))}(s, a_1) \leq -\frac{1}{2} \log \left(1 + \frac{c_m \gamma}{35} \eta (1-\gamma)^2 t_{s-2}(\tau_{s-2}) \right), \quad (188)$$

where the last relation utilizes the bound (40) in Lemma 5. This leads to the advertised bound (45b).

As a result, the claims (45a)-(45b) only need to be justified under the assumption (186).

Organization of the proof. In light of the above basic facts, the subsequent proof focuses on the scenario where (186) is satisfied, namely, the case where

$$t_{s-1}(\tau_s) < t_{\text{ref}}.$$

We shall start by justifying that $\theta^{(t)}(s, a_1)$ has not increased much during $[t_{s-1}(\tau_s), t_{\text{ref}}]$, as detailed in Appendix F.2 and F.3 (focusing on two separate stages respectively). This feature will then be used in Appendix F.4 to establish the claims (45a)-(45b), and in Appendix F.5 to establish the claim (45c).

F.2 Stage I: the duration where $\theta^{(t)}(s, a_2) < \theta^{(t)}(s, a_0)$

Suppose that at the starting point we have $\theta^{(t_{s-1}(\tau_s))}(s, a_2) < \theta^{(t_{s-1}(\tau_s))}(s, a_0)$; otherwise the reader can proceed directly to Stage II in Appendix F.3. The goal is to control the number of iterations taken to achieve $\theta^{(t)}(s, a_2) \geq \theta^{(t)}(s, a_0)$. More specifically, let us define the transition point

$$t_{\text{tran}} := \min \left\{ t \mid \theta^{(t)}(s, a_2) \geq \theta^{(t)}(s, a_0), t \geq t_{s-1}(\tau_s) \right\}. \quad (189)$$

In this subsection, we seek to develop an upper bound on $t_{\text{tran}} - t_{s-1}(\tau_s)$, and to show that $\theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) \leq 1/2$ holds throughout this stage.

Preparation: basic facts and rescaled policies. Before moving forward, we first gather some basic facts. To begin with, from the definition (185) of t_{ref} , we know that the inequality $c_{\text{ref}}(1 - \gamma)\pi^{(t)}(a_0 | s) > \pi^{(t)}(a_1 | s)$ holds true for every $t \in [t_{s-1}(\tau_s), t_{\text{ref}}]$, or equivalently,

$$\theta^{(t)}(s, a_0) > \theta^{(t)}(s, a_1) - \log(c_{\text{ref}}(1 - \gamma)) \quad \text{for all } t \in [t_{s-1}(\tau_s), t_{\text{ref}}]. \quad (190)$$

In the case considered here, we have — according to (190) and (189) — that

$$\theta^{(t)}(s, a_0) > \theta^{(t)}(s, a_1) - \log(c_{\text{ref}}(1 - \gamma)) \quad \text{and} \quad \theta^{(t)}(s, a_0) > \theta^{(t)}(s, a_2) \quad (191)$$

for any t obeying $t_{s-1}(\tau_s) \leq t < \min\{t_{\text{tran}}, t_{\text{ref}}\}$. This means that

$$\theta^{(t)}(s, a_0) = \max_a \theta^{(t)}(s, a) \quad \text{and hence} \quad \pi^{(t)}(a_0 | s) > 1/3 \quad (192)$$

holds for any t obeying $t_{s-1}(\tau_s) \leq t < \min\{t_{\text{tran}}, t_{\text{ref}}\}$, provided that $0 < c_{\text{ref}} < 1$.

Moreover, let us introduce the rescaled policy $\hat{\pi}^{(t)}(a | s)$ as before

$$\hat{\pi}^{(t)}(a | s) := \exp \left(\theta^{(t)}(s, a) - \max_{a' \in \mathcal{A}_s} \theta^{(t)}(s, a') \right).$$

In view of (192), the rescaled policy can therefore be written as

$$\begin{aligned} \hat{\pi}^{(t)}(a_2 | s) &= \exp(\theta^{(t)}(s, a_2) - \theta^{(t)}(s, a_0)) = \exp(2\theta^{(t)}(s, a_2) + \theta^{(t)}(s, a_1)) \\ \hat{\pi}^{(t)}(a_1 | s) &= \exp(\theta^{(t)}(s, a_1) - \theta^{(t)}(s, a_0)) = \exp(2\theta^{(t)}(s, a_1) + \theta^{(t)}(s, a_2)) \end{aligned} \quad (193)$$

for any t with $t_{s-1}(\tau_s) \leq t < \min\{t_{\text{tran}}, t_{\text{ref}}\}$, where we have used the constraint $\sum_a \theta^{(t)}(s, a) = 0$ (see Part (vii) of Lemma 8).

Showing $\theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) \leq 1/2$ by induction. In the following, we seek to prove by induction the following key property

$$\theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) \leq 1/2 \quad (194)$$

for any t that obeys $t_{s-1}(\tau_s) \leq t \leq \min\{t_{\text{tran}}, t_{\text{ref}}\}$ and

$$t - t_{s-1}(\tau_s) \leq \frac{225}{c_p c_m \eta (1 - \gamma)^2 \exp(\theta^{(t_{s-1}(\tau_s))}(s, a_1))} =: \tilde{t}. \quad (195)$$

We shall return to justify (195) for all t within this stage later on. In words, the claim (194) essentially means that $\theta^{(t)}(s, a_1)$ does not deviate much from $\theta^{(t_{s-1}(\tau_s))}(s, a_1)$ during this stage. With regards to the base case where $t = t_{s-1}(\tau_s)$, the hypothesis (194) holds true trivially. Next, assuming that (194) is satisfied for every integer less than or equal to $t - 1$, we intend to establish this hypothesis for the t -th iteration, which is accomplished as follows.

First, Lemma 1 and Lemma 10 tell us that $Q^{(t)}(s, a_1) - V^{(t)}(s) \leq 1$. It then follows that

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} = \frac{1}{1 - \gamma} d_\mu^{(t)}(s) \pi^{(t)}(a_1 | s) \left\{ Q^{(t)}(s, a_1) - V^{(t)}(s) \right\} \leq 14c_m \eta (1 - \gamma) \pi^{(t)}(a_1 | s),$$

which relies on the bound $d_\mu^{(t)}(s) \leq 14c_m(1 - \gamma)^2$ stated in Lemma 3. As a result, it can be derived from the PG update rule (12a) that

$$\begin{aligned} \theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) &= \sum_{j=t_{s-1}(\tau_s)}^{t-1} \eta \frac{\partial V^{(j)}(\mu)}{\partial \theta(s, a_1)} \leq \sum_{j=t_{s-1}(\tau_s)}^{t-1} 14c_m \eta (1 - \gamma) \pi^{(j)}(a_1 | s) \\ &\leq 14c_m \eta (1 - \gamma) (t - t_{s-1}(\tau_s)) \max_{t_{s-1}(\tau_s) \leq j < t} \pi^{(j)}(a_1 | s). \end{aligned} \quad (196)$$

Regarding the term involving $\pi^{(j)}(a_1 | s)$, we observe that for any $t_{s-1}(\tau_s) \leq j < t$,

$$\pi^{(j)}(a_1 | s) \stackrel{(i)}{\leq} \hat{\pi}^{(j)}(a_1 | s) \stackrel{(ii)}{\leq} \exp\left(\frac{3}{2} \theta^{(j)}(s, a_1)\right) \quad (197)$$

$$\stackrel{(iii)}{\leq} \exp\left(\frac{3}{2} (\theta^{(t_{s-1}(\tau_s))}(s, a_1) + 1/2)\right). \quad (198)$$

Here, (i) is a consequence of (132), (ii) holds since (in view of (193), $\theta^{(j)}(s, a_0) \geq 0$, and $\sum_a \theta^{(j)}(s, a) = 0$)

$$\begin{aligned} \hat{\pi}^{(j)}(a_1 | s) &= \exp\left(2\theta^{(j)}(s, a_1) + \theta^{(j)}(s, a_2)\right) \leq \exp\left(2\theta^{(j)}(s, a_1) + 0.5\theta^{(j)}(s, a_2) + 0.5\theta^{(j)}(s, a_0)\right) \\ &= \exp\left(1.5\theta^{(j)}(s, a_1)\right), \end{aligned}$$

whereas (iii) follows from the induction hypothesis (194) for any $t_{s-1}(\tau_s) \leq j < t$. Combine the inequalities (196) and (198) to reach

$$\theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) \leq 14c_m \eta (1 - \gamma) (t - t_{s-1}(\tau_s)) \exp\left(\frac{3}{2} (\theta^{(t_{s-1}(\tau_s))}(s, a_1) + 1/2)\right).$$

Consequently, under the constraint (195), the preceding inequality implies that

$$\begin{aligned} \theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) &\leq 14c_m \eta (1 - \gamma) \frac{225}{c_p c_m \eta (1 - \gamma)^2 \exp(\theta^{(t_{s-1}(\tau_s))}(s, a_1))} \exp\left(\frac{3}{2} (\theta^{(t_{s-1}(\tau_s))}(s, a_1) + 1/2)\right) \\ &= \frac{3150e \exp\left(\frac{1}{2} \theta^{(t_{s-1}(\tau_s))}(s, a_1)\right)}{c_p (1 - \gamma)} \leq \frac{1}{2}, \end{aligned} \quad (199)$$

where the last inequality makes use of (188) and the assumption (44). These allow us to establish the induction hypothesis for the t -th iteration, namely,

$$\theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) \leq 1/2. \quad (200)$$

Validating the constraint (195) and upper bounding $\min\{t_{\text{tran}}, t_{\text{ref}}\} - t_{s-1}(\tau_s)$. It remains to justify the assumed condition (195) for all iteration within this stage. To this end, suppose instead that

$$t_{s-1}(\tau_s) + \tilde{t} \leq \min\{t_{\text{tran}}, t_{\text{ref}}\}, \quad (201)$$

where \tilde{t} is defined in (195). We claim that the following relation is satisfied

$$\hat{\pi}^{(t)}(a_2 | s) - \hat{\pi}^{(t-1)}(a_2 | s) \geq \frac{c_p c_m}{150} \eta(1 - \gamma)^2 \left[\hat{\pi}^{(t-1)}(a_2 | s) \right]^2 \quad (202)$$

for any t obeying $t_{s-1}(\tau_s) \leq t \leq t_{s-1}(\tau_s) + \tilde{t} \leq \min\{t_{\text{tran}}, t_{\text{ref}}\}$. Equipped with this recursive relation, we can invoke Lemma 11 to develop a lower bound on $\hat{\pi}^{(t)}(a_2 | s)$, provided that an initial lower bound is available. In order to do so, in view of the expression (193), we can deduce that

$$\hat{\pi}^{(t_{s-1}(\tau_s))}(a_2 | s) = \exp \left(2\theta^{(t_{s-1}(\tau_s))}(s, a_2) + \theta^{(t_{s-1}(\tau_s))}(s, a_1) \right) \geq \exp \left(\theta^{(t_{s-1}(\tau_s))}(s, a_1) \right),$$

where the last relation is due to the bound $\theta^{(t_{s-1}(\tau_s))}(s, a_2) \geq 0$ (see (43) in Lemma 6). Combining the above two inequalities and applying Lemma 11 (see (69b)), we arrive at $\pi^{(t)}(s, a_2) \geq 1/2$ — and hence $\pi^{(t)}(s, a_2) \geq \pi^{(t)}(s, a_0)$ — as soon as $t - t_{s-1}(\tau_s)$ exceeds

$$\frac{1 + \frac{c_p c_m}{150} \eta(1 - \gamma)^2}{\frac{c_p c_m}{150} \eta(1 - \gamma)^2 \hat{\pi}^{(t_{s-1}(\tau_s))}(a_2 | s)}.$$

This together with the definition of t_{tran} thus indicates that

$$\begin{aligned} t_{\text{tran}} - t_{s-1}(\tau_s) &\leq \frac{1 + \frac{c_p c_m}{150} \eta(1 - \gamma)^2}{\frac{c_p c_m}{150} \eta(1 - \gamma)^2 \hat{\pi}^{(t_{s-1}(\tau_s))}(a_2 | s)} \leq \frac{1 + \frac{c_p c_m}{150} \eta(1 - \gamma)^2}{\frac{c_p c_m}{150} \eta(1 - \gamma)^2 \exp \left(\theta^{(t_{s-1}(\tau_s))}(s, a_1) \right)} \\ &\leq \frac{1.5}{\frac{c_p c_m}{150} \eta(1 - \gamma)^2 \exp \left(\theta^{(t_{s-1}(\tau_s))}(s, a_1) \right)}, \end{aligned}$$

provided that $\frac{c_p c_m}{150} \eta(1 - \gamma)^2 \leq 0.5$. This, however, contradicts the assumption (201). As a consequence, we conclude that $t_{s-1}(\tau_s) + \tilde{t} > \min\{t_{\text{tran}}, t_{\text{ref}}\}$, thus indicating that

$$\min\{t_{\text{tran}}, t_{\text{ref}}\} - t_{s-1}(\tau_s) \leq \tilde{t} \leq \frac{225}{c_p c_m \eta(1 - \gamma)^2 \exp \left(\theta^{(t_{s-1}(\tau_s))}(s, a_1) \right)}. \quad (203)$$

Showing that $t_{\text{tran}} = \min\{t_{\text{tran}}, t_{\text{ref}}\}$. We now justify that $t_{\text{tran}} < t_{\text{ref}}$, so that the upper bound (203) leads to an upper bound on $t_{\text{tran}} - t_{s-1}(\tau_s)$. Suppose instead that

$$t_{\text{tran}} \geq t_{\text{ref}}, \quad \text{or equivalently,} \quad t_{\text{ref}} = \min\{t_{\text{tran}}, t_{\text{ref}}\},$$

and we would like to show that this leads to contradiction. By definition of t_{ref} , we have

$$\theta^{(t_{\text{ref}})}(s, a_0) \leq \theta^{(t_{\text{ref}})}(s, a_1) - \log(c_{\text{ref}}(1 - \gamma)).$$

This further yields

$$\begin{aligned} \max \left\{ \theta^{(t_{\text{ref}})}(s, a_0), \theta^{(t_{\text{ref}})}(s, a_1) \right\} &\leq \theta^{(t_{\text{ref}})}(s, a_1) - \log(c_{\text{ref}}(1 - \gamma)) \\ &\leq \theta^{(t_{s-1}(\tau_s))}(s, a_1) + 1/2 - \log(c_{\text{ref}}(1 - \gamma)) < 0, \end{aligned}$$

where the second inequality arises from (194), and the last one makes use of (188) as long as $t_{s-2}(\tau_{s-2})$. However, this together with the constraint $\sum_a \theta^{(t_{\text{ref}})}(s, a) = 0$ implies that

$$\theta^{(t_{\text{ref}})}(s, a_2) = -\theta^{(t_{\text{ref}})}(s, a_0) - \theta^{(t_{\text{ref}})}(s, a_1) > 0 > \max \left\{ \theta^{(t_{\text{ref}})}(s, a_0), \theta^{(t_{\text{ref}})}(s, a_1) \right\}.$$

which, however, implies that $t_{\text{ref}} > t_{\text{tran}}$ (according to the definition of t_{tran}) and leads to contradiction. As a result, we conclude that

$$t_{\text{tran}} < t_{\text{ref}}, \quad (204)$$

and the bound (203) then indicates that

$$t_{\text{tran}} - t_{s-1}(\tau_s) \leq \frac{225}{c_p c_m \eta(1 - \gamma)^2 \exp \left(\theta^{(t_{s-1}(\tau_s))}(s, a_1) \right)}. \quad (205)$$

F.2.1 Proof of the inequality (202)

From the relation (193), one can deduce that

$$\begin{aligned}
\hat{\pi}^{(t)}(a_2 | s) - \hat{\pi}^{(t-1)}(a_2 | s) &= \exp\left(2\theta^{(t)}(s, a_2) + \theta^{(t)}(s, a_1)\right) - \exp\left(2\theta^{(t-1)}(s, a_2) + \theta^{(t-1)}(s, a_1)\right) \\
&= \hat{\pi}^{(t-1)}(a_2 | s) \left\{ \exp\left(2\theta^{(t)}(s, a_2) - 2\theta^{(t-1)}(s, a_2) + \theta^{(t)}(s, a_1) - \theta^{(t-1)}(s, a_1)\right) - 1 \right\} \\
&\geq \hat{\pi}^{(t-1)}(a_2 | s) \left\{ 2\theta^{(t)}(s, a_2) - 2\theta^{(t-1)}(s, a_2) + \theta^{(t)}(s, a_1) - \theta^{(t-1)}(s, a_1) \right\} \\
&= \hat{\pi}^{(t-1)}(a_2 | s) \cdot \eta \left(2 \frac{\partial V^{(t-1)}(\mu)}{\partial \theta(s, a_2)} + \frac{\partial V^{(t-1)}(\mu)}{\partial \theta(s, a_1)} \right)
\end{aligned} \tag{206}$$

for any t with $t_{s-1}(\tau_s) \leq t \leq \min\{t_{\text{tran}}, t_{\text{ref}}\}$, where the inequality above follows from the elementary fact $e^x - 1 \geq x$ for any $x \in \mathbb{R}$. Therefore, the difference between $\hat{\pi}^{(t)}(a_2 | s)$ and $\hat{\pi}^{(t-1)}(a_2 | s)$ depends on both $\frac{\partial V^{(t-1)}(\mu)}{\partial \theta(s, a_2)}$ and $\frac{\partial V^{(t-1)}(\mu)}{\partial \theta(s, a_1)}$, motivating us to lower bound these two derivatives separately.

Step 1: bounding $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)}$. First, we make the observation that for any $3 \leq s < H$ and any $t \geq t_{s-1}(\tau_s)$,

$$Q^{(t)}(s, a_2) - Q^{(t)}(s, a_0) = \gamma p(V^{(t)}(\overline{s-2}) - \gamma \tau_{s-2}) \geq \gamma p(\gamma^{2s-3} - 1/4 - \gamma \tau_{s-2}) \geq \frac{p}{8} = \frac{c_p(1-\gamma)}{8} \tag{207}$$

holds as long as $\gamma(\gamma^{2s-3} - 1/4 - \gamma \tau_s) \geq 1/8$. Here, the first identity comes from (59) in Lemma 8, and the first inequality holds for any $t \geq t_{s-2}(\gamma^{2s-3} - 1/4)$ — a consequence of the monotonicity property in Lemma 9. As a result, for any t obeying $t_{s-1}(\tau_s) \leq t \leq \min\{t_{\text{tran}}, t_{\text{ref}}\}$ we have

$$\begin{aligned}
Q^{(t)}(s, a_2) - V^{(t)}(s) &= \pi^{(t)}(a_0 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_0) \right) + \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \right) \\
&\geq \frac{c_p(1-\gamma)}{24} - \pi^{(t)}(a_1 | s) \geq \frac{c_p(1-\gamma)}{24} - c_{\text{ref}}(1-\gamma) \geq \frac{c_p(1-\gamma)}{48},
\end{aligned} \tag{208}$$

where the first inequality combines (207) with the facts that $\pi^{(t)}(a_0 | s) \geq 1/3$ (see (192)) and $0 \leq Q^{(t)}(s, a_2), Q^{(t)}(s, a_1) \leq 1$ (see Lemma 1), and the last line holds by observing (see (185))

$$\pi^{(t)}(a_1 | s) \leq c_{\text{ref}}(1-\gamma)\pi^{(t_{s-1}(\tau_s))}(a_0 | s) \leq c_{\text{ref}}(1-\gamma) \quad \text{for all } t \in [t_{s-1}(\tau_s), t_{\text{ref}}]$$

and using the assumption $c_{\text{ref}} \leq c_p/2$. Consequently, for any $t \geq t_{s-1}(\tau_s)$, the gradient w.r.t. $\theta(s, a_2)$ satisfies

$$\begin{aligned}
\frac{\partial V^{(t)}(s)}{\partial \theta(s, a_2)} &= \frac{1}{1-\gamma} d_{\mu}^{(t)}(s) \pi^{(t)}(a_2 | s) \left(Q^{(t)}(s, a_2) - V^{(t)}(s) \right) \\
&\geq \frac{c_p c_m \gamma}{48} (1-\gamma)^2 \pi^{(t)}(a_2 | s) \geq \frac{c_p c_m \gamma}{144} (1-\gamma)^2 \hat{\pi}^{(t)}(a_2 | s) > 0,
\end{aligned} \tag{209}$$

where the first inequality above also makes use of the lower bound in Lemma 2.

In fact, the above lower bound holds true regardless of t , as long as $t \geq t_{s-1}(\tau_s)$ where we have shown that $\frac{\partial V^{(t-1)}(\mu)}{\partial \theta(s, a_2)}$ is bounded from below by 0. One can thus conclude that the iterate $\theta^{(t)}(s, a_2)$ increases with t .

Step 2: bounding $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)}$. Regarding the gradient w.r.t. $\theta(s, a_1)$, we have

$$\begin{aligned}
\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} &= \frac{1}{1-\gamma} d_{\mu}^{(t)}(s) \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_1) - V^{(t)}(s) \right) \\
&= \frac{1}{1-\gamma} d_{\mu}^{(t)}(s) \pi^{(t)}(a_1 | s) \left\{ \pi^{(t)}(a_0 | s) \left(Q^{(t)}(s, a_1) - Q^{(t)}(s, a_0) \right) + \pi^{(t)}(a_2 | s) \left(Q^{(t)}(s, a_1) - Q^{(t)}(s, a_2) \right) \right\} \\
&\geq \frac{1}{1-\gamma} d_{\mu}^{(t)}(s) \pi^{(t)}(a_1 | s) \left(\pi^{(t)}(a_0 | s) + \pi^{(t)}(a_2 | s) \right) (\gamma \tau_s - \gamma^{\frac{1}{2}} \tau_s),
\end{aligned}$$

where the last line follows since (see Lemma 8 and the fact that $t \geq t_{s-1}(\tau_s)$)

$$\max \{Q^{(t)}(s, a_0), Q^\pi(s, a_2)\} \leq \gamma^{\frac{1}{2}} \tau_s, \quad Q^{(t)}(s, a_1) = \gamma V^{(t)}(\overline{s-1}) \geq \gamma \tau_s.$$

In addition, recognizing that $\pi^{(t)}(a_0 | s) + \pi^{(t)}(a_2 | s) \leq 1$ and $d_\mu^{(t)}(s) \leq 14c_m(1-\gamma)^2$ (see Lemma 3), we can continue the above bound to obtain

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} \geq -14c_m(1-\gamma)\pi^{(t)}(a_1 | s)\tau_s\gamma^{\frac{1}{2}} \frac{1-\gamma}{1+\sqrt{\gamma}} \geq -7c_m(1-\gamma)^2 \hat{\pi}^{(t)}(a_1 | s), \quad (210)$$

where the last inequality is due to $\tau_s \leq 1/2$ and $0 < \gamma < 1$ and the bound (132).

Step 3: connecting $\hat{\pi}^{(t)}(a_1 | s)$ with $\hat{\pi}^{(t)}(a_2 | s)$. The above lower bound (210) on $\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)}$ is dependent on $\hat{\pi}^{(t)}(a_1 | s)$. However, the desired lower bound (202) is only a function of $\hat{\pi}^{(t)}(a_2 | s)$. This motivates us to investigate the connection between $\hat{\pi}^{(t)}(a_1 | s)$ and $\hat{\pi}^{(t)}(a_2 | s)$.

To this end, let us write

$$\hat{\pi}^{(t-1)}(a_1 | s) = \hat{\pi}^{(t-1)}(a_2 | s) \exp \left(\theta^{(t-1)}(s, a_1) - \theta^{(t-1)}(s, a_2) \right). \quad (211)$$

As a result, one only needs to control the quantity $\exp \left(\theta^{(t-1)}(s, a_1) - \theta^{(t-1)}(s, a_2) \right)$. In order to do so, we make use of the induction hypothesis (194) for the $(t-1)$ -th iteration to show that

$$\begin{aligned} \exp \left(\theta^{(t-1)}(s, a_1) - \theta^{(t-1)}(s, a_2) \right) &\leq \exp \left(\theta^{(t_{s-1}(\tau_s))}(s, a_1) + 1/2 - \theta^{(t-1)}(s, a_2) \right) \\ &\stackrel{(i)}{\leq} \exp \left(\theta^{(t_{s-1}(\tau_s))}(s, a_1) + 1/2 - \theta^{(t_{s-1}(\tau_s))}(s, a_2) \right) \\ &\stackrel{(ii)}{\leq} \exp \left(\theta^{(t_{s-2}(\tau_{s-2}))}(s, a_1) + 1/2 \right). \end{aligned}$$

Here, (i) follows from the fact that $\theta^{(t)}(s, a_2)$ increases with t (see (209)); and (ii) comes from the inequality (43) in Lemma 6 as well as (188). Recalling Lemma 5, one has

$$\begin{aligned} \exp \left(\theta^{(t-1)}(s, a_1) - \theta^{(t-1)}(s, a_2) \right) &\leq \exp \left(\theta^{(t_{s-2}(\tau_{s-2}))}(s, a_1) + 1/2 \right) \\ &\leq \frac{\exp(1/2)}{\sqrt{1 + \frac{c_m \gamma}{35} \eta (1-\gamma)^2 t_{s-2}(\tau_{s-2})}} \leq \frac{\gamma c_p}{1050}, \end{aligned} \quad (212)$$

where the last inequality is satisfied provided that $t_{s-2}(\tau_{s-2}) \geq \frac{1050^2 e}{\frac{c_m \gamma^3}{35} \eta (1-\gamma)^2 c_p^2}$. Combining (210) with (211) and (212), we arrive at

$$\frac{\partial V^{(t-1)}(\mu)}{\partial \theta(s, a_1)} \geq -\frac{c_p c_m}{150} (1-\gamma)^2 \hat{\pi}^{(t-1)}(a_2 | s). \quad (213)$$

Step 4: combining bounds. Putting the above pieces together and invoking the expression (206) yield for $\gamma > 0.96$,

$$\begin{aligned} \hat{\pi}^{(t)}(a_2 | s) - \hat{\pi}^{(t-1)}(a_2 | s) &\geq \hat{\pi}^{(t-1)}(a_2 | s) \cdot \eta \left(2 \frac{\partial V^{(t-1)}(\mu)}{\partial \theta(s, a_2)} + \frac{\partial V^{(t-1)}(\mu)}{\partial \theta(s, a_1)} \right) \\ &\geq \left[\hat{\pi}^{(t-1)}(a_2 | s) \right]^2 \eta \frac{c_p c_m}{150} (1-\gamma)^2, \end{aligned}$$

which concludes the proof of the advertised bound (202).

F.3 Stage II: the duration where $\theta^{(t)}(s, a_2) \geq \theta^{(t)}(s, a_0)$

We now turn attention to the case where t lies within $[t_{\text{tran}}, t_{\text{ref}})$, which is a non-empty interval according to (204). In this case one has

$$\theta^{(t)}(s, a_2) \geq \theta^{(t)}(s, a_0), \quad \text{or equivalently,} \quad \pi^{(t)}(s, a_2) \geq \pi^{(t)}(s, a_0), \quad (214)$$

as a consequence of the definition (189) of t_{tran} . Again, from the definition (185) of t_{ref} , the inequality $c_{\text{ref}}(1 - \gamma)\pi^{(t)}(a_0 | s) > \pi^{(t)}(a_1 | s)$ holds true for every $t \in [t_{\text{tran}}, t_{\text{ref}})$, or equivalently,

$$\theta^{(t)}(s, a_0) \geq \theta^{(t)}(s, a_1) - \log(c_{\text{ref}}(1 - \gamma)) \quad \text{for all } t \in [t_{\text{tran}}, t_{\text{ref}}). \quad (215)$$

The goal of this subsection is to show that $\theta^{(t)}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) \leq 1/2$ throughout this stage.

Preparation. From the above conditions (214) and (215), we have

$$\pi^{(t)}(s, a_2) \geq \pi^{(t)}(s, a_0) \geq \pi^{(t)}(s, a_1) \quad \text{and hence} \quad \pi^{(t)}(s, a_2) \geq 1/3. \quad (216)$$

We now look at the gradient w.r.t. $\theta(s, a_0)$, for which we first observe that

$$\begin{aligned} Q^{(t)}(s, a_0) - V^{(t)}(s) &= \pi^{(t)}(a_2 | s) \left(Q^{(t)}(s, a_0) - Q^{(t)}(s, a_2) \right) + \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_0) - Q^{(t)}(s, a_1) \right) \\ &\stackrel{(i)}{\leq} -\frac{c_p(1 - \gamma)}{24} + c_{\text{ref}}(1 - \gamma) \stackrel{(ii)}{\leq} -\frac{c_p(1 - \gamma)}{36} < 0. \end{aligned} \quad (217)$$

Here, (i) follows from the inequalities (207) and (216), whereas (ii) holds true as long as $c_{\text{ref}} \leq c_p/72$. Consequently,

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_0)} = \frac{1}{1 - \gamma} d_\mu^{(t)}(s) \pi^{(t)}(a_0 | s) \left(Q^{(t)}(s, a_0) - V^{(t)}(s) \right) < 0,$$

thus indicating that $\theta^{(t)}(s, a_0)$ is decreasing with t .

Key induction hypotheses. Again, we seek to prove by induction that

$$\theta^{(t)}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) \leq 1/2, \quad t \in [t_{\text{tran}}, t_{\text{ref}}). \quad (218)$$

For the base case where $t = t_{\text{tran}}$, this claim trivially holds true. Now suppose that the induction hypothesis (218) is satisfied for every iteration up to $t - 1$, and we would like to establish it for the t -th iteration. Towards this, we find it helpful to introduce another auxiliary induction hypothesis

$$\hat{\pi}^{(i)}(a_0 | s) \leq \frac{1}{1 + \frac{c_p c_m}{288} \eta (1 - \gamma)^2 (i - t_{\text{tran}})} \quad \text{for all } i \in [t_{\text{tran}}, t). \quad (219)$$

As an immediate remark, this hypothesis trivially holds true when $t = t_{\text{tran}} + 1$. In what follows, we shall first establish (218) for the t -th iteration assuming satisfaction of (219), and then use to demonstrate that (219) holds for $i = t$ as well.

Inductive step 1: showing that $\theta^{(t)}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) \leq 1/2$. Towards this, let us introduce for convenience another time instance

$$\tilde{t} := \arg \max_{i: t_{\text{tran}} \leq i < t} \theta^{(i)}(s, a_1), \quad (220)$$

which reflects the time when $\theta^{(i)}(s, a_1)$ reaches its maximum before iteration t . In order to establish the induction hypothesis (218) for the t -th iteration, it is sufficient to demonstrate that

$$\theta^{(\tilde{t})}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) \leq 1/2. \quad (221)$$

As before, let us employ the PG update rule (12a) to expand $\theta^{(\tilde{t})}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1)$ as follows

$$\theta^{(\tilde{t})}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) = \sum_{i=t_{\text{tran}}}^{\tilde{t}-1} \eta \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)}. \quad (222)$$

For each gradient $\frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)}$, invoking Lemma 3, Lemma 1 and Lemma 10 tells us that

$$\frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} = \frac{1}{1-\gamma} d_{\mu}^{(i)}(s) \pi^{(i)}(a_1 | s) (Q^{(i)}(s, a_1) - V^{(i)}(s)) \leq 14c_m(1-\gamma) \pi^{(i)}(a_1 | s). \quad (223)$$

In addition, a little algebra together with (216) leads to

$$\begin{aligned} \pi^{(i)}(a_1 | s) &\leq \hat{\pi}^{(i)}(a_1 | s) = \exp\left(\theta^{(i)}(s, a_1) - \theta^{(i)}(s, a_2)\right) \stackrel{(i)}{=} \exp\left(\frac{3}{2}\theta^{(i)}(s, a_1) + \frac{1}{2}\theta^{(i)}(s, a_0) - \frac{1}{2}\theta^{(i)}(s, a_2)\right) \\ &= \exp\left(\frac{3}{2}\theta^{(i)}(s, a_1)\right) \sqrt{\hat{\pi}^{(i)}(a_0 | s)} \stackrel{(ii)}{\leq} \exp\left(\frac{3}{2}\theta^{(\tilde{t})}(s, a_1)\right) \frac{1}{\sqrt{1 + \frac{c_p c_m}{288} \eta (1-\gamma)^2 (i - t_{\text{tran}})}} \end{aligned}$$

for any i obeying $t_{s-1}(\tau_s) \leq i < \tilde{t}$, where the first inequality comes from (132), (i) makes use of $\sum_a \theta^{(i)}(s, a) = 0$, and (ii) follows from the induction hypothesis (219) along with the definition (220) of \tilde{t} .

Putting the above bounds together with (222) and (223) guarantees that

$$\begin{aligned} \theta^{(\tilde{t})}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) &\leq \sum_{i=t_{\text{tran}}}^{\tilde{t}-1} 14c_m \eta (1-\gamma) \exp\left(\frac{3}{2}\theta^{(\tilde{t})}(s, a_1)\right) \frac{1}{\sqrt{1 + \frac{c_p c_m}{288} \eta (1-\gamma)^2 (i - t_{\text{tran}})}} \\ &= \frac{14c_m \eta \exp\left(\frac{3}{2}\theta^{(\tilde{t})}(s, a_1)\right)}{\sqrt{\frac{c_p c_m}{288} \eta}} \left\{ 1 + \sum_{i=t_{\text{tran}}+1}^{\tilde{t}-1} \frac{1}{\sqrt{i - t_{\text{tran}}}} \right\} \\ &\leq \sqrt{\frac{225792c_m \eta (\tilde{t} - t_{\text{tran}})}{c_p}} \exp\left(\frac{3}{2}\theta^{(\tilde{t})}(s, a_1)\right). \end{aligned} \quad (224)$$

Given that $\theta^{(\tilde{t})}(s, a_0) \geq \theta^{(\tilde{t})}(s, a_1) - \log(c_{\text{ref}}(1-\gamma))$ (see (215)) and $\sum_a \theta^{(\tilde{t})}(s, a) = 0$, one obtains

$$\begin{aligned} \hat{\pi}^{(\tilde{t})}(a_0 | s) &= \exp\left(\theta^{(\tilde{t})}(s, a_0) - \theta^{(\tilde{t})}(s, a_2)\right) = \exp\left(2\theta^{(\tilde{t})}(s, a_0) + \theta^{(\tilde{t})}(s, a_1)\right) \\ &\geq \exp\left(3\theta^{(\tilde{t})}(s, a_1) - 2\log(c_{\text{ref}}(1-\gamma))\right), \end{aligned}$$

which combined with the inequality (219) thus implies that

$$\exp\left(\frac{3}{2}\theta^{(\tilde{t})}(s, a_1)\right) \leq c_{\text{ref}}(1-\gamma) \sqrt{\hat{\pi}^{(\tilde{t})}(a_0 | s)} \leq \frac{c_{\text{ref}}(1-\gamma)}{\sqrt{\frac{c_p c_m}{288} \eta (1-\gamma)^2 (\tilde{t} - t_{\text{tran}})}}. \quad (225)$$

As a consequence of the inequalities (224) and (225), we obtain

$$\begin{aligned} \theta^{(\tilde{t})}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) &\leq \sqrt{\frac{225792c_m \eta (\tilde{t} - t_{\text{tran}})}{c_p}} \frac{c_{\text{ref}}(1-\gamma)}{\sqrt{\frac{c_p c_m}{288} \eta (1-\gamma)^2 (\tilde{t} - t_{\text{tran}})}} \\ &\leq \frac{8064c_{\text{ref}}}{c_p} < \frac{1}{2}, \end{aligned} \quad (226)$$

where the last line holds as long as $c_{\text{ref}} < c_p/16128$. This in turn establishes our induction hypothesis (221) — and hence (220) for the t -th iteration — assuming satisfaction of the hypothesis (219).

Inductive step 2: establishing the upper bound (219). The next step is thus to justify the induction hypothesis (219) when $i = t$. To do so, we first pay attention to the dynamics of $\theta^{(i)}(s, a_0)$ for any $t_{\text{tran}} \leq i \leq t$. Recognizing that $\theta^{(i)}(s, a_2) = \max_a \theta^{(i)}(s, a)$ (see (216)) and $\sum_a \theta^{(i)}(s, a) = 0$, we can express

$$\hat{\pi}^{(i)}(a_0 | s) = \exp \left(\theta^{(i)}(s, a_0) - \theta^{(i)}(s, a_2) \right) = \exp \left(2\theta^{(i)}(s, a_0) + \theta^{(i)}(s, a_1) \right).$$

This allows one to obtain

$$\begin{aligned} \hat{\pi}^{(i)}(a_0 | s) - \hat{\pi}^{(i+1)}(a_0 | s) &= \exp \left(2\theta^{(i)}(s, a_0) + \theta^{(i)}(s, a_1) \right) - \exp \left(2\theta^{(i+1)}(s, a_0) + \theta^{(i+1)}(s, a_1) \right) \\ &= \hat{\pi}^{(i)}(a_0 | s) \left\{ 1 - \exp \left(2\theta^{(i+1)}(s, a_0) - 2\theta^{(i)}(s, a_0) + \theta^{(i+1)}(s, a_1) - \theta^{(i)}(s, a_1) \right) \right\} \\ &= \hat{\pi}^{(i)}(a_0 | s) \left\{ 1 - \exp \left(2\eta \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_0)} + \eta \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} \right) \right\}. \end{aligned} \quad (227)$$

With the above observation in mind, we claim for the moment the following recursive relation

$$\hat{\pi}^{(i)}(a_0 | s) - \hat{\pi}^{(i+1)}(a_0 | s) \geq \frac{c_p c_m}{288} \eta (1 - \gamma)^2 \left[\hat{\pi}^{(i)}(a_0 | s) \right]^2 \quad (228)$$

for any i obeying $t_{\text{tran}} \leq i < t$, whose proof is deferred to the end of this section. If this claim were true, then (67b) in Lemma 11 allows us to conclude the desired bound

$$\hat{\pi}^{(t)}(a_0 | s) \leq \frac{1}{1 + \frac{c_p c_m}{288} \eta (1 - \gamma)^2 (t - t_{\text{tran}})}. \quad (229)$$

Proof of the inequality (228). Combining (217) and the lower bound on $d_\mu^{(i)}(s)$ in Lemma 2, we have

$$\frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_0)} \leq c_m \gamma (1 - \gamma) \pi^{(i)}(a_0 | s) \left(Q^{(i)}(s, a_0) - V^{(i)}(s) \right) \leq -\frac{c_p c_m}{108} (1 - \gamma)^2 \hat{\pi}^{(i)}(a_0 | s),$$

where the last inequality also makes use of (132). In addition, invoking the inequalities (223) and (132) gives

$$\begin{aligned} \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} &\leq 14c_m (1 - \gamma) \pi^{(i)}(a_1 | s) \leq 14c_m (1 - \gamma) \hat{\pi}^{(i)}(a_1 | s) \\ &= 14c_m (1 - \gamma) \hat{\pi}^{(i)}(a_0 | s) \exp \left(\theta^{(i)}(s, a_1) - \theta^{(i)}(s, a_0) \right). \end{aligned} \quad (230)$$

Recall that for any $i \in [t_{\text{tran}}, t_{\text{ref}})$, one has $\theta^{(i)}(s, a_0) \geq \theta^{(i)}(s, a_1) - \log(c_{\text{ref}}(1 - \gamma))$, or equivalently,

$$\exp \left(\theta^{(i)}(s, a_1) - \theta^{(i)}(s, a_0) \right) \leq c_{\text{ref}}(1 - \gamma).$$

It thus follows that

$$\frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} \leq 14c_{\text{ref}} c_m (1 - \gamma)^2 \hat{\pi}^{(i)}(a_0 | s).$$

As a result, the above bounds taken collectively lead to

$$2 \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_0)} + \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} \leq \left[-\frac{c_p c_m}{56} (1 - \gamma)^2 + 14c_{\text{ref}} c_m (1 - \gamma)^2 \right] \hat{\pi}^{(i)}(a_0 | s) \leq -\frac{c_p c_m}{112} (1 - \gamma)^2 \hat{\pi}^{(i)}(a_0 | s),$$

provided that $c_{\text{ref}}/c_p < 1/1568$. In addition, similar to (230), we can easily see that

$$\eta \left| \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} \right| \leq 14\eta c_m (1 - \gamma) \pi^{(i)}(a_1 | s) \leq 14\eta c_m (1 - \gamma) \leq 1/3, \quad (231a)$$

$$\eta \left| \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_0)} \right| \leq 14\eta c_m (1 - \gamma) \pi^{(i)}(a_0 | s) \leq 14\eta c_m (1 - \gamma) \leq 1/3 \quad (231b)$$

as long as $\eta c_m(1 - \gamma) \leq 1/42$.

Substituting the preceding bounds into (227), we immediately arrive at

$$\begin{aligned}\hat{\pi}^{(i)}(a_0 | s) - \hat{\pi}^{(i+1)}(a_0 | s) &= \hat{\pi}^{(i)}(a_0 | s) \left\{ 1 - \exp \left(2\eta \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_0)} + \eta \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} \right) \right\} \\ &\geq \hat{\pi}^{(i)}(a_0 | s) \frac{\eta}{2} \left(-2 \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_0)} - \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} \right) \\ &\geq \frac{\eta c_p c_m}{224} (1 - \gamma)^2 \left[\hat{\pi}^{(i)}(a_0 | s) \right]^2,\end{aligned}$$

where the first inequality holds due to the fact $-1 \leq 2\eta \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_0)} + \eta \frac{\partial V^{(i)}(\mu)}{\partial \theta(s, a_1)} \leq 0$ as well as the elementary inequality $1 - e^x \geq -x/2$ as long as $-1 \leq x \leq 0$. This establishes the inequality (228).

F.4 Proof of the claims (45a) and (45b)

We are now ready to justify the claims (45a) and (45b). Combining (194) and (218), we reach

$$\begin{aligned}\theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) &= \begin{cases} \theta^{(t)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1), & \text{if } t_{s-1}(\tau_s) \leq t \leq t_{\text{tran}} \\ \left(\theta^{(t_{\text{tran}})}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) \right) + \left(\theta^{(t)}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) \right), & \text{if } t_{\text{tran}} \leq t \leq t_{\text{ref}} \end{cases} \\ &\leq \max_{t_{s-1}(\tau_s) \leq i \leq t_{\text{tran}}} \left(\theta^{(i)}(s, a_1) - \theta^{(t_{s-1}(\tau_s))}(s, a_1) \right) + \max_{t_{\text{tran}} \leq i < t_{\text{ref}}} \left(\theta^{(i)}(s, a_1) - \theta^{(t_{\text{tran}})}(s, a_1) \right) \leq 1.\end{aligned}$$

This taken collectively with (188) leads to

$$\theta^{(t_{\text{ref}})}(s, a_1) \leq \theta^{(t_{s-1}(\tau_s))}(s, a_1) + 1 \leq -\frac{1}{2} \log \left(1 + \frac{c_m \gamma}{35} \eta (1 - \gamma) t_{s-2}(\tau_{s-2}) \right) + 1,$$

as claimed in (45b).

In addition, recalling the definition (185) of t_{ref} , we have

$$\theta^{(t_{\text{ref}})}(s, a_0) \leq \theta^{(t_{\text{ref}})}(s, a_1) - \log(c_{\text{ref}}(1 - \gamma)),$$

which clearly satisfies (45a) as long as $c_{\text{ref}} \geq c_p/16128$.

F.5 Proof of the claim (45c)

Finally, we move on to analyze what happens after iteration t_{ref} , for which we focus on tracking the changes of $\hat{\pi}^{(t)}(a_1 | s)$. In this part, let us only consider the set of t satisfying

$$\pi^{(t)}(a_1 | s) \leq \pi^{(t)}(a_2 | s).$$

Note that at time t_{ref} , the inequalities (45a) and (45b) are both satisfied, which together with the property $\pi^{(t)}(a_1 | s) \leq \pi^{(t)}(a_2 | s)$ yield

$$\hat{\pi}^{(t)}(a_1 | s) := \exp \left(\theta^{(t)}(s, a_1) - \max_a \theta^{(t)}(s, a) \right) = \exp \left(\theta^{(t)}(s, a_1) - \theta^{(t)}(s, a_2) \right).$$

Then, if $c_{\text{ref}} < c_p/1000$, we have

$$\begin{aligned}\hat{\pi}^{(t+1)}(a_1 | s) - \hat{\pi}^{(t)}(a_1 | s) &= \exp \left(\theta^{(t+1)}(s, a_1) - \theta^{(t+1)}(s, a_2) \right) - \exp \left(\theta^{(t)}(s, a_1) - \theta^{(t)}(s, a_2) \right) \\ &= \hat{\pi}^{(t)}(a_1 | s) \left\{ \exp \left(\theta^{(t+1)}(s, a_1) - \theta^{(t+1)}(s, a_2) - \theta^{(t)}(s, a_1) + \theta^{(t)}(s, a_2) \right) - 1 \right\} \\ &= \hat{\pi}^{(t)}(a_1 | s) \max \left\{ \exp \left(\eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} - \eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} \right) - 1, 0 \right\} \\ &\leq \hat{\pi}^{(t)}(a_1 | s) \cdot 2\eta \max \left\{ \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} - \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)}, 0 \right\}\end{aligned}$$

$$\leq 56c_m\eta(1-\gamma)^2 \left[\hat{\pi}^{(t)}(a_1 | s) \right]^2. \quad (232)$$

Here, the first inequality holds if $\eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} - \eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} \leq 1$ (given the elementary fact $e^x - 1 \leq 2x$ for any $0 \leq x \leq 1$), and the last line is valid since

$$\begin{aligned} \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} &= \frac{1}{1-\gamma} d_\mu^{(t)}(s) \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_1) - V^{(t)}(s) \right) \stackrel{(i)}{\leq} 14c_m(1-\gamma) \pi^{(t)}(a_1 | s), \\ \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} &= \frac{1}{1-\gamma} d_\mu^{(t)}(s) \left\{ \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \right) + \pi^{(t)}(a_0 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_0) \right) \right\} \\ &\stackrel{(ii)}{\geq} \frac{1}{1-\gamma} d_\mu^{(t)}(s) \pi^{(t)}(a_1 | s) \left(Q^{(t)}(s, a_2) - Q^{(t)}(s, a_1) \right) \stackrel{(iii)}{\geq} -14c_m(1-\gamma) \pi^{(t)}(a_1 | s), \end{aligned}$$

where (ii) holds since $Q^{(t)}(s, a_2) \geq Q^{(t)}(s, a_0)$ (cf. (207)), and (i) and (iii) make use of Lemma 1 and Lemma 3. In addition, these bounds also imply that $\eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_1)} - \eta \frac{\partial V^{(t)}(\mu)}{\partial \theta(s, a_2)} \leq 1$ hold as long as $28\eta c_m(1-\gamma) \leq 1$, thus validating the argument for the first inequality in (232).

Armed with the above recursive relation (232), we can invoke Lemma 11 to show that

$$t_s(\tau_s) - t_{\text{ref}} \geq \frac{\frac{1}{\hat{\pi}^{(t_{\text{ref}})}(a_1 | s)} - \frac{1}{\hat{\pi}^{(t_s(\tau_s))}(a_1 | s)}}{56c_m\eta(1-\gamma)^2} \geq \frac{\frac{1}{\hat{\pi}^{(t_{\text{ref}})}(a_1 | s)} - \frac{2}{1-\gamma}}{56c_m\eta(1-\gamma)^2}, \quad (234)$$

where the last inequality holds since (in view of (132) and (60)).

$$\hat{\pi}^{(t)}(a_1 | s) \geq \pi^{(t)}(a_1 | s) \geq (1-\gamma)/2 \quad \text{for any } t \geq t_s(\tau_s).$$

In order to control $t_s(\tau_s) - t_{\text{ref}}$ via (234), it remains to upper bound $\hat{\pi}^{(t_{\text{ref}})}(a_1 | s)$. Towards this end, it is seen that

$$\begin{aligned} \hat{\pi}^{(t_{\text{ref}})}(a_1 | s) &= \exp \left(\theta^{(t_{\text{ref}})}(s, a_1) - \theta^{(t_{\text{ref}})}(s, a_2) \right) = \exp \left(2\theta^{(t_{\text{ref}})}(s, a_1) + \theta^{(t_{\text{ref}})}(s, a_0) \right) \\ &\leq \exp \left(3\theta^{(t_{\text{ref}})}(s, a_1) - \log \left(\frac{c_p(1-\gamma)}{16128} \right) \right) \\ &\leq \frac{16128e^3}{c_p(1-\gamma) \left(1 + \frac{c_m\gamma}{35} \eta(1-\gamma)^2 t_{s-2}(\tau_{s-2}) \right)^{1.5}} \leq \frac{16128e^3}{c_p(1-\gamma) \left(\frac{c_m\gamma}{35} \eta(1-\gamma)^2 t_{s-2}(\tau_{s-2}) \right)^{1.5}}, \end{aligned} \quad (235)$$

where the first line uses $\sum_a \theta^{(t_{\text{ref}})}(s, a) = 0$, the second line relies on the inequality (45a), and the last one applies the inequality (45b). Substitution into the relation (234) yields

$$t_s(\tau_s) - t_{\text{ref}} \geq 10^{-10} c_p c_m^{0.5} \eta^{0.5} (1-\gamma)^2 \left(t_{s-2}(\tau_{s-2}) \right)^{1.5},$$

thus establishing the advertised bound.