

Breaking the Sample Size Barrier in Model-Based Reinforcement Learning with a Generative Model

Gen Li*
UPenn

Yuting Wei*
UPenn

Yuejie Chi†
CMU

Yuxin Chen*
UPenn

May 2020; Revised: December 2021

Abstract

This paper is concerned with the sample efficiency of reinforcement learning, assuming access to a generative model (or simulator). We first consider γ -discounted infinite-horizon Markov decision processes (MDPs) with state space \mathcal{S} and action space \mathcal{A} . Despite a number of prior works tackling this problem, a complete picture of the trade-offs between sample complexity and statistical accuracy is yet to be determined. In particular, all prior results suffer from a severe sample size barrier, in the sense that their claimed statistical guarantees hold only when the sample size exceeds at least $\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^2}$. The current paper overcomes this barrier by certifying the minimax optimality of two algorithms — a *perturbed* model-based algorithm and a *conservative* model-based algorithm — as soon as the sample size exceeds the order of $\frac{|\mathcal{S}||\mathcal{A}|}{1-\gamma}$ (modulo some log factor). Moving beyond infinite-horizon MDPs, we further study time-inhomogeneous finite-horizon MDPs, and prove that a plain model-based planning algorithm suffices to achieve minimax-optimal sample complexity given any target accuracy level. To the best of our knowledge, this work delivers the first minimax-optimal guarantees that accommodate the entire range of sample sizes (beyond which finding a meaningful policy is information theoretically infeasible).

Keywords: model-based reinforcement learning, minimaxity, policy evaluation, generative model

Contents

1	Introduction	2
2	Problem formulation	5
2.1	Discounted infinite-horizon Markov decision processes	5
2.2	Finite-horizon Markov decision processes	6
2.3	Notation	7
3	Model-based planning in discounted infinite-horizon MDPs	7
3.1	Model-based reinforcement learning: two algorithms	7
3.2	Theoretical guarantees	8
3.3	Comparisons with prior works and implications	10
4	Model-based planning in finite-horizon MDPs	10
4.1	Algorithm: model-based planning	11
4.2	Theoretical guarantees and implications	11
5	Other related works	12

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

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

6	Analysis: infinite-horizon MDPs	12
6.1	Matrix notation and Bellman equations	12
6.2	Analysis: model-based policy evaluation	13
6.3	Analysis: perturbed model-based planning	14
6.3.1	Value function estimation for a policy obeying Bernstein-type conditions	15
6.3.2	Decoupling statistical dependency via (s, a) -absorbing MDPs	16
6.3.3	A tie-breaking argument	17
6.3.4	Proof of Theorem 1	18
6.4	Analysis: conservative model-based planning	19
7	Analysis: finite-horizon MDPs	20
7.1	Matrix notation and Bellman equations	20
7.2	An auxiliary value function sequence obeying Bernstein-type conditions	21
7.3	Proof of Theorem 4	22
8	Discussion	23
A	Preliminary facts	23
B	Proofs of auxiliary lemmas: infinite-horizon MDPs	24
B.1	Proofs of Lemma 1 and Lemma 2	24
B.1.1	Proof of Lemma 11	27
B.2	Proof of Lemma 3	28
B.3	Proof of Lemma 4	29
B.4	Proof of Lemma 5	29
B.5	Proof of Lemma 6	30
B.6	Proof of Lemma 7	33
B.7	Proof of Lemma 8	34
C	Proofs of auxiliary lemmas: finite-horizon MDPs	34
C.1	Proof of Lemma 9	35
C.2	Proof of Lemma 12	37

1 Introduction

Reinforcement learning (RL) (Sutton and Barto, 2018; Szepesvári, 2010), which is frequently modeled as learning and decision making in a Markov decision process (MDP), is garnering growing interest in recent years due to its remarkable success in practice. A core objective of RL is to search for a policy — based on a collection of noisy data samples — that approximately maximizes expected cumulative rewards in an MDP, without direct access to a precise description of the underlying model.¹ In contemporary applications, it is increasingly more common to encounter environments with prohibitively large state and action space, thus exacerbating the challenge of collecting enough samples to learn the model. To enable faithful policy learning in the sample-starved regime (i.e. the regime where the model complexity overwhelms the sample size), it is crucial to obtain a quantitative picture of the fundamental trade-off between sample complexity and statistical accuracy, and to design efficient algorithms that provably achieve the optimal trade-off.

Broadly speaking, there are at least two common algorithmic approaches: a model-based approach and a model-free one. The model-based approach decouples model estimation and policy learning tasks; more specifically, one first estimates the unknown model using the data samples in hand, and then leverages the fitted model to perform planning — a task that can be accomplished by resorting to Bellman’s principle of optimality (Bellman, 1952). A notable advantage of model-based algorithms is their flexibility: the learned model can be adapted to perform new ad-hoc tasks without revisiting the data samples. In comparison, the model-free approach attempts to compute the optimal policy (and the optimal value function) without learning the model explicitly, which lends itself well to scenarios when a model is difficult to estimate or

¹Here and throughout, the “model” refers to the transition kernel and the rewards of the MDP taken collectively.

changes on the fly. Characterizing the sample efficiency of both approaches has been the focal point of a large body of recent works, e.g. Agarwal et al. (2020); Azar et al. (2013); Jin et al. (2018); Kearns and Singh (1999); Li et al. (2021a); Sidford et al. (2018a,b); Tu and Recht (2019); Wainwright (2019a,b).

In this paper, we pursue a comprehensive understanding of model-based RL, given access to a generative model — that is, a simulator that produces samples based on the transition kernel of the true MDP for each state-action pair (Kakade, 2003; Kearns and Singh, 1999). To allow for more precise discussions, we first look at an infinite-horizon discounted MDP with state space \mathcal{S} , action space \mathcal{A} and discount factor $0 < \gamma < 1$, and pay particular attention to the scenarios where the sizes of the state/action spaces and the effective horizon $\frac{1}{1-\gamma}$ are all quite large. We obtain N samples per state-action pair by querying the generative model. For an *arbitrary* target accuracy level $\varepsilon > 0$, a desired model-based planning algorithm should return an ε -optimal policy with a minimal number of calls to the generative model. Particular emphasis is placed on the sub-linear sampling scenario, in which the total sample size is smaller than the total number $|\mathcal{S}|^2|\mathcal{A}|$ of model parameters (so that it is in general infeasible to estimate the model accurately).

Motivation: sample size barriers. Several prior works were dedicated to investigating model-based RL for γ -discounted infinite-horizon MDPs with a generative model, which uncovered the minimax optimality of this approach for an already wide regime (Agarwal et al., 2020; Azar et al., 2013). However, the results therein often suffered from a sample complexity barrier that prevents us from obtaining a complete trade-off curve between sample complexity and statistical accuracy. For instance, the state-of-the-art result Agarwal et al. (2020) required the total sample size to at least exceed $\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^2}$ (up to some log factor), thus restricting the validity of the theory for broader contexts. In truth, this is not merely an issue for model-based planning; the same barrier already showed up when analyzing the simpler task of model-based policy evaluation (Agarwal et al., 2020; Pananjady and Wainwright, 2019). Furthermore, a even more severe barrier emerged in prior theory for model-free methods; for instance, Sidford et al. (2018b); Wainwright (2019b) required the sample size to exceed $\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3}$ modulo some log factor. In stark contrast, however, no lower bounds developed thus far preclude us from attaining reasonable statistical accuracy when going below the aforementioned sample complexity barriers, thus resulting in a gap between upper and lower bounds in this sample-starved regime. Noteworthy, such a sample size barrier is not only present for discounted infinite-horizon MDPs; the situation is similar for finite-horizon MDPs (Yin et al., 2021).

Our contributions. The current paper seeks to achieve optimal sample complexity even below the aforementioned sample size barrier. For γ -discounted infinite-horizon MDPs, we propose two model-based algorithms: (i) *perturbed model-based planning*: which performs planning based on an empirical MDP learned from samples with mild random *reward perturbation*; and (ii) *conservative model-based planning*: which computes approximately optimal policies for the empirical MDP without reward perturbation. These two proposed algorithms provably find an ε -optimal policy with an order of $\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\varepsilon^2}$ samples (up to log factor), thereby matching the minimax lower bound (Azar et al., 2013). Our result accommodates the full range of accuracy level ε (namely, $\varepsilon \in (0, \frac{1}{1-\gamma}]$), thus unveiling the minimaxity of our algorithms as soon as the sample size exceeds $\frac{|\mathcal{S}||\mathcal{A}|}{1-\gamma}$ (modulo some log factor). Encouragingly, this covers the *full* range of sample sizes that enable one to find a policy strictly better than a random guess. See Table 1 for detailed comparisons with prior literature. Along the way, we also derive minimax-optimal statistical guarantees for policy evaluation, which strengthen state-of-the-art results by broadening the applicable sample size range.

Moving beyond discounted infinite-horizon MDPs, we further characterize the sample efficiency of model-based planning for time-inhomogeneous finite-horizon MDPs, which provably achieves minimax-optimal sample complexity as well for the full range of target accuracy levels (Domingues et al., 2021). No reward perturbation or conservative action selection is needed for this finite-horizon scenario. See Table 2 for detailed comparisons with prior literature.

On the technical side, our theory for infinite-horizon MDPs is established upon a novel combination of several key ideas: (1) a high-order expansion of the estimation error for value functions, coupled with fine-grained analysis for each term in the expansion; (2) the construction of auxiliary leave-one-out type (state-action-absorbing) MDPs — motivated by Agarwal et al. (2020) — that help decouple the complicated statistical dependency between the empirically optimal policy (as opposed to value functions) and data samples; (3) a tie-breaking argument guaranteeing that the empirically optimal policy is sufficiently separated

from all other policies under reward perturbation. The case with finite-horizon MDPs is also established based on certain high-order expansion of the value estimation errors, in addition to careful variance control for the terms in the expansion.

Algorithm	Sample size range	Sample complexity	ε -range
Phased Q-learning Kearns and Singh (1999)	$[\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^5}, \infty)$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^7\varepsilon^2}$	$(0, \frac{1}{1-\gamma}]$
Empirical QVI Azar et al. (2013)	$[\frac{ \mathcal{S} ^2 \mathcal{A} }{(1-\gamma)^2}, \infty)$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3\varepsilon^2}$	$(0, \frac{1}{\sqrt{(1-\gamma) \mathcal{S} }}]$
Sublinear randomized value iteration Sidford et al. (2018b)	$[\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^2}, \infty)$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^4\varepsilon^2}$	$(0, \frac{1}{1-\gamma}]$
Variance-reduced QVI Sidford et al. (2018a)	$[\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3}, \infty)$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3\varepsilon^2}$	$(0, 1]$
Randomized primal-dual method Wang (2019)	$[\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^2}, \infty)$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^4\varepsilon^2}$	$(0, \frac{1}{1-\gamma}]$
Empirical MDP + planning Agarwal et al. (2020)	$[\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^2}, \infty)$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3\varepsilon^2}$	$(0, \frac{1}{\sqrt{1-\gamma}}]$
<i>Perturbed</i> empirical MDP + planning This paper	$[\frac{ \mathcal{S} \mathcal{A} }{1-\gamma}, \infty)$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3\varepsilon^2}$	$(0, \frac{1}{1-\gamma}]$
Empirical MDP + <i>conservative</i> planning This paper	$[\frac{ \mathcal{S} \mathcal{A} }{1-\gamma}, \infty)$	$\frac{ \mathcal{S} \mathcal{A} }{(1-\gamma)^3\varepsilon^2}$	$(0, \frac{1}{1-\gamma}]$

Table 1: Comparisons with prior results (up to log factors) regarding finding an ε -optimal policy in a γ -discounted infinite-horizon MDP with a generative model. The sample size range and the ε -range stand for the range of sample size and optimality gap (e.g. ε -accuracy) for the claimed sample complexity to hold. Note that the results in [Kearns and Singh \(1999\)](#); [Wang \(2019\)](#) only hold for a restricted family of MDPs satisfying certain ergodicity assumptions. In addition, [Azar et al. \(2013\)](#) (resp. [Wainwright \(2019b\)](#)) showed that empirical QVI (resp. variance-reduced Q-learning) finds an ε -optimal Q-function estimate with sample complexity $\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\varepsilon^2}$ ($\varepsilon \in (0, 1]$) in a sample size range $[\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3}, \infty)$, which did not translate directly to an ε -optimal policy.

Algorithm	Sample size range	Sample complexity	ε -range
Sublinear randomized value iteration Sidford et al. (2018b)	$[\mathcal{S} \mathcal{A} H^3, \infty)$	$\frac{ \mathcal{S} \mathcal{A} H^5}{\varepsilon^2}$	$(0, H]$
Variance-reduced QVI Sidford et al. (2018a)	$[\mathcal{S} \mathcal{A} H^4, \infty)$	$\frac{ \mathcal{S} \mathcal{A} H^4}{\varepsilon^2}$	$(0, 1]$
Empirical MDP + planning Yin et al. (2021)	$[\mathcal{S} \mathcal{A} H^3, \infty)$	$\frac{ \mathcal{S} \mathcal{A} H^4}{\varepsilon^2}$	$(0, \sqrt{H}]$
Empirical MDP + planning This paper	$[\mathcal{S} \mathcal{A} H^2, \infty)$	$\frac{ \mathcal{S} \mathcal{A} H^4}{\varepsilon^2}$	$(0, H]$

Table 2: Comparisons with prior results (up to log factors) regarding finding an ε -optimal policy in a time-inhomogeneous finite-horizon MDP with a generative model. The sample size range and the ε -range stand for the range of sample size and optimality gap (e.g. ε -accuracy) for the claimed sample complexity to hold. The results in [Sidford et al. \(2018a,b\)](#) were originally stated for the time-homogeneous case; we translate them into the time-inhomogeneous case with an additional factor of H . In addition, [Li et al. \(2021b\)](#) proved that Q-learning finds an ε -optimal Q-function estimate with sample complexity $\frac{|\mathcal{S}||\mathcal{A}|H^4}{\varepsilon^2}$ ($\varepsilon \in (0, 1]$) in a sample size range $[|\mathcal{S}||\mathcal{A}|H^4, \infty)$, which did not translate directly to an ε -optimal policy.

2 Problem formulation

The current paper studies both discounted infinite-horizon MDPs and finite-horizon MDPs, which will be introduced separately in the sequel. Here and throughout, we adopt the standard notation $[H] := \{1, \dots, H\}$.

2.1 Discounted infinite-horizon Markov decision processes

Models and background. Consider a discounted infinite-horizon MDP represented by a quintuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$, where $\mathcal{S} := \{1, 2, \dots, |\mathcal{S}|\}$ denotes a finite set of states, $\mathcal{A} := \{1, 2, \dots, |\mathcal{A}|\}$ is a finite set of actions, $\gamma \in (0, 1)$ stands for the discount factor, and $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ represents the reward function, namely, $r(s, a)$ is the immediate reward received upon executing action a while in state s (here and throughout, we consider the normalized setting where the rewards lie within $[0, 1]$). In addition, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ represents the probability transition kernel of the MDP, where $P(s' | s, a)$ denotes the probability of transiting from state s to state s' when action a is executed, and $\Delta(\mathcal{S})$ denotes the probability simplex over \mathcal{S} .

A deterministic policy (or action selection rule) is a mapping $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maps a state to an action. The value function $V^\pi : \mathcal{S} \rightarrow \mathbb{R}$ of a policy π is defined by

$$\forall s \in \mathcal{S} : \quad V^\pi(s) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s^t, a^t) \mid s^0 = s \right], \quad (1)$$

which is the expected discounted total reward starting from the initial state $s^0 = s$; here, the sample trajectory $\{(s^t, a^t)\}_{t \geq 0}$ is generated based on the transition kernel (namely, $s^{t+1} \sim P(\cdot | s^t, a^t)$), with the actions taken according to policy π (namely, $a^t = \pi(s^t)$ for all $t \geq 0$). It is easily seen that $0 \leq V^\pi(s) \leq \frac{1}{1-\gamma}$. The corresponding action-value function (or Q-function) $Q^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ of a policy π is defined by

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A} : \quad Q^\pi(s, a) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s^t, a^t) \mid s^0 = s, a^0 = a \right], \quad (2)$$

where the actions are taken according to the policy π after the initial action (i.e. $a^t = \pi(s^t)$ for all $t \geq 1$). It is well-known that there exists an optimal policy, denoted by π^* , that simultaneously maximizes $V^\pi(s)$ (resp. $Q^\pi(s, a)$) for all states $s \in \mathcal{S}$ (resp. state-action pairs $(s, a) \in (\mathcal{S} \times \mathcal{A})$) (Sutton and Barto, 2018). The corresponding value function $V^* := V^{\pi^*}$ (resp. action-value function $Q^* := Q^{\pi^*}$) is called the optimal value function (resp. optimal action-value function).

A generative model and an empirical MDP. The current paper focuses on a stylized generative model (also called a simulator) as studied in Kakade (2003); Kearns et al. (2002). Assuming access to this generative model, we collect N independent samples

$$s_{s,a}^i \stackrel{\text{i.i.d.}}{\sim} P(\cdot | s, a), \quad i = 1, \dots, N$$

for each state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, which allows us to construct an empirical transition kernel \widehat{P} as follows

$$\forall s' \in \mathcal{S}, \quad \widehat{P}(s' | s, a) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{s_{s,a}^i = s'\}, \quad (3)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function. In words, $\widehat{P}(s' | s, a)$ counts the empirical frequency of transitions from (s, a) to state s' . The total sample size should therefore be understood as $N^{\text{total}} := N|\mathcal{S}||\mathcal{A}|$. This leads to an empirical MDP $\widehat{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, \widehat{P}, r, \gamma)$ constructed from the data samples. We can define the value function and the action-value function of a policy π for $\widehat{\mathcal{M}}$ analogously, which we shall denote by \widehat{V}^π and \widehat{Q}^π , respectively. The optimal policy of $\widehat{\mathcal{M}}$ is denoted by $\widehat{\pi}^*$, with the optimal value function and Q-function denoted by $\widehat{V}^* := \widehat{V}^{\widehat{\pi}^*}$ and $\widehat{Q}^* := \widehat{Q}^{\widehat{\pi}^*}$, respectively.

Learning the optimal policy via model-based planning. Given a few data samples in hand, the task of policy learning seeks to identify a policy that (approximately) maximizes the expected discounted reward given the data samples. Specifically, for any target level $\varepsilon > 0$, the aim is to compute an ε -accurate policy π_{est} obeying

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A} : \quad V^{\pi_{\text{est}}}(s) \geq V^*(s) - \varepsilon, \quad Q^{\pi_{\text{est}}}(s, a) \geq Q^*(s, a) - \varepsilon. \quad (4)$$

Naturally, one would hope to accomplish these tasks with as few samples as possible. Recall that for the normalized reward setting with $0 \leq r \leq 1$, the value function and Q-function fall within the range $[0, \frac{1}{1-\gamma}]$; this means that the range of the target accuracy level ε should be set to $\varepsilon \in [0, \frac{1}{1-\gamma}]$. The model-based approach typically starts by constructing an empirical MDP $\widehat{\mathcal{M}}$ based on all collected samples, and then “plugs in” this empirical model directly into the Bellman recursion to perform policy evaluation or planning, with prominent examples including Q-value iteration (QVI) and policy iteration (PI) (Bertsekas, 2017).

Aside: policy evaluation. A related task is policy evaluation, which aims to compute or approximate the value function V^π under a given policy π . To be precise, for any target level $\varepsilon > 0$, the goal is to find an ε -accurate estimate V_{est}^π such that

$$\forall s \in \mathcal{S} : \quad |V_{\text{est}}^\pi(s) - V^\pi(s)| \leq \varepsilon. \quad (5)$$

2.2 Finite-horizon Markov decision processes

Models and background. Another type of models considered in this paper is a finite-horizon MDP, which can be represented and denoted by $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \{P_h\}_{h=1}^H, \{r_h\}_{h=1}^H, H)$. Here, \mathcal{S} and \mathcal{A} denote respectively the state space and the action space as before, and H represents the horizon length of the MDP. For any $1 \leq h \leq H$, we let $P_h : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ denote the probability transition kernel at step h , that is, $P_h(s' | s, a)$ is the probability of transiting to s' from (s, a) at step h ; $r_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ indicates the reward function at step h , namely, $r_h(s, a)$ is the immediate reward gained at step h in response to (s, a) . As before, we assume normalized rewards throughout the paper, so that all the $r_h(s, a)$'s reside within the interval $[0, 1]$.

Let $\pi = \{\pi_h\}_{1 \leq h \leq H}$ represent a deterministic policy, such that for any $1 \leq h \leq H$ and any $s \in \mathcal{S}$, $\pi_h(s)$ specifies the action selected at step h in state s . Note that π could be non-stationary, meaning that the π_h 's might be different across different time steps h . The value function and the Q-function associated with policy π are defined respectively by

$$V_h^\pi(s) := \mathbb{E} \left[\sum_{k=h}^H r_k(s_k, a_k) \mid s_h = s \right]$$

for all $s \in \mathcal{S}$ and all $1 \leq h \leq H$, and

$$Q_h^\pi(s, a) := \mathbb{E} \left[\sum_{k=h}^H r_k(s_k, a_k) \mid s_h = s, a_h = a \right]$$

for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ and all $1 \leq h \leq H$. As usual, the expectations are taken over the randomness of the MDP trajectory $\{(s_k, a_k)\}_{1 \leq k \leq H}$ induced by the transition kernel $\{P_h\}_{h=1}^H$ when policy π is adopted. With slight abuse of notation, we let $Q_{H+1}^\pi(s, a) = 0$ for every $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $V_{H+1}^\pi(s) = 0$ for every $s \in \mathcal{S}$. In view of the assumed range of the immediate rewards, it is easily seen that

$$0 \leq V_h^\pi(s) \leq H \quad \text{and} \quad 0 \leq Q_h^\pi(s, a) \leq H$$

for any π , any state-action pair (s, a) , and any step h . Akin to the infinite-horizon counterpart, the optimal value functions $\{V_h^*\}_{1 \leq h \leq H}$ and optimal Q-functions $\{Q_h^*\}_{1 \leq h \leq H}$ are defined respectively by

$$V_h^*(s) := \max_{\pi} V_h^\pi(s) \quad \text{and} \quad Q_h^*(s, a) := \max_{\pi} Q_h^\pi(s, a)$$

for any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ and any $1 \leq h \leq H$. It is well known that there exists at least one policy that allows one to simultaneously achieve the optimal value function and optimal Q-functions for all state-action pairs and all time steps. Throughout this paper, we shall denote by $\pi^* = \{\pi_h^*\}_{1 \leq h \leq H}$ an optimal policy.

A generative model and an empirical MDP. Similar to the infinite-horizon setting, we assume access to a generative model, which is able to generate N independent samples for each triple $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ as follows

$$s_h^i(s, a) \stackrel{\text{i.i.d.}}{\sim} P_h(\cdot | s, a), \quad i = 1, \dots, N.$$

The empirical transition kernel $\{\widehat{P}_h\}_{h=1}^H$ is thus given by

$$\forall s' \in \mathcal{S}, \quad \widehat{P}_h(s' | s, a) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{s_h^i(s, a) = s'\}, \quad (6)$$

which records the empirical frequency of transitions from (s, a) to state s' at step h . This gives rise to a total sample size $N^{\text{total}} := NH|\mathcal{S}||\mathcal{A}|$. We shall let $\widehat{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, \{\widehat{P}_h\}_{h=1}^H, \{r_h\}_{h=1}^H, H)$ represent the empirical MDP constructed from the data samples. The value function and the Q-function of a policy π for $\widehat{\mathcal{M}}$ can be defined analogously, which shall be denoted by $\{\widehat{V}_h^\pi\}_{h=1}^H$ and $\{\widehat{Q}_h^\pi\}_{h=1}^H$, respectively. We denote by $\widehat{\pi}^*$ the optimal policy of $\widehat{\mathcal{M}}$, and the resulting optimal value function and Q-function are denoted by $\widehat{V}_h^* := \widehat{V}_h^{\widehat{\pi}^*}$ and $\widehat{Q}_h^* := \widehat{Q}_h^{\widehat{\pi}^*}$, respectively.

Learning the optimal policy via model-based planning. Given the data samples in hand, the task of policy learning in the finite-horizon case can be defined similarly as the infinite-horizon counterpart. Specifically, for any target level $\varepsilon > 0$, the aim of policy learning is to compute an ε -accurate policy π_{est} obeying

$$\forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]: \quad V_h^{\pi_{\text{est}}}(s) \geq V_h^*(s) - \varepsilon, \quad Q_h^{\pi_{\text{est}}}(s, a) \geq Q_h^*(s, a) - \varepsilon. \quad (7)$$

With the normalized range of the reward function, it is easily seen that the value function and the Q-function reside within the interval $[0, H]$, thus implying that the range of the target accuracy level should be $\varepsilon \in [0, H]$.

2.3 Notation

Let $\mathcal{X} := (|\mathcal{S}|, |\mathcal{A}|, \frac{1}{1-\gamma}, \frac{1}{\varepsilon})$. The notation $f(\mathcal{X}) = O(g(\mathcal{X}))$ means there exists a universal constant $C_1 > 0$ such that $f \leq C_1 g$, whereas the notation $f(\mathcal{X}) = \Omega(g(\mathcal{X}))$ means $g(\mathcal{X}) = O(f(\mathcal{X}))$. In addition, the notation $\widetilde{O}(\cdot)$ (resp. $\widetilde{\Omega}(\cdot)$) is defined in the same way as $O(\cdot)$ (resp. $\Omega(\cdot)$) except that it ignores logarithmic factors.

For any vector $\mathbf{a} = [a_i]_{1 \leq i \leq n} \in \mathbb{R}^n$, we overload the notation $\sqrt{\cdot}$ and $|\cdot|$ in an entry-wise manner such that $\sqrt{\mathbf{a}} := [\sqrt{a_i}]_{1 \leq i \leq n}$ and $|\mathbf{a}| := [|a_i|]_{1 \leq i \leq n}$. For any vectors $\mathbf{a} = [a_i]_{1 \leq i \leq n}$ and $\mathbf{b} = [b_i]_{1 \leq i \leq n}$, the notation $\mathbf{a} \geq \mathbf{b}$ (resp. $\mathbf{a} \leq \mathbf{b}$) means $a_i \geq b_i$ (resp. $a_i \leq b_i$) for all $1 \leq i \leq n$, and we let $\mathbf{a} \circ \mathbf{b} := [a_i b_i]_{1 \leq i \leq n}$ represent the Hadamard product. Additionally, we denote by $\mathbf{1}$ the all-one vector, and \mathbf{I} the identity matrix. For any matrix \mathbf{A} , we define the norm $\|\mathbf{A}\|_1 := \max_i \sum_j |A_{i,j}|$.

3 Model-based planning in discounted infinite-horizon MDPs

As summarized in Table 1, the theory of all prior works required the sample size per state-action pair to at least exceed $N \geq \Omega(\frac{1}{(1-\gamma)^2})$. In order to break this sample size barrier, we develop two model-based algorithms that provably overcome such a sample size barrier.

3.1 Model-based reinforcement learning: two algorithms

Algorithm 1: perturbed model-based planning. The first algorithm applies model-based planning to an empirical MDP with *randomly perturbed rewards*. Specifically, for each state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, we randomly perturb the immediate reward by

$$r_p(s, a) = r(s, a) + \zeta(s, a), \quad \zeta(s, a) \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, \xi), \quad (8)$$

where $\text{Unif}(0, \xi)$ denotes the uniform distribution between 0 and some parameter $\xi > 0$ (to be specified momentarily).² For any policy π , we denote by \widehat{V}_p^π the corresponding value function of the perturbed

²Note that perturbation is only invoked when running the planning algorithms and does not require collecting new samples.

empirical MDP $\widehat{\mathcal{M}}_p = (\mathcal{S}, \mathcal{A}, \widehat{P}, r_p, \gamma)$ with the probability transition kernel \widehat{P} (cf. (3)) and the perturbed reward function r_p . Let $\widehat{\pi}_p^*$ represent the optimal policy of $\widehat{\mathcal{M}}_p$, i.e.

$$\widehat{\pi}_p^* := \arg \max_{\pi} \widehat{V}_p^{\pi}. \quad (9)$$

Algorithm 2: conservative model-based planning. An alternative approach that eliminates the need of reward perturbation is to select *approximately optimal* actions for the empirical MDP instead of the absolute optimal actions. To be precise, denote by \widehat{Q}^* (resp. \widehat{V}^*) the optimal action-value (resp. value) function of the empirical MDP $\widehat{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, \widehat{P}, r, \gamma)$ with the probability transition kernel \widehat{P} (cf. (3)) and the original reward function r . By producing a random draw from $\varsigma \sim \text{Unif}(0, \xi)$ (with ξ specified shortly), we can generate the following policy $\widehat{\pi}_c$:

$$\forall s \in \mathcal{S}: \quad \widehat{\pi}_c(s) := \min \{a \in \mathcal{A} : \widehat{Q}^*(s, a) > \widehat{V}^*(s) - \varsigma\}. \quad (10)$$

Note that there is an index assigned to each action as $\mathcal{A} = \{1, \dots, |\mathcal{A}|\}$ which induces a natural order for all actions. In words, this approach is more conservative and does not stick to the optimal actions w.r.t. the empirical MDP; instead, the policy $\widehat{\pi}_c$ picks out — for each state $s \in \mathcal{S}$ — the smallest indexed action that is within a gap of ς from optimal.

3.2 Theoretical guarantees

Indeed, both the above-mentioned approaches result in a value function (resp. Q-function) that well approximates the true optimal value function V^* (resp. optimal Q-function Q^*). We start by presenting our results for the perturbed model-based approach.

Theorem 1 (Perturbed model-based planning). *There exist some universal constants $c_0, c_1 > 0$ such that: for any $\delta > 0$ and any $0 < \varepsilon \leq \frac{1}{1-\gamma}$, the policy $\widehat{\pi}_p^*$ defined in (9) obeys*

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad V^{\widehat{\pi}_p^*}(s) \geq V^*(s) - \varepsilon \quad \text{and} \quad Q^{\widehat{\pi}_p^*}(s, a) \geq Q^*(s, a) - \gamma\varepsilon, \quad (11)$$

with probability at least $1 - \delta$, provided that the perturbation size is $\xi = \frac{c_1(1-\gamma)\varepsilon}{|\mathcal{S}|^5|\mathcal{A}|^5}$ and that the sample size per state-action pair exceeds

$$N \geq \frac{c_0 \log\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)\varepsilon\delta}\right)}{(1-\gamma)^3\varepsilon^2}. \quad (12)$$

In addition, both the empirical QVI and PI algorithms w.r.t. $\widehat{\mathcal{M}}_p$ (cf. (Azar et al., 2013, Algorithms 1-2)) are able to recover $\widehat{\pi}_p^*$ perfectly within $O\left(\frac{1}{1-\gamma} \log\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)\varepsilon\delta}\right)\right)$ iterations.

Remark 1. Theorem 1 holds unchanged if ξ is taken to be $\frac{c_1(1-\gamma)\varepsilon}{|\mathcal{S}|^\alpha|\mathcal{A}|^\alpha}$ for any fixed constant $\alpha \geq 1$. This paper picks the specific choice $\alpha = 5$ merely to convey that a very small degree of perturbation suffices for our purpose.

Remark 2. Perturbation brings a side benefit: one can recover the optimal policy $\widehat{\pi}_p^*$ of the perturbed empirical MDP $\widehat{\mathcal{M}}_p$ *exactly* in a small number of iterations without incurring further optimization errors. To give a flavor of the overall computational complexity, let us take QVI for example Azar et al. (2013). Recall that each iteration of QVI takes time proportional to the time taken to read \widehat{P} (which is a matrix with at most $N|\mathcal{S}||\mathcal{A}|$ nonzeros), hence the resulting computational complexity can be as low as $O\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^4\varepsilon^2} \log^2\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)\varepsilon\delta}\right)\right)$.

Further, similar performance guarantees can be established for the conservative model-based approach without reward perturbation, as stated below.

Theorem 2 (Conservative model-based planning). *Under the same assumptions of Theorem 1 (including both the sample size and the choice of ξ), the policy $\widehat{\pi}_c$ defined in (10) achieves*

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad V^{\widehat{\pi}_c}(s) \geq V^*(s) - \varepsilon \quad \text{and} \quad Q^{\widehat{\pi}_c}(s, a) \geq Q^*(s, a) - \gamma\varepsilon \quad (13)$$

with probability at least $1 - \delta$.

In a nutshell, the above theorems demonstrate that: both model-based algorithms we introduce succeed in finding an ε -optimal policy as soon as the total sample complexity exceeds the order of $\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3\varepsilon^2}$ (modulo some log factor). It is worth emphasizing that, compared to prior literature, our result imposes no restriction on the range of ε and, in particular, we allow the accuracy level ε to go all the way up to $\frac{1}{1-\gamma}$. Our result is particularly useful in the regime with small-to-moderate sample sizes, since its validity is guaranteed as long as

$$N \geq \tilde{\Omega}\left(\frac{1}{1-\gamma}\right). \quad (14)$$

Tackling the sample-limited regime (in particular, the scenario when $N \in [\frac{1}{1-\gamma}, \frac{1}{(1-\gamma)^2}]$) requires us to develop new analysis frameworks beyond prior theory, which we shall discuss in detail momentarily.

We remark that the work [Azar et al. \(2013\)](#) established a minimax lower bound of the same order as (12) (up to some log factor) in the regime $\varepsilon = O(1)$. A closer inspection of their analysis, however, reveals that their argument and bound hold true as long as $\varepsilon = O(\frac{1}{1-\gamma})$. This in turn corroborates the *minimax optimality* of our perturbed model-based approach for the full ε -range (which is previously unavailable), and demonstrates the information-theoretic infeasibility to learn a policy strictly better than a random guess if $N \leq \tilde{O}(\frac{1}{1-\gamma})$. Put another way, the condition (14) contains the full range of “meaningful” sample sizes.

Finally, we single out an intermediate result in the analysis of our theorems concerning model-based policy evaluation, which might be of interest on its own. Specifically, for any fixed policy π independent of the data, this task concerns value function estimation via the plug-in estimate \hat{V}^π (i.e. the value function of the empirical \mathcal{M} under this policy). However simple as this might seem, existing theoretical underpinnings of this approach remain suboptimal, unless the sample size is already sufficiently large. Our result is the following, which does not require enforcing reward perturbation.

Theorem 3 (Model-based policy evaluation). *Fix any policy π . There exists some universal constant $c_0 > 0$ such that: for any $0 < \delta < 1$ and any $0 < \varepsilon \leq \frac{1}{1-\gamma}$, one has*

$$\forall s \in \mathcal{S} : \quad |\hat{V}^\pi(s) - V^\pi(s)| \leq \varepsilon \quad (15)$$

with probability at least $1 - \delta$, provided that the sample size per state-action pair exceeds

$$N \geq c_0 \frac{\log\left(\frac{|\mathcal{S}|\log\frac{e}{1-\gamma}}{\delta}\right)}{(1-\gamma)^3\varepsilon^2}. \quad (16)$$

In words, this theorem reveals that \hat{V}^π begins to outperform a random guess as soon as $N \geq \tilde{\Omega}\left(\frac{1}{1-\gamma}\right)$. The sample complexity bound (16) enjoys *full coverage* of the ε -range $(0, \frac{1}{1-\gamma}]$, and matches the minimax lower bound derived in ([Pananjady and Wainwright, 2019](#), Theorem 2(b)) up to only a $\log \log \frac{1}{1-\gamma}$ factor. In addition, a recent line of work investigated instance-dependent guarantees for policy evaluation ([Khamarui et al. \(2020\)](#); [Pananjady and Wainwright \(2019\)](#)). While this is not our focus, our analysis does uncover an instance-dependent bound with a broadened sample size range. See Lemma 1 and the discussion thereafter.

Numerical experiments. To further demonstrate the effectiveness of our model-based approach, we conduct a series of numerical experiments (motivated by [Azar et al. \(2013\)](#)), focusing on the effect of the discount complexity $\frac{1}{1-\gamma}$ upon the sample complexity. More specifically, consider the following MDP $(\mathcal{S}, \mathcal{A}, P, r, \gamma)$, where $\mathcal{S} = \{0, 1\}$, $\mathcal{A}_0 = \{0\}$, $\mathcal{A}_1 = \{0, 1\}$, $P(0|0, 0) = 1$, $P(1|1, 0) = p$, $P(1|1, 1) = q$ and $r(0, 0) = 0$, $r(1, 0) = r(1, 1) = 1$ (see Fig. 1 for an illustration), and we take the quantities p and q to be $p = \gamma + \frac{2\gamma(1-\gamma)^2\varepsilon}{(1+\gamma)^2}$ and $q = \gamma - \frac{2\gamma(1-\gamma)^2\varepsilon}{(1+\gamma)^2}$ as suggested by [Azar et al. \(2013\)](#). For each point $(N, \frac{1}{1-\gamma}, \varepsilon)$, we perform 100 Monte Carlo simulations, and we declare N to successfully attain ε accuracy if this accuracy is achieved in at least 95% of the trials. In Fig. 1, the solid lines depict the minimum sample size per state-action pair needed to achieve ε accuracy numerically as the discount complexity varies. As can be seen, the numerical sample complexity per state-action pair scales on the order of $\frac{1}{(1-\gamma)^3\varepsilon^2}$ for varying choices of ε , which is consistent with our theory.

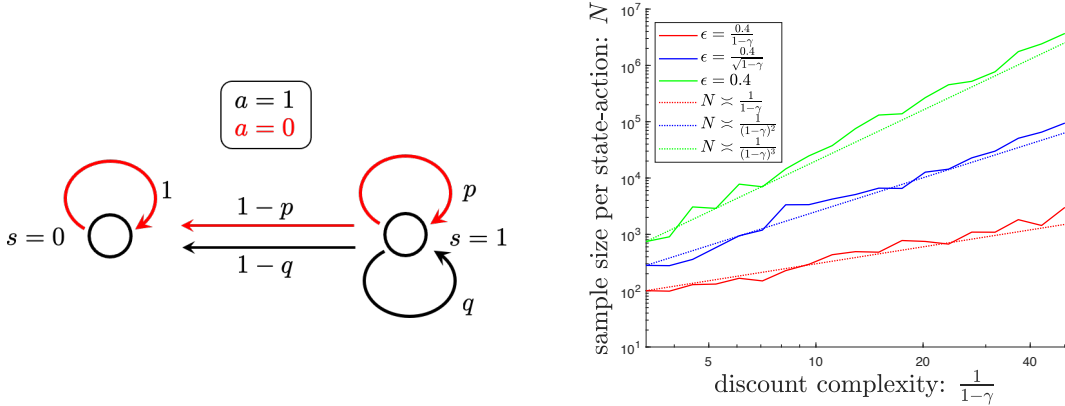


Figure 1: Numerical sample complexity per state-action pair N vs. discount complexity $\frac{1}{1-\gamma}$.

3.3 Comparisons with prior works and implications

In order to discuss the novelty of our results in context, we take a moment to compare them with prior theory. See Table 1 for a more complete list of comparisons.

Prior bounds for planning and policy learning. None of the prior results with a generative model (including both model-based or model-free approaches) was capable of efficiently finding the desired policy while accommodating the full sample size range (14). For instance, the state-of-the-art analysis for the model-based approach Agarwal et al. (2020) required the sample size to at least exceed

$$N \geq \tilde{\Omega}\left(\frac{1}{(1-\gamma)^2}\right), \quad (17)$$

whereas the theory for the variance-reduced model-free approach Sidford et al. (2018a); Wainwright (2019b) imposed the sample size requirement

$$N \geq \tilde{\Omega}\left(\frac{1}{(1-\gamma)^3}\right). \quad (18)$$

In fact, it was previously unknown what is achievable in the sample size range $N \in [\frac{1}{1-\gamma}, \frac{1}{(1-\gamma)^2}]$. In contrast, our results confirm the minimax-optimal statistical performance of the model-based approach with full coverage of the ε -range and the sample size range.

Remark 3. We briefly point out why the sample size barrier (17) appeared in the analysis of Agarwal et al. (2020). Take Agarwal et al. (2020) Section 4.3 for example: the contraction factor $\gamma \sqrt{\frac{8 \log(|\mathcal{S}||\mathcal{A}|/(1-\gamma)\delta)}{N}} \frac{1}{1-\gamma}$ therein needs to be smaller than 1, thereby requiring $N \geq \tilde{\Omega}((1-\gamma)^{-2})$.

Prior bounds for policy evaluation. Regarding value function estimation for any fixed policy π , the prior results Agarwal et al. (2020); Azar et al. (2013); Pananjady and Wainwright (2019) for the plug-in approach all operated under the assumption that $N \geq \tilde{\Omega}(\frac{1}{(1-\gamma)^2})$, which is more stringent than our result by a factor of at least $\frac{1}{1-\gamma}$. In addition, our sample complexity matches the state-of-the-art guarantees in the regime where $\varepsilon \leq \frac{1}{\sqrt{1-\gamma}}$ (Agarwal et al., 2020; Pananjady and Wainwright, 2019), while extending them to the range $\varepsilon \in [\frac{1}{\sqrt{1-\gamma}}, \frac{1}{1-\gamma}]$ uncovered in these previous papers.

4 Model-based planning in finite-horizon MDPs

Moving beyond discounted infinite-horizon MDPs, our theoretical framework is also able to accommodate finite-horizon MDPs, which we detail in this section.

4.1 Algorithm: model-based planning

The algorithm considered in this section is model-based planning (without reward perturbation). Specifically, this model-based approach returns a policy $\hat{\pi}^* = \{\hat{\pi}_h^*\}_{1 \leq h \leq H}$ by means of the following two steps:

- 1) Construct the empirical MDP $\widehat{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, \{\widehat{P}_h\}_{1 \leq h \leq H}, \{r_h\}_{1 \leq h \leq H}, H)$ based on the data samples in hand (see (6) for the computation of the empirical transition kernel \widehat{P}_h);
- 2) Run a classical dynamic programming algorithm (Bertsekas, 2017) to find an optimal policy $\hat{\pi}^*$ of the empirical MDP $\widehat{\mathcal{M}}$.

Note that $\hat{\pi}_h^*$ is an optimal policy of $\widehat{\mathcal{M}}$ at step h , computed by the dynamic programming algorithm calculated backward from $h = H$. Since $\hat{\pi}_h^*$ is calculated solely based on what happens after step h , $\hat{\pi}_h^*$ is independent of the empirical transitions $\{\widehat{P}_j\}_{1 \leq j < h}$.

It is noteworthy that, in contrast to the infinite-horizon counterpart in Section 3, we do not need to enforce random reward perturbation for this finite-horizon case.

4.2 Theoretical guarantees and implications

The model-based algorithm described above turns out to be nearly minimax optimal, as asserted by the following theorem.

Theorem 4 (Model-based planning). *There exist some universal constants $c_0, c_1 > 0$ such that: for any $\delta > 0$ and any $0 < \varepsilon \leq H$, the aforementioned policy $\hat{\pi}^*$ returned by model-based planning obeys*

$$\forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H] : \quad V_h^{\hat{\pi}^*}(s) \geq V_h^*(s) - \varepsilon \quad \text{and} \quad Q_h^{\hat{\pi}^*}(s, a) \geq Q_h^*(s, a) - \varepsilon \quad (19)$$

with probability at least $1 - \delta$, provided that the sample size for every triple (s, a, h) exceeds

$$N \geq \frac{c_0 H^3 \log\left(\frac{H|\mathcal{S}||\mathcal{A}|}{\delta}\right)}{\varepsilon^2}. \quad (20)$$

Akin to the discounted infinite-horizon scenario, the model-based approach manages to achieve ε -accuracy as long as the sample size per (s, a, h) exceeds the order of

$$\frac{H^3}{\varepsilon^2} \quad (\text{up to some log factor}).$$

This result, which is valid for the full ε range $(0, H]$, is reminiscent of the bound (12), except that the effective horizon $\frac{1}{1-\gamma}$ needs to be replaced by the horizon length H . Given that there are in total $|\mathcal{S}||\mathcal{A}|H$ different combinations of (s, a, h) , the total sample complexity is on the order of $\tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^4}{\varepsilon^2}\right)$.

The quadruple scaling H^4 of this total sample complexity — as opposed to the cubic scaling in the discounted infinite-horizon case — is due to time inhomogeneity; that is, the P_h 's might be different across h , resulting in an additional H factor. Again, our result kicks in as soon as the sample size satisfies

$$N \geq \tilde{\Omega}(H), \quad (21)$$

improving upon the sample size requirement

$$N \geq \tilde{\Omega}(H^2) \quad (22)$$

in the state-of-the-art analysis for the model-based approach Yin et al. (2021).

5 Other related works

Classical analyses of reinforcement learning algorithms have largely focused on asymptotic performance (e.g. Jaakkola et al. (1994); Szepesvári (1998); Tsitsiklis and Van Roy (1997); Tsitsiklis (1994)). Leveraging the toolkit of concentration inequalities, a number of recent papers have shifted attention towards understanding the performance in the non-asymptotic and finite-time settings. A highly incomplete list includes Azar et al. (2017); Beck and Srikant (2012); Bhandari et al. (2018); Bradtke and Barto (1996); Cai et al. (2019); Chen et al. (2020); Dalal et al. (2018); Even-Dar and Mansour (2003); Fan et al. (2019); Gupta et al. (2019); Jin et al. (2018); Kaledin et al. (2020); Kearns and Singh (1999); Khamaru et al. (2020); Lakshminarayanan and Szepesvari (2018); Li et al. (2021a,c,d); Mou et al. (2020); Qu and Wierman (2020); Shah and Xie (2018); Sidford et al. (2018a); Srikant and Ying (2019); Strehl et al. (2006); Wainwright (2019b); Xu and Gu (2020); Xu et al. (2019), a large fraction of which is concerned with model-free algorithms.

The generative model (or simulator) adopted in this paper was first proposed in Kearns and Singh (1999), which has been invoked in Agarwal et al. (2020); Azar et al. (2012, 2013); Kakade (2003); Kearns et al. (2002); Kearns and Singh (1999); Khamaru et al. (2020); Lattimore and Hutter (2012); Pananjady and Wainwright (2019); Sidford et al. (2018a,b); Wainwright (2019b); Wang (2019); Yang and Wang (2019), to name just a few. In particular, Azar et al. (2013) developed the minimax lower bound on the sample complexity $N = \Omega\left(\frac{|\mathcal{S}||\mathcal{A}|\log(|\mathcal{S}||\mathcal{A}|)}{(1-\gamma)^3\varepsilon^2}\right)$ necessary for finding an ε -optimal policy, and showed that, for any $\varepsilon \in (0, 1)$, a model-based approach (e.g. applying QVI or PI to the empirical MDP) can estimate the optimal Q-function to within an ε -accuracy given near-minimal samples. Note, however, that directly translating this result to the policy guarantees leads to an additional factor of $\frac{1}{1-\gamma}$ in estimation accuracy and of $\frac{1}{(1-\gamma)^2}$ in sample complexity. In light of this, Azar et al. (2013) further showed that a near-optimal sample complexity is possible for policy learning if the sample size is at least on the order of $\frac{|\mathcal{S}|^2|\mathcal{A}|}{(1-\gamma)^2}$ which, however, is no longer sub-linear in the model complexity. A recent breakthrough Agarwal et al. (2020) substantially improved the model-based guarantee with the aid of auxiliary state-absorbing MDPs, extending the range of sample complexity to $\left[\frac{|\mathcal{S}||\mathcal{A}|\log(|\mathcal{S}||\mathcal{A}|)}{(1-\gamma)^2}, \infty\right)$. Our analysis is motivated in part by Agarwal et al. (2020), but also relies on several other novel techniques to complete the picture.

Finally, we remark that the construction of state-absorbing MDPs or state-action-absorbing MDPs falls under the category of “leave-one-out” type analysis, which is particularly effective in decoupling complicated statistical dependency in various statistical estimation problems Agarwal et al. (2020); Chen et al. (2021, 2019); El Karoui (2015); Ma et al. (2020); Pananjady and Wainwright (2019). The application of such an analysis framework to MDPs should be attributed to Agarwal et al. (2020). Other applications to Markov chains include Chen et al. (2019); Pananjady and Wainwright (2019).

6 Analysis: infinite-horizon MDPs

This section presents the key ideas for proving our main results, following an introduction of some convenient matrix notation.

6.1 Matrix notation and Bellman equations

It is convenient to present our proof based on some matrix notation for MDPs. Denoting by $e_1, \dots, e_{|\mathcal{S}|} \in \mathbb{R}^{|\mathcal{S}|}$ the standard basis vectors, we can define:

- $\mathbf{r} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$: a vector representing the reward function r (so that $r_{(s,a)} = r(s, a)$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$).
- $\mathbf{V}^\pi \in \mathbb{R}^{|\mathcal{S}|}$: a vector representing the value function V^π (so that $V_s^\pi = V^\pi(s)$ for all $s \in \mathcal{S}$).
- $\mathbf{Q}^\pi \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$: a vector representing the Q-function Q^π (so that $Q_{(s,a)}^\pi = Q^\pi(s, a)$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$).
- $\mathbf{V}^* \in \mathbb{R}^{|\mathcal{S}|}$ and $\mathbf{Q}^* \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$: representing the optimal value function V^* and optimal Q-function Q^* .
- $\mathbf{P} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times |\mathcal{S}|}$: a matrix representing the probability transition kernel P , where the (s, a) -th row of \mathbf{P} is a probability vector representing $P(\cdot|s, a)$. Denote $\mathbf{P}_{s,a}$ as the (s, a) -th row of the transition matrix \mathbf{P} .

- $\mathbf{\Pi}^\pi \in \{0, 1\}^{|\mathcal{S}| \times |\mathcal{S}| \times |\mathcal{A}|}$: a projection matrix associated with a given policy π taking the following form

$$\mathbf{\Pi}^\pi = \begin{pmatrix} \mathbf{e}_{\pi(1)}^\top & & & \\ & \mathbf{e}_{\pi(2)}^\top & & \\ & & \ddots & \\ & & & \mathbf{e}_{\pi(|\mathcal{S}|)}^\top \end{pmatrix}. \quad (23)$$

- $\mathbf{P}^\pi \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}| \times |\mathcal{S}| \times |\mathcal{A}|}$ and $\mathbf{P}_\pi \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$: two *square* probability transition matrices induced by the policy π over the state-action pairs and the states respectively, defined by

$$\mathbf{P}^\pi := \mathbf{P}\mathbf{\Pi}^\pi \quad \text{and} \quad \mathbf{P}_\pi := \mathbf{\Pi}^\pi \mathbf{P}. \quad (24)$$

- $\mathbf{r}_\pi \in \mathbb{R}^{|\mathcal{S}|}$: a reward vector restricted to the actions chosen by the policy π , namely, $r_\pi(s) = r(s, \pi(s))$ for all $s \in \mathcal{S}$ (or simply, $\mathbf{r}_\pi = \mathbf{\Pi}^\pi \mathbf{r}$).

Armed with the above matrix notation, we can write, for any policy π , the *Bellman consistency equation* as

$$\mathbf{Q}^\pi = \mathbf{r} + \gamma \mathbf{P}\mathbf{V}^\pi = \mathbf{r} + \gamma \mathbf{P}^\pi \mathbf{Q}^\pi, \quad (25)$$

which implies that

$$\mathbf{Q}^\pi = (\mathbf{I} - \gamma \mathbf{P}^\pi)^{-1} \mathbf{r}; \quad (26)$$

$$\mathbf{V}^\pi = \mathbf{r}_\pi + \gamma \mathbf{P}_\pi \mathbf{V}^\pi \quad \text{and} \quad \mathbf{V}^\pi = (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}_\pi. \quad (27)$$

For a vector $\mathbf{V} = [V_i]_{1 \leq i \leq |\mathcal{S}|} \in \mathbb{R}^{|\mathcal{S}|}$, we define the vector $\text{Var}_{\mathbf{P}}(\mathbf{V}) \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$ whose entries are given by

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad [\text{Var}_{\mathbf{P}}(\mathbf{V})]_{(s,a)} := \sum_{s' \in \mathcal{S}} P(s'|s, a) V_{s'}^2 - \left(\sum_{s' \in \mathcal{S}} P(s'|s, a) V_{s'} \right)^2,$$

i.e. the variance of \mathbf{V} w.r.t. $P(\cdot|s, a)$. This can be expressed using our matrix notation as follows

$$\text{Var}_{\mathbf{P}}(\mathbf{V}) = \mathbf{P}(\mathbf{V} \circ \mathbf{V}) - (\mathbf{P}\mathbf{V}) \circ (\mathbf{P}\mathbf{V}). \quad (28)$$

Similarly, for any given policy π we define

$$\text{Var}_{\mathbf{P}_\pi}(\mathbf{V}) = \mathbf{\Pi}^\pi \text{Var}_{\mathbf{P}}(\mathbf{V}) = \mathbf{P}_\pi(\mathbf{V} \circ \mathbf{V}) - (\mathbf{P}_\pi \mathbf{V}) \circ (\mathbf{P}_\pi \mathbf{V}) \in \mathbb{R}^{|\mathcal{S}|}. \quad (29)$$

We shall also define $\widehat{\mathbf{V}}^\pi$, $\widehat{\mathbf{Q}}^\pi$, $\widehat{\mathbf{V}}^*$, $\widehat{\mathbf{Q}}^*$, $\widehat{\mathbf{P}}$, $\widehat{\mathbf{P}}^\pi$, $\widehat{\mathbf{P}}_\pi$, $\text{Var}_{\widehat{\mathbf{P}}}(\mathbf{V})$, $\text{Var}_{\widehat{\mathbf{P}}_\pi}(\mathbf{V})$ w.r.t. the empirical MDP $\widehat{\mathcal{M}}$ in an analogous fashion.

6.2 Analysis: model-based policy evaluation

We start with the simpler task of policy evaluation, which also plays a crucial role in the analysis of planning. To establish our guarantees in Theorem 3, we aim to prove the following result. Here, we recall that the true value function under a policy π and the model-based empirical estimate are given respectively by

$$\mathbf{V}^\pi = (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}_\pi \quad \text{and} \quad \widehat{\mathbf{V}}^\pi = (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \mathbf{r}_\pi. \quad (30)$$

Lemma 1. *Fix any policy π . Consider any $0 < \delta < 1$, and suppose $N \geq \frac{32e^2}{1-\gamma} \log\left(\frac{4|\mathcal{S}| \log(\frac{e}{1-\gamma})}{\delta}\right)$. Then with probability at least $1 - \delta$, the vectors defined in (30) obey*

$$\begin{aligned} \|\widehat{\mathbf{V}}^\pi - \mathbf{V}^\pi\|_\infty &\leq 4\gamma \sqrt{\frac{2 \log\left(\frac{4|\mathcal{S}| \log(\frac{e}{1-\gamma})}{\delta}\right)}{N}} \left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^\pi]} \right\|_\infty + \frac{2\gamma \log\left(\frac{4|\mathcal{S}| \log(\frac{e}{1-\gamma})}{\delta}\right)}{(1-\gamma)N} \|\mathbf{V}^\pi\|_\infty \\ &\leq 6 \sqrt{\frac{2 \log\left(\frac{4|\mathcal{S}| \log(\frac{e}{1-\gamma})}{\delta}\right)}{N(1-\gamma)^3}}. \end{aligned} \quad (31)$$

Proof. The key proof idea is to resort to a high-order successive expansion of $\widehat{\mathbf{V}}^\pi - \mathbf{V}^\pi$, followed by fine-grained analysis of each term up to a certain logarithmic order. See Appendix B.1. \square

Clearly, Theorem 3 is a straightforward consequence of Lemma 1. Further, we strengthen the result by providing an additional instance-dependent bound (see the first line of (31) that depends on the true instance $\mathbf{P}_\pi, \mathbf{V}^\pi$), which is often tighter than the worst-case bound stated in the second line of (31). Our contribution can be better understood when compared with Pananjady and Wainwright (2019). Assuming that there is no noise in the rewards, our instance-dependent guarantee matches Pananjady and Wainwright (2019, Theorem 1(a)) up to some $\log \log \frac{1}{1-\gamma}$ factor, while being capable of covering the full sample size range $N \geq \widetilde{\Omega}(\frac{1}{1-\gamma})$. In contrast, Pananjady and Wainwright (2019, Theorem 1) is only valid when $N \geq \widetilde{\Omega}(\frac{1}{(1-\gamma)^2})$.

Proof ideas. We now briefly and informally describe the key proof ideas. As a starting point, the elementary identities (30) allow us to obtain

$$\begin{aligned} \widehat{\mathbf{V}}^\pi - \mathbf{V}^\pi &= (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \mathbf{r}_\pi - \mathbf{V}^\pi \\ &= (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} (\mathbf{I} - \gamma \mathbf{P}_\pi) \mathbf{V}^\pi - (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi) \mathbf{V}^\pi \\ &= \gamma (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^\pi. \end{aligned} \quad (32)$$

Due to the complicated dependency between $(\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1}$ and $(\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^\pi$, a natural strategy is to control these two terms separately and then to combine bounds; see Agarwal et al. (2020, Lemma 5) for an introduction. This simple approach, however, leads to sub-optimal statistical guarantees.

In order to refine the statistical analysis, we propose to further expand (32) in a similar way to deduce

$$\begin{aligned} (32) &= \gamma (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^\pi + \gamma \left\{ (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} - (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \right\} (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^\pi \\ &= \gamma (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^\pi + \gamma^2 (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^\pi, \end{aligned} \quad (33)$$

where the last line holds due to the same reason as (32) (basically it can be seen by replacing \mathbf{r}_π with $(\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^\pi$ in (32)). This can be viewed as a “second-order” expansion, with (32) being a “first-order” counterpart. The advantage is that: the first term in (33) becomes easier to cope with than its counterpart (32), owing to the independence between $(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1}$ and $(\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^\pi$. However, the second term in (33) remains difficult to control optimally. To remedy this issue, we shall continue to expand it to higher order (up to some logarithmic order), which eventually allows for optimal control of the estimation error.

Another crucial issue is that: in order to obtain fine-grained analyses on each term in the expansion (except for the first-order term), a common approach is to combine the Bernstein inequality with a classical entrywise bound on a quantity taking the form $(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}(\mathbf{V})}$ (which dates back to Azar et al. (2013)). Such a classical bound in prior literature, however, is not sufficiently tight for our purpose, which calls for refinement; see Lemma 11. Details are deferred to Appendix B.1.

6.3 Analysis: perturbed model-based planning

This subsection moves on to establishing our theory for model-based planning (cf. Theorem 1) and outlines the key ideas. In what follows, we shall start by analyzing the unperturbed version, which will elucidate the role of reward perturbation in our analysis.

We first make note of the following elementary decomposition:

$$\begin{aligned} \mathbf{V}^* - \mathbf{V}^{\widehat{\pi}^*} &= (\widehat{\mathbf{V}}^{\widehat{\pi}^*} - \mathbf{V}^{\widehat{\pi}^*}) + (\widehat{\mathbf{V}}^{\pi^*} - \widehat{\mathbf{V}}^{\widehat{\pi}^*}) + (\mathbf{V}^* - \widehat{\mathbf{V}}^{\pi^*}) \\ &\leq (\widehat{\mathbf{V}}^{\widehat{\pi}^*} - \mathbf{V}^{\widehat{\pi}^*}) + (\mathbf{V}^{\pi^*} - \widehat{\mathbf{V}}^{\pi^*}), \end{aligned} \quad (34)$$

where the inequality follows from the optimality of $\widehat{\pi}^*$ w.r.t. $\widehat{\mathbf{V}}$ (so that $\widehat{\mathbf{V}}^{\pi^*} \leq \widehat{\mathbf{V}}^{\widehat{\pi}^*}$) and the definition $\mathbf{V}^* = \mathbf{V}^{\pi^*}$. This leaves us with two terms to control.

Step 1: bounding $\|\mathbf{V}^{\pi^*} - \widehat{\mathbf{V}}^{\pi^*}\|_\infty$. Given that π^* is independent of the data, we can carry out this step using Lemma 1. Specifically, taking $\pi = \pi^*$ in Lemma 1 yields that, with probability at least $1 - \delta$,

$$\|\widehat{\mathbf{V}}^{\pi^*} - \mathbf{V}^{\pi^*}\|_\infty \leq 6\sqrt{\frac{2\log\left(\frac{4|\mathcal{S}|\log\frac{e}{1-\gamma}}{\delta}\right)}{N(1-\gamma)^3}}. \quad (35)$$

Step 2: bounding $\|\widehat{\mathbf{V}}^{\widehat{\pi}^*} - \mathbf{V}^{\widehat{\pi}^*}\|_\infty$. Extending the result in Step 1 to $\|\widehat{\mathbf{V}}^{\widehat{\pi}^*} - \mathbf{V}^{\widehat{\pi}^*}\|_\infty$ is considerably more challenging, primarily due to the complicated statistical dependency between $(\mathbf{V}^{\widehat{\pi}^*}, \widehat{\mathbf{V}}^{\widehat{\pi}^*})$ and the data matrix $\widehat{\mathbf{P}}$. The recent work Agarwal et al. (2020) developed a clever “leave-one-out” type argument by constructing some auxiliary state-absorbing MDPs to decouple the statistical dependency when $\varepsilon < 1/\sqrt{1-\gamma}$. However, their argument falls short of accommodating the full range of ε . To address this challenge, our analysis consists of the following two steps, both of which require new ideas beyond Agarwal et al. (2020).

- *Decoupling statistical dependency between $\widehat{\pi}^*$ and $\widehat{\mathbf{P}}$.* Instead of attempting to decouple the statistical dependency between $\widehat{\mathbf{V}}^{\widehat{\pi}^*}$ and $\widehat{\mathbf{P}}$ as in Agarwal et al. (2020), we focus on decoupling the statistical dependency between the policy $\widehat{\pi}^*$ and $\widehat{\mathbf{P}}$. If this can be achieved, then the proof strategy adopted in Step 1 for a fixed policy becomes applicable (see Section 6.3.1). A key ingredient of this step lies in the construction of a collection of auxiliary state-action-absorbing MDPs (motivated by Agarwal et al. (2020)), which allows us to get hold of $\|\mathbf{V}^{\widehat{\pi}^*} - \widehat{\mathbf{V}}^{\widehat{\pi}^*}\|_\infty$. See Section 6.3.2 for details, with a formal bound delivered in Lemma 5.
- *Tie-breaking via reward perturbation.* A shortcoming of the above-mentioned approach, however, is that it relies crucially on the separability of $\widehat{\pi}^*$ from other policies; in other words, the proof might fail if $\widehat{\pi}^*$ is non-unique or not sufficiently distinguishable from others. Consequently, it remains to ensure that the optimal policy $\widehat{\pi}^*$ stands out from all the rest for all MDPs of interest. As it turns out, this can be guaranteed with high probability by slightly perturbing the reward function so as to break the ties. See Section 6.3.3 for details.

In the sequel, we shall flesh out these key ideas.

6.3.1 Value function estimation for a policy obeying Bernstein-type conditions

Before discussing how to decouple statistical dependency, we record a useful result that plays an important role in the analysis. Specifically, Lemma 1 can be generalized beyond the family of fixed policies (namely, those independent of $\widehat{\mathbf{P}}$), as long as a certain Bernstein-type condition — to be formalized in (37) — is satisfied. To make it precise, we need to introduce a set of auxiliary vectors as follows

$$\begin{aligned} \mathbf{r}^{(0)} &:= \mathbf{r}_\pi, & \mathbf{V}^{(0)} &:= (\mathbf{I} - \gamma\mathbf{P}_\pi)^{-1}\mathbf{r}^{(0)}, \\ \mathbf{r}^{(l)} &:= \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(l-1)}]}, & \mathbf{V}^{(l)} &:= (\mathbf{I} - \gamma\mathbf{P}_\pi)^{-1}\mathbf{r}^{(l)}, \quad l \geq 1. \end{aligned} \quad (36)$$

Our generalization of Lemma 1 is as follows, which does *not* require statistical independence between the policy π and the data $\widehat{\mathbf{P}}$. Here, we remind the reader of the notation $|\mathbf{z}| := [z_1, \dots, z_n]^\top$ and $\sqrt{\mathbf{z}} := [\sqrt{z_1}, \dots, \sqrt{z_n}]^\top$ for any vector $\mathbf{z} \in \mathbb{R}^n$.

Lemma 2. *Suppose that there exists some quantity $\beta_1 > 0$ such that $\{\mathbf{V}^{(l)}\}$ (cf. (36)) obeys*

$$\left| (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi)\mathbf{V}^{(l)} \right| \leq \sqrt{\frac{\beta_1}{N}} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(l)}]} + \frac{\beta_1 \|\mathbf{V}^{(l)}\|_\infty}{N} \mathbf{1}, \quad \text{for all } 0 \leq l \leq \log\left(\frac{e}{1-\gamma}\right). \quad (37)$$

Suppose that $N > \frac{16e^2}{1-\gamma}\beta_1$. Then the vectors $\mathbf{V}^\pi = (\mathbf{I} - \gamma\mathbf{P}_\pi)^{-1}\mathbf{r}_\pi$ and $\widehat{\mathbf{V}}^\pi = (\mathbf{I} - \gamma\widehat{\mathbf{P}}_\pi)^{-1}\mathbf{r}_\pi$ satisfy

$$\|\widehat{\mathbf{V}}^\pi - \mathbf{V}^\pi\|_\infty \leq \frac{6}{1-\gamma} \sqrt{\frac{\beta_1}{N(1-\gamma)}}. \quad (38)$$

While the Bernstein-type condition (37) clearly holds for some reasonably small β_1 if π is independent of $\widehat{\mathbf{P}}$, it might remain valid if π exhibits fairly “weak” statistical dependency on the data samples. This is a key step that paves the way for our subsequent analysis of $\widehat{\pi}^*$.

6.3.2 Decoupling statistical dependency via (s, a) -absorbing MDPs

We are now positioned to demonstrate how to control $\|\widehat{\mathbf{V}}^{\widehat{\pi}^*} - \mathbf{V}^{\widehat{\pi}^*}\|_\infty$ w.r.t. the optimal policy $\widehat{\pi}^*$ to $\widehat{\mathbf{V}}$. A crucial technical challenge lies in how to decouple the complicated statistical dependency between the optimal policy $\widehat{\pi}^*$ and the $\widehat{\mathbf{V}}^*$ (which heavily relies on the data samples). Towards this, we resort to a leave-one-row-out argument built upon a collection of auxiliary MDPs, largely motivated by the novel construction in (Agarwal et al., 2020, Section 4.2). In comparison to Agarwal et al. (2020) that introduces state-absorbing MDPs (so that a state s is absorbing regardless of the subsequent actions chosen), our construction is a set of state-action-absorbing MDPs, in which a state s is absorbing only when a designated action a is always executed at the state s .

Construction of (s, a) -absorbing MDPs. For each state-action pair (s, a) and each scalar u with $|u| \leq 1/(1-\gamma)$, we construct an auxiliary MDP $\mathcal{M}_{s,a,u}$ — it is identical to the original \mathcal{M} except that it is absorbing in state s if we always choose action a in state s . More specifically, the probability transition kernel associated with $\mathcal{M}_{s,a,u}$ (denoted by $P_{\mathcal{M}_{s,a,u}}$) can be specified by

$$\begin{aligned} P_{\mathcal{M}_{s,a,u}}(s | s, a) &= 1, \\ P_{\mathcal{M}_{s,a,u}}(s' | s, a) &= 0, & \text{for all } s' \neq s, \\ P_{\mathcal{M}_{s,a,u}}(\cdot | s', a') &= P_{\mathcal{M}}(\cdot | s', a'), & \text{for all } (s', a') \neq (s, a), \end{aligned} \quad (39)$$

where $P_{\mathcal{M}}$ is the probability transition kernel w.r.t. the original \mathcal{M} . Meanwhile, the instant reward received at (s, a) in $\mathcal{M}_{s,a,u}$ is set to be u , while the rewards at all other state-action pairs stay unchanged. We can define $\widehat{\mathcal{M}}_{s,a,u}$ analogously (so that its probability transition matrix is identical to $\widehat{\mathbf{P}}$ except that the (s, a) -th row becomes absorbing). The main advantage of this construction is that: for any fixed u , the MDP $\widehat{\mathcal{M}}_{s,a,u}$ is statistically independent of $\widehat{\mathbf{P}}_{s,a}$ (the row of $\widehat{\mathbf{P}}$ corresponding to the state-action pair (s, a) , determined by the samples collected for the (s, a) pair).

To streamline notation, we let $\mathbf{Q}_{s,a,u}^\pi$ represent the Q-function of $\mathcal{M}_{s,a,u}$ under a policy π , denote by $\pi_{s,a,u}^*$ the optimal policy associated with $\mathcal{M}_{s,a,u}$, and let $\mathbf{Q}_{s,a,u}^*$ be the Q-function under this optimal policy $\pi_{s,a,u}^*$. The notation $\mathbf{V}_{s,a,u}^\pi$ and $\mathbf{V}_{s,a,u}^*$ regarding value functions, as well as their counterparts (i.e. $\widehat{\mathbf{Q}}_{s,a,u}^\pi$, $\widehat{\mathbf{Q}}_{s,a,u}^*$, $\widehat{\mathbf{V}}_{s,a,u}^\pi$, $\widehat{\mathbf{V}}_{s,a,u}^*$, $\widehat{\pi}_{s,a,u}^*$) in the empirical MDP $\widehat{\mathcal{M}}$, can be defined in an analogous fashion.

Remark 4. The careful reader will remark that the instant reward u is constrained to reside within $[-\frac{1}{1-\gamma}, \frac{1}{1-\gamma}]$ rather than the usual range $[0, 1]$. Fortunately, none of the subsequent steps that involve u requires u to lie within $[0, 1]$.

Intimate connections between the auxiliary MDPs and the original MDP. In the following, we introduce a result that connects the Q-function and the value function of the absorbing MDP with those of the original MDP. The idea is motivated by Agarwal et al. (2020, Lemma 7) and its proof is deferred to Appendix B.2.

Lemma 3. *Setting $u^* := r(s, a) + \gamma(\mathbf{P}\mathbf{V}^*)_{s,a} - \gamma V^*(s)$, one has*

$$\mathbf{Q}_{s,a,u^*}^* = \mathbf{Q}^* \quad \text{and} \quad \mathbf{V}_{s,a,u^*}^* = \mathbf{V}^*. \quad (40)$$

Remark 5. Lemma 3 does not rely on the particular form of \mathbf{P} , and can be directly generalized to the empirical model $\widehat{\mathbf{P}}$ and the auxiliary MDPs built upon $\widehat{\mathbf{P}}$.

In words, by properly setting the instant reward $u = u^*$ (which can be easily shown to reside within $[-\frac{1}{1-\gamma}, \frac{1}{1-\gamma}]$), one guarantees that the (s, a) -absorbing MDP and the original MDP have the same Q-function and value function under the respective optimal policies.

Representing $\widehat{\pi}^*$ via a small set of policies independent of $\widehat{\mathbf{P}}_{s,a}$. With Lemma 3 in place, it is tempting to use $\widehat{\mathcal{M}}_{s,a,\widehat{u}^*}$ with $\widehat{u}^* := r(s, a) + \gamma(\widehat{\mathbf{P}}\widehat{\mathbf{V}}^*)_{s,a} - \gamma\widehat{V}^*(s)$ to replace the original $\widehat{\mathcal{M}}$. The rationale is simple: given that the probability transition matrix of $\widehat{\mathcal{M}}_{s,a,\widehat{u}^*}$ does not rely upon $\widehat{\mathbf{P}}_{s,a}$, the statistical dependency between $\widehat{\mathcal{M}}_{s,a,\widehat{u}^*}$ and $\widehat{\mathbf{P}}_{s,a}$ is now fully embedded into a single parameter \widehat{u}^* . This motivates us

to decouple the statistical dependency effectively by constructing an epsilon-net (see, e.g., [Vershynin \(2018\)](#)) w.r.t. this single parameter. The aim is to locate a point u_0 over a small fixed set such that (i) it is close to \hat{u}^* , and (ii) its associated optimal policy is identical to the original optimal policy $\hat{\pi}^*$.

It turns out that this aim can be accomplished as long as the original Q-function \hat{Q}^* satisfies a sort of separation condition (which indicates that there is no tie when it comes to the optimal policy). To make it precise, given any $0 < \omega < 1$, our separation condition is characterized through the following event

$$\mathcal{B}_\omega := \left\{ \hat{Q}^*(s, \hat{\pi}^*(s)) - \max_{a: a \neq \hat{\pi}^*(s)} \hat{Q}^*(s, a) \geq \omega \text{ for all } s \in \mathcal{S} \right\}. \quad (41)$$

Clearly, on the event \mathcal{B}_ω , the optimal policy $\hat{\pi}^*$ is unique, since for each s the action $\hat{\pi}^*(s)$ results in a strictly higher Q-value compared to any other action. With this separation condition in mind, our result is stated below. Here and throughout, we define an epsilon-net of the interval $[-\frac{1}{1-\gamma}, \frac{1}{1-\gamma}]$ as follows

$$\mathcal{N}_\epsilon := \left\{ -n_\epsilon \epsilon, \dots, -\epsilon, 0, \epsilon, \dots, n_\epsilon \epsilon \right\}, \quad \text{for the largest integer } n_\epsilon \text{ obeying } n_\epsilon \epsilon < \frac{1}{1-\gamma}, \quad (42)$$

which has cardinality at most $\frac{2}{(1-\gamma)\epsilon}$.

Lemma 4. *Consider any $\omega > 0$, and suppose the event \mathcal{B}_ω (cf. (41)) holds. Then for any pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, there exists a point $u_0 \in \mathcal{N}_{(1-\gamma)\omega/4}$, such that*

$$\hat{\pi}^* = \hat{\pi}_{s,a,u_0}^*. \quad (43)$$

Proof. See Appendix B.3. □

Deriving an optimal error bound under the separation condition. Armed with the above bounds, we are ready to derive the desired error bound by combining Lemma 2 and Lemma 4.

Lemma 5. *Given $0 < \omega < 1$ and $\delta > 0$, suppose that \mathcal{B}_ω (defined in (41)) occurs with probability at least $1 - \delta$. Then with probability at least $1 - 3\delta$,*

$$\|\widehat{\mathbf{V}}^{\hat{\pi}^*} - \mathbf{V}^{\hat{\pi}^*}\|_\infty \leq 6 \sqrt{\frac{2 \log \left(\frac{32|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3 \omega \delta} \right)}{N(1-\gamma)^3}} \quad \text{and} \quad \mathbf{V}^* - \mathbf{V}^{\hat{\pi}^*} \leq 12 \sqrt{\frac{2 \log \left(\frac{32|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^3 \omega \delta} \right)}{N(1-\gamma)^3}} \mathbf{1}, \quad (44)$$

provided that $N \geq \frac{c_0 \log \left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)\delta\omega} \right)}{1-\gamma}$ for some sufficiently large constant $c_0 > 0$.

Proof. See Appendix B.4. □

6.3.3 A tie-breaking argument

Unfortunately, the separation condition specified in \mathcal{B}_ω (cf. (41)) does not always hold. In order to accommodate all possible MDPs of interest without imposing such a special separation condition, we put forward a perturbation argument allowing one to generate a new MDP that (i) satisfies the separation condition, and that (ii) is sufficiently close to the original MDP.

Specifically, let us represent the proposed reward perturbation (8) in a vector form as follows

$$\mathbf{r}_p := \mathbf{r} + \boldsymbol{\zeta}, \quad (45)$$

where $\boldsymbol{\zeta} = [\zeta(s, a)]_{(s,a) \in \mathcal{S} \times \mathcal{A}}$ is an $|\mathcal{S}||\mathcal{A}|$ -dimensional vector composed of independent entries with each $\zeta(s, a) \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, \xi)$. We aim to show that: by randomly perturbing the reward function, we can “break the tie” in the Q-function and ensure sufficient separation of Q-values associated with different actions.

To formalize our result, we find it convenient to introduce additional notation. Denote by π_p^* the optimal policy of the MDP $\mathcal{M}_p = (\mathcal{S}, \mathcal{A}, \mathbf{P}, \mathbf{r}_p, \gamma)$, and Q_p^* its optimal state-action value function. We can define \hat{Q}_p^* and $\hat{\pi}_p^*$ analogously for the MDP $\widehat{\mathcal{M}}_p = (\mathcal{S}, \mathcal{A}, \widehat{\mathbf{P}}, \mathbf{r}_p, \gamma)$. Our result is phrased as follows.

Lemma 6. Consider the perturbed reward vector defined in expression (45). With probability at least $1 - \delta$,

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A} \text{ with } a \neq \pi_p^*(s) : \quad Q_p^*(s, \pi_p^*(s)) - Q_p^*(s, a) > \frac{\xi \delta (1 - \gamma)}{3|\mathcal{S}||\mathcal{A}|^2}. \quad (46)$$

This result holds unchanged if (Q_p^*, π_p^*) is replaced by $(\widehat{Q}_p^*, \widehat{\pi}_p^*)$.

Proof. See Appendix B.5. □

Lemma 6 reveals that at least a polynomially small degree of separation ($\omega = \frac{\xi \delta (1 - \gamma)}{3|\mathcal{S}||\mathcal{A}|^2}$) arises upon random perturbation (with size ξ) of the reward function. As we shall see momentarily, this level of separation suffices for our purpose.

6.3.4 Proof of Theorem 1

Let us consider the randomly perturbed reward function as in (45). For any policy π , we denote by \mathbf{V}_p^π (resp. $\widehat{\mathbf{V}}_p^\pi$) the corresponding value function vector in the MDP with probability transition matrix \mathbf{P} (resp. $\widehat{\mathbf{P}}$) and reward vector \mathbf{r}_p . Note that π_p^* (resp. $\widehat{\pi}_p^*$) denotes the optimal policy that maximizes \mathbf{V}_p^π (resp. $\widehat{\mathbf{V}}_p^\pi$).

In view of Lemma 6, with probability at least $1 - \delta$ one has the separation

$$\left| \widehat{Q}_p^*(s, \widehat{\pi}_p^*(s)) - \widehat{Q}_p^*(s, a) \right| > \frac{\xi \delta (1 - \gamma)}{3|\mathcal{S}||\mathcal{A}|^2} \quad (47)$$

uniformly over all s and $a \neq \widehat{\pi}_p^*(s)$. With this separation in place, taking $\omega := \frac{\xi \delta (1 - \gamma)}{3|\mathcal{S}||\mathcal{A}|^2}$ in Lemma 5 yields

$$\left\| \mathbf{V}_p^{\pi_p^*} - \mathbf{V}_p^{\widehat{\pi}_p^*} \right\|_\infty \leq 12 \sqrt{\frac{2 \log \left(\frac{96|\mathcal{S}|^2|\mathcal{A}|^3}{(1-\gamma)^4 \xi \delta^2} \right)}{N(1-\gamma)^3}}. \quad (48)$$

In addition, the value functions under any policy π obeys

$$\mathbf{V}^\pi - \mathbf{V}_p^\pi = \mathbf{\Pi}^\pi \left((\mathbf{I} - \mathbf{P}^\pi)^{-1} \mathbf{r} - (\mathbf{I} - \mathbf{P}^\pi)^{-1} \mathbf{r}_p \right),$$

which taken collectively with the facts $\|\mathbf{r} - \mathbf{r}_p\|_\infty \leq \xi$ and $\|(\mathbf{I} - \gamma \mathbf{P}^\pi)^{-1}\|_1 \leq \frac{1}{1-\gamma}$ gives

$$\left\| \mathbf{V}^\pi - \mathbf{V}_p^\pi \right\|_\infty \leq \|(\mathbf{I} - \gamma \mathbf{P}^\pi)^{-1}\|_1 \|\mathbf{r} - \mathbf{r}_p\|_\infty \leq \frac{1}{1-\gamma} \xi.$$

Specializing the above relation to π^* and $\widehat{\pi}_p^*$ gives

$$\left\| \mathbf{V}^{\pi^*} - \mathbf{V}_p^{\pi^*} \right\|_\infty \leq \frac{1}{1-\gamma} \xi \quad \text{and} \quad \left\| \mathbf{V}^{\widehat{\pi}_p^*} - \mathbf{V}_p^{\widehat{\pi}_p^*} \right\|_\infty \leq \frac{1}{1-\gamma} \xi. \quad (49)$$

Now let us consider the following decomposition

$$\begin{aligned} \mathbf{V}^{\widehat{\pi}_p^*} - \mathbf{V}^* &= (\mathbf{V}_p^{\widehat{\pi}_p^*} - \mathbf{V}_p^{\widehat{\pi}_p^*}) + (\mathbf{V}_p^{\widehat{\pi}_p^*} - \mathbf{V}_p^{\pi_p^*}) + (\mathbf{V}_p^{\pi_p^*} - \mathbf{V}_p^{\pi^*}) + (\mathbf{V}_p^{\pi^*} - \mathbf{V}^*) \\ &\geq (\mathbf{V}_p^{\widehat{\pi}_p^*} - \mathbf{V}_p^{\widehat{\pi}_p^*}) + (\mathbf{V}_p^{\widehat{\pi}_p^*} - \mathbf{V}_p^{\pi_p^*}) + (\mathbf{V}_p^{\pi^*} - \mathbf{V}^*), \end{aligned}$$

where the last step follows from the optimality of π_p^* w.r.t. \mathbf{V}_p . Taking this collectively with the inequalities (48) and (49), one shows that with probability greater than $1 - 3\delta$,

$$\mathbf{V}_p^{\widehat{\pi}_p^*} - \mathbf{V}^* \geq - \left(\frac{2}{1-\gamma} \xi + 12 \sqrt{\frac{2 \log \left(\frac{96|\mathcal{S}|^2|\mathcal{A}|^3}{\xi(1-\gamma)^4 \delta^2} \right)}{N(1-\gamma)^3}} \right) \mathbf{1}$$

By taking $\xi = \frac{(1-\gamma)\varepsilon}{3|\mathcal{S}|^5|\mathcal{A}|^5}$ and $N \geq \frac{c_0 \log(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)\delta\varepsilon})}{(1-\gamma)^3\varepsilon^2}$ for some constant $c_0 > 0$ large enough, we can ensure that $\mathbf{0} \geq \mathbf{V}^{\hat{\pi}_p^*} - \mathbf{V}^* \geq -\varepsilon \mathbf{1}$ as claimed. Regarding the Q-functions, the Bellman equation gives

$$\mathbf{Q}^{\hat{\pi}_p^*} - \mathbf{Q}^* = \mathbf{r} + \gamma \mathbf{P} \mathbf{V}^{\hat{\pi}_p^*} - (\mathbf{r} + \gamma \mathbf{P} \mathbf{V}^*) = \gamma \mathbf{P} (\mathbf{V}^{\hat{\pi}_p^*} - \mathbf{V}^*).$$

Consequently, one has

$$\mathbf{Q}^{\hat{\pi}_p^*} - \mathbf{Q}^* \geq -(\gamma \|\mathbf{P}\|_1 \|\mathbf{V}^{\hat{\pi}_p^*} - \mathbf{V}^*\|_\infty) \mathbf{1} \geq -\gamma \varepsilon \mathbf{1}.$$

Finally, we demonstrate that both the empirical QVI and PI w.r.t. $\widehat{\mathcal{M}}_p$ are guaranteed to find $\hat{\pi}_p^*$ in a few iterations. Suppose for the moment that we can obtain a policy π_k obeying

$$\|\widehat{\mathbf{Q}}_p^{\pi_k} - \widehat{\mathbf{Q}}_p^*\|_\infty < \frac{\xi \delta (1-\gamma)}{8|\mathcal{S}||\mathcal{A}|^2}. \quad (50)$$

Then for any $s \in \mathcal{S}$ and any action $a \neq \hat{\pi}_p^*(s)$ one has

$$\begin{aligned} & \widehat{\mathbf{Q}}_p^{\pi_k}(s, \hat{\pi}_p^*(s)) - \widehat{\mathbf{Q}}_p^{\pi_k}(s, a) \\ &= \widehat{\mathbf{Q}}_p^*(s, \hat{\pi}_p^*(s)) - \widehat{\mathbf{Q}}_p^*(s, a) + \left(\widehat{\mathbf{Q}}_p^{\pi_k}(s, \hat{\pi}_p^*(s)) - \widehat{\mathbf{Q}}_p^*(s, \hat{\pi}_p^*(s)) \right) - \left(\widehat{\mathbf{Q}}_p^{\pi_k}(s, a) - \widehat{\mathbf{Q}}_p^*(s, a) \right) \\ &\geq \widehat{\mathbf{Q}}_p^*(s, \hat{\pi}_p^*(s)) - \widehat{\mathbf{Q}}_p^*(s, a) - 2\|\widehat{\mathbf{Q}}_p^{\pi_k} - \widehat{\mathbf{Q}}_p^*\|_\infty \\ &> \frac{\xi \delta (1-\gamma)}{4|\mathcal{S}||\mathcal{A}|^2} - 2 \cdot \frac{\xi \delta (1-\gamma)}{8|\mathcal{S}||\mathcal{A}|^2} = 0, \end{aligned}$$

where the last line results from (47) and (50). In other words, we can perfectly recover the policy $\hat{\pi}_p^*$ from the estimate $\widehat{\mathbf{Q}}_p^{\pi_k}$, provided that (50) is satisfied. In addition, it has been shown that (Azar et al., 2013, Lemma 2) the greedy policy induced by k -th iteration of both algorithms — denoted by π_k — satisfies $\|\widehat{\mathbf{Q}}_p^{\pi_k} - \widehat{\mathbf{Q}}_p^*\|_\infty \leq \frac{2\gamma^{k+1}}{(1-\gamma)^2}$. Taking $\xi = \frac{c_1(1-\gamma)\varepsilon}{|\mathcal{S}|^5|\mathcal{A}|^5}$ and $k = \frac{c_2}{1-\gamma} \log(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)\delta\varepsilon})$ for some constant $c_2 > 0$ large enough, one guarantees that π_k satisfies (50), which in turn ensures perfect recovery of $\hat{\pi}_p^*$.

6.4 Analysis: conservative model-based planning

In view of the conservative model-based planning (10), we begin with the following decomposition:

$$\mathbf{V}^* - \mathbf{V}^{\hat{\pi}_c} = (\mathbf{V}^{\pi^*} - \widehat{\mathbf{V}}^{\pi^*}) + (\widehat{\mathbf{V}}^{\pi^*} - \widehat{\mathbf{V}}^{\hat{\pi}_c}) + (\widehat{\mathbf{V}}^{\hat{\pi}_c} - \mathbf{V}^{\hat{\pi}_c}). \quad (51)$$

In order to control the second term on the right-hand side of the above identity, we resort to the following lemma, whose proof is postponed to Section B.6.

Lemma 7. *It holds that*

$$\widehat{\mathbf{V}}^{\pi^*} - \widehat{\mathbf{V}}^{\hat{\pi}_c} \leq \frac{\xi}{1-\gamma} \mathbf{1}. \quad (52)$$

Combining Lemma 7 with (51), we arrive at

$$\begin{aligned} \mathbf{V}^* - \mathbf{V}^{\hat{\pi}_c} &= (\mathbf{V}^{\pi^*} - \widehat{\mathbf{V}}^{\pi^*}) + (\widehat{\mathbf{V}}^{\pi^*} - \widehat{\mathbf{V}}^{\hat{\pi}_c}) + (\widehat{\mathbf{V}}^{\hat{\pi}_c} - \mathbf{V}^{\hat{\pi}_c}) \\ &\leq (\mathbf{V}^{\pi^*} - \widehat{\mathbf{V}}^{\pi^*}) + \frac{\xi}{1-\gamma} \mathbf{1} + (\widehat{\mathbf{V}}^{\hat{\pi}_c} - \mathbf{V}^{\hat{\pi}_c}). \end{aligned} \quad (53)$$

Clearly, the first term of (53) has already been controlled in Section 6.2, while the second term of (53) is extremely small when we take $\xi = O(\frac{(1-\gamma)\varepsilon}{|\mathcal{S}|^5|\mathcal{A}|^5})$. It thus suffices to bound the third term of (53), which again requires decoupling the statistical dependence between $\hat{\pi}_c$ and $\widehat{\mathbf{P}}$.

Representing $\hat{\pi}_c$ via a small set of policies independent of $\hat{P}_{s,a}$. Akin to our analysis for the perturbed model-based planning algorithm in Section 6.3, a key step lies in demonstrating the connection between $\hat{\pi}_c$ and a reasonably small collection of leave-one-out auxiliary MDPs. Towards this, we are in need of the following lemma, which characterizes certain “stability” of our conservative model-based strategy and lies at the core of our analysis. The proof is deferred to Section B.7.

Lemma 8. *Consider any given Q-function $\hat{Q} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ and its associated value function $\hat{V} : \mathcal{S} \rightarrow \mathbb{R}$ (i.e. $\hat{V}(s) = \max_a \hat{Q}(s, a)$ for all s). Generate an independent random variable $\varsigma \sim \text{Unif}(0, \xi)$. Then with probability at least $1 - \delta$,*

$$\forall s \in \mathcal{S}, \quad \{a \in \mathcal{A} : \hat{Q}(s, a) > \hat{V}(s) - \varsigma\} = \{a \in \mathcal{A} : Q(s, a) > V(s) - \varsigma\} \quad (54)$$

holds simultaneously for all Q-function $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ (and its associated value function $V : \mathcal{S} \rightarrow \mathbb{R}$) obeying

$$\max_{s,a} |Q(s, a) - \hat{Q}(s, a)| \leq \frac{\xi \delta}{8|\mathcal{S}||\mathcal{A}|}. \quad (55)$$

As an important implication of this lemma, the policy $\hat{\pi}_c$ computed in (10) remains unchanged upon slight perturbation of the Q-function estimates. Armed with Lemma 8 and the leave-one-out auxiliary MDPs $\{\widehat{\mathcal{M}}_{s,a,u}\}$ constructed in Section 6.3, our analysis proceeds as follows.

- For each $(s, a) \in \mathcal{S} \times \mathcal{A}$, there exists a point $u_0 \in \mathcal{N}_{(1-\gamma)\omega}$ (cf. (42)) such that the optimal Q-function of $\widehat{\mathcal{Q}}_{s,a,u_0}^*$ of $\widehat{\mathcal{M}}_{s,a,u_0}$ obeys

$$\|\widehat{\mathcal{Q}}_{s,a,u_0}^* - \widehat{Q}^*\|_\infty \leq \omega. \quad (56)$$

This is a fact that has already been established in the proof of Lemma 4; see (89).

- Define a conservative policy for $\widehat{\mathcal{M}}_{s,a,u_0}$ as follows:

$$\forall s' \in \mathcal{S} : \quad \hat{\pi}_{s,a,u_0,c}(s') = \min \{a' \in \mathcal{A} : \widehat{\mathcal{Q}}_{s,a,u_0}^*(s', a') > \widehat{V}_{s,a,u_0}^*(s') - \varsigma\},$$

where $\widehat{\mathcal{Q}}_{s,a,u_0}^*$ and \widehat{V}_{s,a,u_0}^* denote the optimal Q-function and optimal value function of $\widehat{\mathcal{M}}_{s,a,u_0}$, respectively. Taking $\omega = \frac{\xi \delta}{8|\mathcal{S}||\mathcal{A}|}$ and invoking Lemma 8 and (56), we arrive at

$$\hat{\pi}_c = \hat{\pi}_{s,a,u_0,c}. \quad (57)$$

This result (57) parallels Lemma 4 for the perturbed model-based planning algorithm, revealing that $\hat{\pi}_c$ is representable using a policy independent of the randomness associated with (s, a) . The remaining proof of Theorem 2 then follows from an identical argument as in the proof of Theorem 1, and is hence omitted here.

7 Analysis: finite-horizon MDPs

In this section, we outline the proof of Theorem 4. We shall start by introducing a set of convenient matrix notation before embarking on the main proof.

7.1 Matrix notation and Bellman equations

Akin to the infinite-horizon case, we introduce some matrix notation for finite-horizon MDPs. Analogous to Section 6.1, we introduce the following set of notation.

- $\mathbf{r}_h \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$: a vector representing the reward function r_h at step h .
- $\mathbf{V}_h^\pi \in \mathbb{R}^{|\mathcal{S}|}$: a vector representing the value function V_h^π of π at step h .
- $\mathbf{V}_h^* \in \mathbb{R}^{|\mathcal{S}|}$: a vector representing the optimal value function V_h^* at step h .
- $\mathbf{Q}_h^\pi \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$: a vector representing the Q-function Q_h^π of π at step h .

- $\mathbf{Q}_h^* \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$: a vector representing the optimal Q-function Q_h^* at step h .
- $\mathbf{P}_h \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times |\mathcal{S}|}$: a matrix representing the probability transition kernel P_h at step h .
- $\mathbf{P}_{h,\pi} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$: a submatrix of \mathbf{P}_h , which consists of the rows with indices coming from $\{(s, \pi_h(s)) \mid s \in \mathcal{S}\}$.
- $\mathbf{r}_h^\pi \in \mathbb{R}^{|\mathcal{S}|}$: a subvector of \mathbf{r}_h , which consists of the rows with indices coming from $\{(s, \pi_h(s)) \mid s \in \mathcal{S}\}$.

Armed with the above notation, the Bellman equation here is given by

$$\mathbf{Q}_h^\pi = \mathbf{r}_h + \mathbf{P}_h \mathbf{V}_{h+1}^\pi, \quad 1 \leq h \leq H, \quad (58)$$

where we recall that for all $s \in \mathcal{S}$

$$V_h^\pi(s) = Q_h^\pi(s, \pi_h(s)) \quad \text{and} \quad V_{H+1}^\pi(s) = 0. \quad (59)$$

This also allows one to derive

$$\mathbf{V}_h^\pi = \mathbf{r}_h^\pi + \mathbf{P}_{h,\pi} \mathbf{V}_{h+1}^\pi. \quad (60)$$

We shall also define $\widehat{\mathbf{V}}_h^\pi, \widehat{\mathbf{Q}}_h^\pi, \widehat{\mathbf{V}}_h^*, \widehat{\mathbf{Q}}_h^*, \widehat{\mathbf{P}}_h, \widehat{\mathbf{P}}_{h,\pi}$ w.r.t. the empirical MDP $\widehat{\mathcal{M}}$ in an analogous fashion.

7.2 An auxiliary value function sequence obeying Bernstein-type conditions

Similar to the infinite-horizon case (in particular, Section 6.3.1), we find it convenient to introduce a collection of auxiliary vectors as follows. For any $l \geq 0$, define

$$\mathbf{V}_{H+1}^{(l)} := \mathbf{0} \quad \text{and} \quad \widehat{\mathbf{V}}_{H+1}^{(l)} := \mathbf{0}; \quad (61)$$

and for any $1 \leq h \leq H$ and any policy π , define the following sequences recursively:

$$\begin{aligned} \mathbf{r}_h^{(0)} &:= \mathbf{r}_h^\pi, & \mathbf{V}_h^{(0)} &:= \mathbf{r}_h^{(0)} + \mathbf{P}_{h,\pi} \mathbf{V}_{h+1}^{(0)}, & \widehat{\mathbf{V}}_h^{(0)} &:= \mathbf{r}_h^{(0)} + \widehat{\mathbf{P}}_{h,\pi} \widehat{\mathbf{V}}_{h+1}^{(0)}, \\ \mathbf{r}_h^{(l)} &:= \sqrt{\text{Var}_{\mathbf{P}_{h,\pi}}[\mathbf{V}_{h+1}^{(l-1)}]}, & \mathbf{V}_h^{(l)} &:= \mathbf{r}_h^{(l)} + \mathbf{P}_{h,\pi} \mathbf{V}_{h+1}^{(l)}, & \widehat{\mathbf{V}}_h^{(l)} &:= \mathbf{r}_h^{(l)} + \widehat{\mathbf{P}}_{h,\pi} \widehat{\mathbf{V}}_{h+1}^{(l)}, \end{aligned} \quad l \geq 1. \quad (62)$$

As can be easily verified, $\{\mathbf{V}_h^{(0)}\}$ coincides with the value function of policy π in the true MDP \mathcal{M} , while $\{\widehat{\mathbf{V}}_h^{(0)}\}$ corresponds to the value function of policy π in the empirical MDP $\widehat{\mathcal{M}}$.

As it turns out, if the above auxiliary sequence satisfies certain Bernstein-type conditions, then we can establish a useful upper bound on the entrywise difference between $\mathbf{V}_h^{(0)}$ and $\widehat{\mathbf{V}}_h^{(0)}$, as stated below. The proof of this lemma is deferred to Section C.1.

Lemma 9. *Suppose that there exists some quantity $\beta_1 > 0$ such that the sequence $\{\mathbf{V}_{h+1}^{(l)}\}$ constructed in (62) obeys*

$$\left| (\widehat{\mathbf{P}}_{h,\pi} - \mathbf{P}_{h,\pi}) \mathbf{V}_{h+1}^{(l)} \right| \leq \sqrt{\frac{\beta_1}{N}} \sqrt{\text{Var}_{\mathbf{P}_{h,\pi}}[\mathbf{V}_{h+1}^{(l)}]} + \frac{\beta_1 \|\mathbf{V}_{h+1}^{(l)}\|_\infty}{N} \mathbf{1}, \quad \forall 0 \leq l \leq \log_2 H, 1 \leq h \leq H-1. \quad (63)$$

In addition, assume that $N > 12H\beta_1$. Then we have

$$\|\widehat{\mathbf{V}}_h^{(0)} - \mathbf{V}_h^{(0)}\|_\infty \leq 6H \sqrt{\frac{3\beta_1 H}{N}} \quad (64)$$

for all $1 \leq h \leq H$.

7.3 Proof of Theorem 4

Let us begin with the following elementary decomposition:

$$\begin{aligned} \mathbf{V}_h^* - \mathbf{V}_h^{\hat{\pi}^*} &= (\widehat{\mathbf{V}}_h^{\hat{\pi}^*} - \mathbf{V}_h^{\hat{\pi}^*}) + (\widehat{\mathbf{V}}_h^{\pi^*} - \widehat{\mathbf{V}}_h^{\hat{\pi}^*}) + (\mathbf{V}_h^* - \widehat{\mathbf{V}}_h^{\pi^*}) \\ &\leq (\widehat{\mathbf{V}}_h^{\hat{\pi}^*} - \mathbf{V}_h^{\hat{\pi}^*}) + (\mathbf{V}_h^{\pi^*} - \widehat{\mathbf{V}}_h^{\pi^*}). \end{aligned} \quad (65)$$

Here, the inequality follows from the definition $\mathbf{V}_h^* = \mathbf{V}_h^{\pi^*}$, as well as the fact that $\widehat{\mathbf{V}}_h^{\pi^*} \leq \widehat{\mathbf{V}}_h^{\hat{\pi}^*}$ (since $\hat{\pi}^*$ is the optimal policy of the empirical MDP). In light of (65), there are two terms that need to be controlled.

We intend to bound both $\widehat{\mathbf{V}}_h^{\hat{\pi}^*} - \mathbf{V}_h^{\hat{\pi}^*}$ and $\mathbf{V}_h^{\pi^*} - \widehat{\mathbf{V}}_h^{\pi^*}$ by means of Lemma 9. Towards this, we first note that for any policy π , the associated value functions of the MDP and the empirical MDP obey the following Bellman equations:

$$\mathbf{V}_h^\pi := \mathbf{r}_h^\pi + \mathbf{P}_{h,\pi} \mathbf{V}_{h+1}^\pi, \quad \widehat{\mathbf{V}}_h^\pi := \mathbf{r}_h^\pi + \widehat{\mathbf{P}}_{h,\pi} \widehat{\mathbf{V}}_{h+1}^\pi,$$

along with the boundary conditions $\mathbf{V}_{H+1}^\pi = \widehat{\mathbf{V}}_{H+1}^\pi = \mathbf{0}$. This indicates that the vector $\mathbf{V}_h^{(0)}$ (resp. $\widehat{\mathbf{V}}_h^{(0)}$) constructed in (62) is precisely the value function of policy π at step h in the true MDP (resp. empirical MDP). As a result, in order to invoke Lemma 9, it is sufficient to verify the Bernstein-type condition (63) w.r.t. policies π^* and $\hat{\pi}^*$ for some sufficiently small quantity β_1 .

- Let us begin with the optimal policy π^* , which is fixed and statistically independent of the data samples. As a result, if we take $\pi = \pi^*$ during the construction (62), then it is clearly seen that $\widehat{\mathbf{P}}_h$ is statistically independent of $\mathbf{V}_{h+1}^{(l)}$. Applying the Bernstein inequality together with the union bound then guarantees that: with probability exceeding $1 - \delta$,

$$\left| (\widehat{\mathbf{P}}_h - \mathbf{P}_h) \mathbf{V}_{h+1}^{(l)} \right| \leq \sqrt{\frac{\beta_1}{N}} \sqrt{\text{Var}_{\mathbf{P}_h} [\mathbf{V}_{h+1}^{(l)}]} + \frac{\beta_1 \|\mathbf{V}_{h+1}^{(l)}\|_\infty}{N} \mathbf{1} \quad (66)$$

holds uniformly over all $0 \leq l \leq \log_2 H$, $1 \leq h \leq H - 1$, $(s, a) \in \mathcal{S} \times \mathcal{A}$, where β_1 is given by

$$\beta_1 := 4 \log \left(\frac{H|\mathcal{S}||\mathcal{A}|}{\delta} \right). \quad (67)$$

Armed with Condition (66), we can readily invoke Lemma 9 to reach

$$\|\widehat{\mathbf{V}}_h^{\pi^*} - \mathbf{V}_h^{\pi^*}\|_\infty = \|\widehat{\mathbf{V}}_h^{(0)} - \mathbf{V}_h^{(0)}\|_\infty \leq 6H \sqrt{\frac{3\beta_1 H}{N}} \leq 6H \sqrt{\frac{12H \log \left(\frac{H|\mathcal{S}||\mathcal{A}|}{\delta} \right)}{N}}$$

with probability at least $1 - \delta$.

- Next, we move on to the policy $\hat{\pi}^*$ by taking $\pi = \hat{\pi}^*$ during the construction (62). Note that $\mathbf{V}_{h+1}^{(l)}$ depends only on $\hat{\pi}_j^*$ ($j \geq h + 1$). In view of our assumption on $\hat{\pi}^*$ (i.e., it is computed backward via dynamic programming), $\hat{\pi}_i^*$ is independent of any $\widehat{\mathbf{P}}_j$ with $j < i$, and hence $\mathbf{V}_{h+1}^{(l)}$ is statistically independent from $\widehat{\mathbf{P}}_h$. Consequently, the preceding bounds (66) and (67) continue to hold. All of this immediately results in

$$\|\widehat{\mathbf{V}}_h^{\hat{\pi}^*} - \mathbf{V}_h^{\hat{\pi}^*}\|_\infty \leq 6H \sqrt{\frac{12H \log \left(\frac{H|\mathcal{S}||\mathcal{A}|}{\delta} \right)}{N}}$$

with probability exceeding $1 - \delta$.

Substituting the above bounds into (65), we arrive at

$$\mathbf{0} \leq \mathbf{V}_h^* - \mathbf{V}_h^{\hat{\pi}^*} \leq 12H \sqrt{\frac{12H \log \left(\frac{H|\mathcal{S}||\mathcal{A}|}{\delta} \right)}{N}} \mathbf{1} \quad (68)$$

with probability greater than $1 - 2\delta$, provided that $N \geq 48H \log \left(\frac{H|\mathcal{S}||\mathcal{A}|}{\delta} \right)$. By taking the right-hand side of (68) to be smaller than $\varepsilon \mathbf{1}$, we immediately conclude the proof.

8 Discussion

This paper has demonstrated that (some variants of) model-based planning algorithms achieve the minimax sample complexity in the presence of a generative model, as soon as the sample size exceeds the order of $\frac{|\mathcal{S}||\mathcal{A}|}{1-\gamma}$ for γ -discounted infinite horizon MDPs and $|\mathcal{S}||\mathcal{A}|H^2$ for time-inhomogeneous finite-horizon MDPs (modulo some log factor). Compared to prior literature, our result has considerably broadened the sample size range, allowing us to pin down a complete trade-off curve between sample complexity and statistical accuracy.

The present work opens up several directions for future investigation, which we discuss in passing below.

- *Is perturbation or conservative action selection necessary for infinite-horizon MDPs?* The planning algorithm analyzed here for infinite-horizon MDPs is either applied to a perturbed variant of the empirical MDP (as in perturbed model-based planning) or run in a conservative manner (as in conservative model-based planning). This, however, gives rise to a natural question regarding the necessity of perturbation or conservative action selection: can we achieve optimal performance directly using *plain* model-based planning on the empirical MDP? While we conjecture that the answer is affirmative, settling this conjecture requires new techniques beyond the analysis framework of this paper.
- *Improved analysis for model-free algorithms.* As mentioned previously, a even more severe sample complexity barrier is present in all prior theory regarding model-free approaches (e.g. Li et al. (2021d); Sidford et al. (2018a); Wainwright (2019b)). Our analysis might shed light on how to overcome such barriers for model-free approaches.
- *Time-homogeneous finite-horizon MDPs.* When it comes to finite-horizon MDPs, the present work concentrates on time-inhomogeneous MDPs where the probability transition kernels may vary across time steps. Another important scenario is concerned with time-homogeneous MDPs, where $P_1 = P_2 = \dots = P_H$. It remains unclear how to develop tight sample analysis for time-homogeneous MDPs due to the lack of statistical independence across time steps (namely, we shall use all samples to estimate the kernels across time steps as they are identical).
- *Beyond the tabular setting.* The current paper focuses on the tabular setting with finite state and action spaces. While we improve the sample size range, the sample complexities might still be prohibitively large when $|\mathcal{S}|$ and $|\mathcal{A}|$ are enormous. Therefore, it is desirable to further investigate settings where low-complexity function approximation is employed to improve the efficiency (e.g. Wang et al. (2021); Yang and Wang (2019)).

Acknowledgements

Y. Wei is supported in part by the the NSF grants CCF-2106778, DMS-2147546/2015447 and CAREER award DMS-2143215. Y. Chi is supported in part by the ONR grants N00014-18-1-2142 and N00014-19-1-2404, the NSF grants 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 grants FA9550-19-1-0030 and FA9550-22-1-0198, the ONR grant N00014-22-1-2354, and the NSF grants CCF-2221009, CCF-1907661, DMS-2014279, IIS-2218713 and IIS-2218773. We thank Qiwen Cui for pointing out an issue in Section B.5 in an early version of this paper, and thank Shicong Cen, Chen Cheng and Cong Ma for numerous discussions. Part of this work was done while G. Li, Y. Wei and Y. Chen were visiting the Simons Institute for the Theory of Computing.

A Preliminary facts

We begin by recording a few elementary facts about \mathbf{P}^π and \mathbf{P}_π (see definitions in (24)). These are standard results and we omit the proofs for brevity.

Lemma 10. *For any policy π , any probability transition matrix $\mathbf{P} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times |\mathcal{S}|}$ and any $0 < \gamma < 1$, one has*

- (a) $(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} = \sum_{i=0}^{\infty} (\gamma \mathbf{P}_\pi)^i$;
- (b) All entries of the matrix $(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1}$ are non-negative;
- (c) $\|(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1}\|_1 \leq 1/(1 - \gamma)$;
- (d) $(1 - \gamma)(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{1} = \mathbf{1}$;
- (e) For any non-negative vectors $\mathbf{0} \leq \mathbf{r}_1 \leq \mathbf{r}_2$ of compatible dimension, one has $\mathbf{0} \leq (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}_1 \leq (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}_2$.

The above results continue to hold if \mathbf{P}_π is replaced by \mathbf{P}^π .

B Proofs of auxiliary lemmas: infinite-horizon MDPs

B.1 Proofs of Lemma 1 and Lemma 2

Auxiliary notation and preliminaries. Before proceeding, we define several $|\mathcal{S}|$ -dimensional auxiliary vectors $\mathbf{r}^{(i)}$, $\mathbf{V}^{(i)}$, $\widehat{\mathbf{V}}^{(i)}$ ($1 \leq i \leq m$) recursively as follows

$$\begin{aligned}
\mathbf{r}^{(0)} &:= \mathbf{r}_\pi, & \mathbf{V}^{(0)} &:= (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}^{(0)}, & \widehat{\mathbf{V}}^{(0)} &:= (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \mathbf{r}^{(0)}, \\
\mathbf{r}^{(1)} &:= \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(0)}]}, & \mathbf{V}^{(1)} &:= (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}^{(1)}, & \widehat{\mathbf{V}}^{(2)} &:= (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \mathbf{r}^{(1)}, \\
&\vdots & & \vdots & & \vdots \\
\mathbf{r}^{(m)} &:= \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(m-1)}]}, & \mathbf{V}^{(m)} &:= (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}^{(m)}, & \widehat{\mathbf{V}}^{(m)} &:= (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \mathbf{r}^{(m)},
\end{aligned} \tag{69}$$

where m will be specified momentarily.

A crucial quantity that appears repeatedly in analyzing the above terms is $\|(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}(\mathbf{V})}\|_\infty$, whose importance was already made apparent in the work [Azar et al. \(2013\)](#). A widely used upper bound on this quantity, originally due to ([Azar et al., 2013](#), Lemma 8), is given by

$$\left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}(\mathbf{V})} \right\|_\infty \leq \frac{2 \log 2}{\gamma(1 - \gamma)^{1.5}} \|\mathbf{r}\|_\infty. \tag{70}$$

This bound turns out to be loose for our purpose, and we develop an improved bound as follows, whose proof is deferred to Appendix B.1.1.

Lemma 11. Consider any policy π and any probability transition matrix $\mathbf{P} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times |\mathcal{S}|}$. Let \mathbf{V} be a vector obeying $\mathbf{V} = (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}_\pi$ for some $|\mathcal{S}|$ -dimensional vector $\mathbf{r}_\pi \geq \mathbf{0}$. For any $0 < \gamma < 1$, one has

$$\left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}(\mathbf{V})} \right\|_\infty \leq \frac{4}{\gamma \sqrt{1 - \gamma}} \|\mathbf{V}\|_\infty. \tag{71}$$

Remark 6. In comparison to the bound (70) derived in ([Azar et al., 2013](#), Lemma 8), Lemma 11 offers an improved upper bound stated directly in terms of the properties of \mathbf{V} rather than those of \mathbf{r} .

As it turns out, Lemma 11 allows us to obtain an entrywise bound for $\mathbf{V}^{(l)}$ ($1 \leq l \leq m$). To begin with, the first term $\mathbf{V}^{(1)}$ satisfies

$$\|\mathbf{V}^{(1)}\|_\infty = \left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(0)}]} \right\|_\infty \tag{72}$$

since $\mathbf{r}^{(1)} = \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(0)}]}$. Next, for any $l > 1$ one has

$$\begin{aligned}
\|\mathbf{V}^{(l)}\|_\infty &= \|(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}^{(l)}\|_\infty = \left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(l-1)}]} \right\|_\infty \\
&\leq \frac{4}{\gamma \sqrt{1 - \gamma}} \|\mathbf{V}^{(l-1)}\|_\infty,
\end{aligned}$$

where the second identity results from the definition of $\mathbf{r}^{(l)}$, and the last inequality comes from Lemma 11. As a consequence, applying this inequality recursively gives

$$\|\mathbf{V}^{(l)}\|_\infty \leq \left(\frac{4}{\gamma \sqrt{1 - \gamma}} \right)^{l-1} \|\mathbf{V}^{(1)}\|_\infty. \tag{73}$$

Main proof. Equipped with the above facts, we are now in a position to prove the lemmas, for which we start with the more general one — Lemma 2. Consider any $0 \leq l \leq m$. We first observe that

$$\begin{aligned}\widehat{\mathbf{V}}^{(l)} - \mathbf{V}^{(l)} &= (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \mathbf{r}^{(l)} - \mathbf{V}^{(l)} \\ &= (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} (\mathbf{I} - \gamma \mathbf{P}_\pi) \mathbf{V}^{(l)} - (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi) \mathbf{V}^{(l)} \\ &= \gamma (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^{(l)},\end{aligned}\tag{74}$$

where the second line follows since, by definition, $(\mathbf{I} - \gamma \mathbf{P}_\pi) \mathbf{V}^{(l)} = \mathbf{r}^{(l)}$. Suppose that there exists some quantity $\beta_1 > 0$ such that the following condition

$$\left| (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^{(l)} \right| \leq \sqrt{\frac{\beta_1}{N}} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(l)}]} + \frac{\|\mathbf{V}^{(l)}\|_\infty \beta_1}{N} \mathbf{1}\tag{75}$$

holds uniformly for all $0 \leq l \leq m$. Then this combined with (74) gives

$$\begin{aligned}\|\widehat{\mathbf{V}}^{(l)} - \mathbf{V}^{(l)}\|_\infty &= \gamma \left\| (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^{(l)} \right\|_\infty \\ &\stackrel{(i)}{\leq} \gamma \left\| (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \left| (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^{(l)} \right| \right\|_\infty \\ &\stackrel{(ii)}{\leq} \gamma \sqrt{\frac{\beta_1}{N}} \left\| (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(l)}]} \right\|_\infty + \frac{\gamma \|\mathbf{V}^{(l)}\|_\infty \beta_1}{N} \left\| (\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1} \mathbf{1} \right\|_\infty.\end{aligned}$$

Here, (i) follows since $(\mathbf{I} - \gamma \widehat{\mathbf{P}}_\pi)^{-1}$ is a non-negative matrix, (ii) comes from (75) and the triangle inequality. Recalling the definition of $\mathbf{r}^{(l)}$ and $\widehat{\mathbf{V}}^{(l)}$ and invoking Lemma 10(d), we can further bound the above as

$$\begin{aligned}\|\widehat{\mathbf{V}}^{(l)} - \mathbf{V}^{(l)}\|_\infty &\leq \gamma \sqrt{\frac{\beta_1}{N}} \|\widehat{\mathbf{V}}^{(l+1)}\|_\infty + \frac{\gamma \|\mathbf{V}^{(l)}\|_\infty \beta_1}{(1-\gamma)N} \\ &\leq \gamma \sqrt{\frac{\beta_1}{N}} \|\widehat{\mathbf{V}}^{(l+1)} - \mathbf{V}^{(l+1)}\|_\infty + \gamma \sqrt{\frac{\beta_1}{N}} \|\mathbf{V}^{(l+1)}\|_\infty + \frac{\gamma \beta_1}{(1-\gamma)N} \|\mathbf{V}^{(l)}\|_\infty,\end{aligned}\tag{76}$$

where the last inequality is a consequence of the triangle inequality.

The above inequality (76) provides a recursive relation that in turn allows for an effective upper bound. Specifically, combining the inequalities (73) and (76) leads to

$$\begin{aligned}\|\widehat{\mathbf{V}}^{(0)} - \mathbf{V}^{(0)}\|_\infty &\leq \gamma \sqrt{\frac{\beta_1}{N}} \|\widehat{\mathbf{V}}^{(1)} - \mathbf{V}^{(1)}\|_\infty + \gamma \sqrt{\frac{\beta_1}{N}} \|\mathbf{V}^{(1)}\|_\infty + \frac{\gamma \beta_1}{(1-\gamma)N} \|\mathbf{V}^{(0)}\|_\infty \\ &=: b_1 \|\widehat{\mathbf{V}}^{(1)} - \mathbf{V}^{(1)}\|_\infty + b_1 \|\mathbf{V}^{(1)}\|_\infty + \frac{\gamma \beta_1}{(1-\gamma)N} \|\mathbf{V}^{(0)}\|_\infty,\end{aligned}$$

and for $l \geq 1$,

$$\begin{aligned}\|\widehat{\mathbf{V}}^{(l)} - \mathbf{V}^{(l)}\|_\infty &\leq \gamma \sqrt{\frac{\beta_1}{N}} \|\widehat{\mathbf{V}}^{(l+1)} - \mathbf{V}^{(l+1)}\|_\infty + \left(4 \sqrt{\frac{\beta_1}{(1-\gamma)N}} + \frac{\gamma \beta_1}{(1-\gamma)N} \right) \left(\frac{4}{\gamma \sqrt{1-\gamma}} \right)^{l-1} \|\mathbf{V}^{(1)}\|_\infty \\ &=: b_1 \|\widehat{\mathbf{V}}^{(l+1)} - \mathbf{V}^{(l+1)}\|_\infty + b_2 b_3^{l-1} \|\mathbf{V}^{(1)}\|_\infty.\end{aligned}$$

Here for notational simplicity, we introduce

$$b_1 := \gamma \sqrt{\frac{\beta_1}{N}}, \quad b_2 := 4 \sqrt{\frac{\beta_1}{(1-\gamma)N}} + \frac{\gamma \beta_1}{(1-\gamma)N}, \quad b_3 := \frac{4}{\gamma \sqrt{1-\gamma}}.$$

Invoking the above recursive relation, we can arrange terms to reach

$$\|\widehat{\mathbf{V}}^{(0)} - \mathbf{V}^{(0)}\|_\infty \leq \underbrace{b_1^m \|\widehat{\mathbf{V}}^{(m)} - \mathbf{V}^{(m)}\|_\infty}_{=:\alpha_1} + \underbrace{\left(b_1 b_2 \sum_{l=0}^{m-2} (b_1 b_3)^l + b_1 \right)}_{=:\alpha_2} \|\mathbf{V}^{(1)}\|_\infty + \frac{\gamma \beta_1}{(1-\gamma)N} \|\mathbf{V}^{(0)}\|_\infty.\tag{77}$$

Controlling the quantity α_2 . Now it suffices to control the two terms on the right-hand side of the inequality (77) separately, towards which we shall start with the quantity α_2 . Assuming that $N \geq 64\beta_1/(1-\gamma)$, one can easily verify $b_1b_3 \leq 1/2$. The summation of the geometric sequence thus gives

$$\begin{aligned}\alpha_2 &:= b_1b_2 \sum_{l=0}^{m-2} (b_1b_3)^l + b_1 \leq \frac{b_1b_2}{1-b_1b_3} + b_1 \leq 2b_1b_2 + b_1 \\ &= \gamma\sqrt{\frac{\beta_1}{N}} \left\{ 1 + 8\sqrt{\frac{\beta_1}{(1-\gamma)N}} + \frac{2\gamma\beta_1}{(1-\gamma)N} \right\} \leq 3\gamma\sqrt{\frac{\beta_1}{N}},\end{aligned}\tag{78}$$

where the last step holds with the assumption $N \geq \frac{64\beta_1}{1-\gamma}$.

Controlling the quantity α_1 . Next, we proceed to the quantity α_1 , which requires the control of $\|\widehat{\mathbf{V}}^{(m)} - \mathbf{V}^{(m)}\|_\infty$. In view of the identity (74), we obtain

$$\begin{aligned}\|\widehat{\mathbf{V}}^{(m)} - \mathbf{V}^{(m)}\|_\infty &= \gamma\|(\mathbf{I} - \gamma\widehat{\mathbf{P}}_\pi)^{-1}(\mathbf{P}_\pi - \widehat{\mathbf{P}}_\pi)\mathbf{V}^{(m)}\|_\infty \\ &\leq \gamma\left\|(\mathbf{I} - \gamma\widehat{\mathbf{P}}_\pi)^{-1}\left(\sqrt{\frac{\beta_1}{N}}\sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(m)}]} + \frac{\|\mathbf{V}^{(m)}\|_\infty\beta_1}{N}\mathbf{1}\right)\right\|_\infty,\end{aligned}$$

where the last inequality follows from the Bernstein-type condition (75) and the fact that $(\mathbf{I} - \gamma\widehat{\mathbf{P}}_\pi)^{-1}$ has non-negative entries. By virtue of the simple relation $\sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(m)}]} \leq \|\mathbf{V}^{(m)}\|_\infty$ and the fact that $\|(\mathbf{I} - \gamma\widehat{\mathbf{P}}_\pi)^{-1}\|_1 \leq \frac{1}{1-\gamma}$ (cf. Lemma 10(c)), it is further guaranteed that

$$\begin{aligned}\|\widehat{\mathbf{V}}^{(m)} - \mathbf{V}^{(m)}\|_\infty &\leq \gamma\|(\mathbf{I} - \gamma\widehat{\mathbf{P}}_\pi)^{-1}\|_1 \cdot \left\| \sqrt{\frac{\beta_1}{N}}\sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(m)}]} + \frac{\|\mathbf{V}^{(m)}\|_\infty\beta_1}{N}\mathbf{1} \right\|_\infty, \\ &\leq \frac{\gamma}{1-\gamma} \left(\sqrt{\frac{\beta_1}{N}} + \frac{\beta_1}{N} \right) \|\mathbf{V}^{(m)}\|_\infty,\end{aligned}$$

which combined with the bound (73) yields

$$\|\widehat{\mathbf{V}}^{(m)} - \mathbf{V}^{(m)}\|_\infty \leq \frac{\gamma}{1-\gamma} \left(\sqrt{\frac{\beta_1}{N}} + \frac{\beta_1}{N} \right) \left(\frac{4}{\gamma\sqrt{1-\gamma}} \right)^{m-1} \|\mathbf{V}^{(1)}\|_\infty.$$

Putting the above bounds together yields

$$\begin{aligned}\alpha_1 &:= b_1^m \|\widehat{\mathbf{V}}^{(m)} - \mathbf{V}^{(m)}\|_\infty \leq \left(\gamma\sqrt{\frac{\beta_1}{N}} \right)^m \frac{\gamma}{1-\gamma} \left(\sqrt{\frac{\beta_1}{N}} + \frac{\beta_1}{N} \right) \left(\frac{4}{\gamma\sqrt{1-\gamma}} \right)^{m-1} \|\mathbf{V}^{(1)}\|_\infty \\ &= \left(\sqrt{\frac{16\beta_1}{N(1-\gamma)}} \right)^{m-1} \left(\sqrt{\frac{\beta_1}{N}} + 1 \right) \frac{\gamma^2\beta_1}{(1-\gamma)N} \|\mathbf{V}^{(1)}\|_\infty \\ &\leq \left(\frac{1}{e} \right)^{m-1} \frac{1.1\gamma^2\beta_1}{(1-\gamma)N} \|\mathbf{V}^{(1)}\|_\infty,\end{aligned}\tag{79}$$

where the last inequality holds provided that $N > \frac{16e^2}{1-\gamma}\beta_1$.

Putting all this together. Combining the inequalities (77), (78) and (79) gives

$$\|\widehat{\mathbf{V}}^{(0)} - \mathbf{V}^{(0)}\|_\infty \leq \frac{\gamma\beta_1}{(1-\gamma)N} \|\mathbf{V}^{(0)}\|_\infty + \left\{ 3\gamma\sqrt{\frac{\beta_1}{N}} + \left(\frac{1}{e} \right)^{m-1} \frac{1.1\gamma^2\beta_1}{(1-\gamma)N} \right\} \|\mathbf{V}^{(1)}\|_\infty.$$

To finish up, set $m = \log(\frac{e}{1-\gamma})$ and assume that $N > \frac{16e^2}{1-\gamma}\beta_1$. Recognizing that $\mathbf{V}^{(0)} = \mathbf{V}^\pi$ and $\widehat{\mathbf{V}}^{(0)} = \widehat{\mathbf{V}}^\pi$, we arrive at

$$\begin{aligned} \|\widehat{\mathbf{V}}^\pi - \mathbf{V}^\pi\|_\infty &= \|\widehat{\mathbf{V}}^{(0)} - \mathbf{V}^{(0)}\|_\infty \leq \frac{\gamma\beta_1}{(1-\gamma)N} \|\mathbf{V}^{(0)}\|_\infty + \left\{ 3\gamma\sqrt{\frac{\beta_1}{N}} + \left(\frac{1}{e}\right)^{m-1} \frac{1.1\gamma^2\beta_1}{(1-\gamma)N} \right\} \|\mathbf{V}^{(1)}\|_\infty \\ &\leq \frac{\gamma\beta_1}{(1-\gamma)N} \|\mathbf{V}^\pi\|_\infty + 4\gamma\sqrt{\frac{\beta_1}{N}} \left\| (\mathbf{I} - \gamma\mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^\pi]} \right\|_\infty, \end{aligned} \quad (80)$$

provided that $N \geq \frac{16e^2\beta_1}{1-\gamma}$.

Proof of Lemma 2. Invoking the inequality (70) to bound the second term of (80), we reach

$$\|\widehat{\mathbf{V}}^\pi - \mathbf{V}^\pi\|_\infty \leq \frac{\gamma\beta_1}{(1-\gamma)N} \|\mathbf{V}^\pi\|_\infty + 8\log 2 \sqrt{\frac{\beta_1}{N(1-\gamma)^3}} \|\mathbf{r}\|_\infty \leq 6\sqrt{\frac{\beta_1}{N(1-\gamma)^3}} \|\mathbf{r}\|_\infty, \quad (81)$$

where the last inequality uses the elementary fact that $\|\mathbf{V}^\pi\|_\infty \leq \frac{1}{1-\gamma} \|\mathbf{r}\|_\infty$ and the assumption that $N > \frac{16e^2}{1-\gamma}\beta_1$. We complete the proof of Lemma 2.

Proof of Lemma 1. Finally, to establish Lemma 1, we observe that: for any fixed policy π , the vector $\mathbf{V}^{(l)}$ is independent of $\widehat{\mathbf{P}}_\pi$. The Bernstein inequality (e.g. (Agarwal et al., 2020, Lemma 6)) then reveals that with probability at least $1 - \delta$,

$$\left| (\widehat{\mathbf{P}}_\pi - \mathbf{P}_\pi) \mathbf{V}^{(l)} \right| \leq \sqrt{\frac{2 \log\left(\frac{4m|S|}{\delta}\right)}{N}} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^{(l)}]} + \frac{\|\mathbf{V}^{(l)}\|_\infty \log\left(\frac{4m|S|}{\delta}\right)}{N} \mathbf{1} \quad (82)$$

holds uniformly for all $0 \leq l \leq m$. This means that we can take $\beta_1 := 2 \log\left(\frac{4m|S|}{\delta}\right)$ with $m = \log(\frac{e}{1-\gamma})$ for this case. Combining this with the inequality (80), we derive the advertised instance-dependent bound

$$\|\widehat{\mathbf{V}}^\pi - \mathbf{V}^\pi\|_\infty \leq \frac{2\gamma \log\left(\frac{4|S| \log \frac{e}{1-\gamma}}{\delta}\right)}{(1-\gamma)N} \|\mathbf{V}^\pi\|_\infty + 4\gamma \sqrt{\frac{2 \log\left(\frac{4|S| \log \frac{e}{1-\gamma}}{\delta}\right)}{(1-\gamma)N}} \left\| (\mathbf{I} - \gamma\mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}[\mathbf{V}^\pi]} \right\|_\infty.$$

Further, this taken collectively with (70) and the crude bound $\|\mathbf{V}^\pi\|_\infty \leq \frac{1}{1-\gamma} \|\mathbf{r}\|_\infty$ gives

$$\begin{aligned} \|\widehat{\mathbf{V}}^\pi - \mathbf{V}^\pi\|_\infty &\leq \frac{2\gamma \log\left(\frac{4|S| \log \frac{e}{1-\gamma}}{\delta}\right)}{(1-\gamma)^2 N} \|\mathbf{r}\|_\infty + 8\log 2 \sqrt{\frac{2 \log\left(\frac{4|S| \log \frac{e}{1-\gamma}}{\delta}\right)}{(1-\gamma)^3 N}} \|\mathbf{r}\|_\infty \\ &\leq 6\sqrt{\frac{2 \log\left(\frac{4|S| \log \frac{e}{1-\gamma}}{\delta}\right)}{(1-\gamma)^3 N}} \|\mathbf{r}\|_\infty, \end{aligned}$$

with the proviso that $N \geq \frac{32e^2}{1-\gamma} \log\left(\frac{4|S| \log \frac{e}{1-\gamma}}{\delta}\right)$.

B.1.1 Proof of Lemma 11

To begin with, we make the observation that

$$\begin{aligned} \text{Var}_{\mathbf{P}_\pi}(\mathbf{V}) &= \mathbf{P}_\pi(\mathbf{V} \circ \mathbf{V}) - (\mathbf{P}_\pi \mathbf{V}) \circ (\mathbf{P}_\pi \mathbf{V}) \\ &= \mathbf{P}_\pi(\mathbf{V} \circ \mathbf{V}) - \frac{1}{\gamma^2} (\mathbf{V} - \mathbf{r}_\pi) \circ (\mathbf{V} - \mathbf{r}_\pi) \\ &= \mathbf{P}_\pi(\mathbf{V} \circ \mathbf{V}) - \frac{1}{\gamma^2} \mathbf{V} \circ \mathbf{V} + \frac{2}{\gamma^2} \mathbf{V} \circ \mathbf{r}_\pi - \frac{1}{\gamma^2} \mathbf{r}_\pi \circ \mathbf{r}_\pi \end{aligned} \quad (83)$$

$$\leq \frac{1}{\gamma^2}(\gamma^2 \mathbf{P}_\pi - \mathbf{I})(\mathbf{V} \circ \mathbf{V}) + \frac{2}{\gamma^2} \mathbf{V} \circ \mathbf{r}_\pi, \quad (84)$$

where the identity (83) makes use of the relation $\mathbf{V} = \mathbf{r}_\pi + \gamma \mathbf{P}_\pi \mathbf{V}$. In addition, one can deduce that

$$\begin{aligned} \left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}(\mathbf{V})} \right\|_\infty &= \frac{1}{1-\gamma} \left\| (1-\gamma)(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}(\mathbf{V})} \right\|_\infty \\ &\stackrel{(i)}{\leq} \frac{1}{\sqrt{1-\gamma}} \left\| \sqrt{(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \text{Var}_{\mathbf{P}_\pi}(\mathbf{V})} \right\|_\infty \\ &\stackrel{(ii)}{\leq} \frac{1}{\sqrt{1-\gamma}} \sqrt{\|2(\mathbf{I} - \gamma^2 \mathbf{P}_\pi)^{-1} \text{Var}_{\mathbf{P}_\pi}(\mathbf{V})\|_\infty}. \end{aligned}$$

Here, (i) comes from Jensen's inequality (so that $\mathbb{E}[\sqrt{v}] \leq \sqrt{\mathbb{E}[v]}$) recognizing that each row of $(1-\gamma)(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1}$ is a probability distribution, and Lemma 10(d), (ii) is an elementary fact established in (Agarwal et al., 2020, Lemma 4). Combining Lemma 10(e) and the inequality (84) further yields

$$\begin{aligned} \left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}(\mathbf{V})} \right\|_\infty &\leq \frac{1}{\sqrt{1-\gamma}} \sqrt{\left\| 2(\mathbf{I} - \gamma^2 \mathbf{P}_\pi)^{-1} \left(\frac{1}{\gamma^2} (\gamma^2 \mathbf{P}_\pi - \mathbf{I})(\mathbf{V} \circ \mathbf{V}) + \frac{2}{\gamma^2} \mathbf{V} \circ \mathbf{r}_\pi \right) \right\|_\infty} \\ &\leq \frac{1}{\gamma \sqrt{1-\gamma}} \sqrt{2 \|\mathbf{V} \circ \mathbf{V}\|_\infty} + \frac{2}{\gamma \sqrt{1-\gamma}} \sqrt{\|(\mathbf{I} - \gamma^2 \mathbf{P}_\pi)^{-1} (\mathbf{V} \circ \mathbf{r}_\pi)\|_\infty}. \quad (85) \end{aligned}$$

where the last step arises from the triangle inequality and $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$. This leaves us with two terms to deal with.

Regarding the first term of (85), we observe that $\|(\mathbf{V} \circ \mathbf{V})\|_\infty = \|\mathbf{V}\|_\infty^2$. When it comes to the second term of (85), it is seen that

$$\begin{aligned} \left\| (\mathbf{I} - \gamma^2 \mathbf{P}_\pi)^{-1} (\mathbf{V} \circ \mathbf{r}_\pi) \right\|_\infty &\leq \left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} (\mathbf{V} \circ \mathbf{r}_\pi) \right\|_\infty \\ &\leq \|\mathbf{V}\|_\infty \left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \mathbf{r}_\pi \right\|_\infty \\ &= \|\mathbf{V}\|_\infty^2. \end{aligned}$$

Here, the first inequality holds true since $(\mathbf{I} - \gamma^2 \mathbf{P}_\pi)^{-1} = \sum_{i=0}^{\infty} \gamma^{2i} \mathbf{P}_\pi^i \leq \sum_{i=0}^{\infty} \gamma^i \mathbf{P}_\pi^i = (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1}$, while the second line holds since \mathbf{V} , \mathbf{r} and $(\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1}$ are all non-negative. Substitution into (85) thus yields

$$\left\| (\mathbf{I} - \gamma \mathbf{P}_\pi)^{-1} \sqrt{\text{Var}_{\mathbf{P}_\pi}(\mathbf{V})} \right\|_\infty \leq \frac{\sqrt{2}}{\gamma \sqrt{1-\gamma}} \|\mathbf{V}\|_\infty + \frac{2}{\gamma \sqrt{1-\gamma}} \|\mathbf{V}\|_\infty \leq \frac{4}{\gamma \sqrt{1-\gamma}} \|\mathbf{V}\|_\infty$$

as claimed.

B.2 Proof of Lemma 3

To establish Lemma 3, it suffices to check that \mathbf{V}^* and \mathbf{Q}^* satisfy the Bellman optimality equations underlying \mathcal{M}_{s,a,u^*} . Towards this end, we study the absorbing state-action pair (s, a) and other pairs separately. For notational simplicity, we shall let \mathbf{P}^{abs} and $r^{\text{abs}}(\cdot, \cdot)$ denote respectively the probability transition matrix and the reward function associated with \mathcal{M}_{s,a,u^*} .

First, we observe that, by construction,

$$r^{\text{abs}}(s, a) + \gamma(\mathbf{P}^{\text{abs}} \mathbf{V}^*)_{s,a} = u^* + \gamma V^*(s).$$

Recall that \mathbf{V}^* satisfies the Bellman optimality equation w.r.t. the original MDP, namely, $Q^*(s, a) = r(s, a) + \gamma(\mathbf{P} \mathbf{V}^*)_{s,a}$. This together with our choice of u^* gives

$$u^* + \gamma V^*(s) = Q^*(s, a) - \gamma V^*(s) + \gamma V^*(s) = Q^*(s, a).$$

Putting the above identities together, we arrive at

$$Q^*(s, a) = r^{\text{abs}}(s, a) + \gamma(\mathbf{P}^{\text{abs}} \mathbf{V}^*)_{s,a}. \quad (86)$$

Next, consider any state-action pair $(s', a') \neq (s, a)$. Recalling again the properties of \mathcal{M}_{s,a,u^*} , we reach

$$r^{\text{abs}}(s', a') + \gamma(\mathbf{P}^{\text{abs}}\mathbf{V}^*)_{s',a'} = r(s', a') + \gamma(\mathbf{P}\mathbf{V}^*)_{s',a'} = Q^*(s', a'). \quad (87)$$

Here the last identity is due to the Bellman equation w.r.t. the original MDP. Combining (86) and (87) implies that \mathbf{V}^* and \mathbf{Q}^* satisfy Bellman's optimality equations in \mathcal{M}_{s,a,u^*} , thus concluding the proof.

B.3 Proof of Lemma 4

Our first observation is that $\widehat{\mathbf{Q}}_{s,a,u}^*$ satisfies Lipschitz continuity w.r.t. u in the sense that

$$\|\widehat{\mathbf{Q}}_{s,a,u}^* - \widehat{\mathbf{Q}}_{s,a,u'}^*\|_\infty \leq \frac{1}{1-\gamma}|u - u'|. \quad (88)$$

The proof of this relation is identical to that of (Agarwal et al., 2020, Lemma 8); we omit here for brevity. In view of Lemma 3, if we set $\widehat{u}^* := r(s, a) + \gamma(\widehat{\mathbf{P}}\widehat{\mathbf{V}}^*)_{s,a} - \gamma\widehat{\mathbf{V}}^*(s)$, then one has

$$\widehat{\mathbf{Q}}_{s,a,\widehat{u}^*}^* = \widehat{\mathbf{Q}}^* \quad \text{and} \quad \widehat{\mathbf{V}}_{s,a,\widehat{u}^*}^* = \widehat{\mathbf{V}}^*.$$

In addition, there exists a point u_0 in the epsilon-net $\mathcal{N}_{(1-\gamma)\omega/4}$ such that $|\widehat{u}^* - u_0| \leq (1-\gamma)\omega/4$, which combined with the Lipschitz continuity property (88) gives

$$\|\widehat{\mathbf{Q}}^* - \widehat{\mathbf{Q}}_{s,a,u_0}^*\|_\infty = \|\widehat{\mathbf{Q}}_{s,a,\widehat{u}^*}^* - \widehat{\mathbf{Q}}_{s,a,u_0}^*\|_\infty \leq \frac{1}{1-\gamma}|\widehat{u}^* - u_0| \leq \frac{\omega}{4}. \quad (89)$$

Additionally, for any $s' \in \mathcal{S}$ and any $a_1, a_2 \in \mathcal{A}$ with $a_1 \neq a_2$, we have the following decomposition

$$\begin{aligned} & \widehat{\mathbf{Q}}_{s,a,u_0}^*(s', a_1) - \widehat{\mathbf{Q}}_{s,a,u_0}^*(s', a_2) \\ &= \widehat{\mathbf{Q}}^*(s', a_1) - \widehat{\mathbf{Q}}^*(s', a_2) + \widehat{\mathbf{Q}}_{s,a,u_0}^*(s', a_1) - \widehat{\mathbf{Q}}^*(s', a_1) - \left(\widehat{\mathbf{Q}}_{s,a,u_0}^*(s', a_2) - \widehat{\mathbf{Q}}^*(s', a_2) \right) \\ &\geq \widehat{\mathbf{Q}}^*(s', a_1) - \widehat{\mathbf{Q}}^*(s', a_2) - 2\|\widehat{\mathbf{Q}}^* - \widehat{\mathbf{Q}}_{s,a,u_0}^*\|_\infty \\ &\geq \widehat{\mathbf{Q}}^*(s', a_1) - \widehat{\mathbf{Q}}^*(s', a_2) - \frac{\omega}{2}, \end{aligned} \quad (90)$$

where the last inequality invokes the inequality (89). Moreover, our separation condition defined in (41) requires that: for any $s' \in \mathcal{S}$ and any $a_2 \neq \widehat{\pi}^*(s')$, one has $\widehat{\mathbf{Q}}^*(s', \widehat{\pi}^*(s')) - \widehat{\mathbf{Q}}^*(s', a_2) \geq \omega$, which together with (90) reveals that

$$\widehat{\mathbf{Q}}_{s,a,u_0}^*(s', \widehat{\pi}^*(s')) - \widehat{\mathbf{Q}}_{s,a,u_0}^*(s', a_2) \geq \widehat{\mathbf{Q}}^*(s', \widehat{\pi}^*(s')) - \widehat{\mathbf{Q}}^*(s', a_2) - \frac{\omega}{2} \geq \frac{\omega}{2}. \quad (91)$$

Given that $\widehat{\pi}_{s,a,u_0}^*(s') := \arg \max_{a'} \widehat{\mathbf{Q}}_{s,a,u_0}^*(s', a')$, it is seen from (91) that

$$\widehat{\pi}_{s,a,u_0}^*(s') = \widehat{\pi}^*(s'),$$

which holds true for all $s' \in \mathcal{S}$. This concludes the proof.

B.4 Proof of Lemma 5

We start by bounding $\|\widehat{\mathbf{V}}^{\widehat{\pi}^*} - \mathbf{V}^{\widehat{\pi}^*}\|_\infty$. Recall the definition of the series $\{\mathbf{V}^{(l)}\}$ in (69). Throughout this proof, we shall write $\mathbf{V}_\pi^{(l)}$ instead in order to make apparent the dependency on the policy π .

For each state-action pair (s, a) , let us construct the epsilon-net $\mathcal{N}_{(1-\gamma)\omega/4}$ as in the expression (42). For every $u \in \mathcal{N}_{(1-\gamma)\omega/4}$, recall that $\widehat{\pi}_{s,a,u}^*$ is defined as the optimal policy with respect to the (s, a) -absorbing MDP $\widehat{\mathcal{M}}_{s,a,u}$. By construction, the set of policies $\widehat{\pi}_{s,a,u}^*$ ($u \in \mathcal{N}_{(1-\gamma)\omega/4}$) is independent of $\widehat{\mathbf{P}}_{s,a}$. The Bernstein inequality (e.g. (Agarwal et al., 2020, Lemma 6)) taken together with the union bound thus guarantees that with probability at least $1 - \delta$,

$$\left| (\widehat{\mathbf{P}} - \mathbf{P})_{s,a} \mathbf{V}_{\widehat{\pi}_{s,a,u}^*}^{(l)} \right| \leq \sqrt{\frac{\beta_1}{N}} \sqrt{(\text{Var}_{\mathbf{P}}[\mathbf{V}_{\widehat{\pi}_{s,a,u}^*}^{(l)}])_{s,a}} + \frac{\|\mathbf{V}_{\widehat{\pi}_{s,a,u}^*}^{(l)}\|_\infty \beta_1}{N} \quad (92)$$

holds uniformly over all $0 \leq l \leq \log \frac{e}{1-\gamma}$, $u \in \mathcal{N}_{(1-\gamma)\omega/4}$, $(s, a) \in \mathcal{S} \times \mathcal{A}$. Here, β_1 is given by

$$\beta_1 := 2 \log \left(\frac{4 \log \left(\frac{e}{1-\gamma} \right) |\mathcal{N}_{(1-\gamma)\omega/4}| |\mathcal{S}| |\mathcal{A}|}{\delta} \right) \leq 2 \log \left(\frac{32}{(1-\gamma)^2 \omega \delta} |\mathcal{S}| |\mathcal{A}| \log \left(\frac{e}{1-\gamma} \right) \right),$$

where we have used the fact $|\mathcal{N}_{(1-\gamma)\omega/4}| \leq \frac{8}{(1-\gamma)^2 \omega}$. In addition, for any $0 < \omega < 1$, Lemma 4 guarantees that for each state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, there exists a point $u_0 \in \mathcal{N}_{(1-\gamma)\omega/4}$ such that $\hat{\pi}^* = \hat{\pi}_{s,a,u_0}^*$. Invoking this important fact, we obtain

$$\begin{aligned} \left| (\hat{\mathbf{P}} - \mathbf{P})_{s,a} \mathbf{V}_{\hat{\pi}^*}^{(l)} \right| &= \left| (\hat{\mathbf{P}} - \mathbf{P})_{s,a} \mathbf{V}_{\hat{\pi}_{s,a,u_0}^*}^{(l)} \right| \\ &\leq \sqrt{\frac{\beta_1}{N}} \sqrt{(\text{Var}_{\mathbf{P}} [\mathbf{V}_{\hat{\pi}_{s,a,u_0}^*}^{(l)}])_{s,a}} + \frac{\|\mathbf{V}_{\hat{\pi}_{s,a,u_0}^*}^{(l)}\|_{\infty} \beta_1}{N} \\ &= \sqrt{\frac{\beta_1}{N}} \sqrt{(\text{Var}_{\mathbf{P}} [\mathbf{V}_{\hat{\pi}^*}^{(l)}])_{s,a}} + \frac{\|\mathbf{V}_{\hat{\pi}^*}^{(l)}\|_{\infty} \beta_1}{N}. \end{aligned}$$

The above inequality further allows us to deduce that, with probability $1 - \delta$,

$$\begin{aligned} \left| (\hat{\mathbf{P}}_{\hat{\pi}^*} - \mathbf{P}_{\hat{\pi}^*}) \mathbf{V}_{\hat{\pi}^*}^{(l)} \right| &= \left| \mathbf{\Pi}^{\hat{\pi}^*} (\hat{\mathbf{P}} - \mathbf{P}) \mathbf{V}_{\hat{\pi}^*}^{(l)} \right| \\ &\leq \sqrt{\frac{\beta_1}{N}} \sqrt{\mathbf{\Pi}^{\hat{\pi}^*} \text{Var}_{\mathbf{P}} [\mathbf{V}_{\hat{\pi}^*}^{(l)}]} + \frac{\|\mathbf{V}_{\hat{\pi}^*}^{(l)}\|_{\infty} \beta_1}{N} \mathbf{1} \\ &= \sqrt{\frac{\beta_1}{N}} \sqrt{\text{Var}_{\mathbf{P}_{\hat{\pi}^*}} [\mathbf{V}_{\hat{\pi}^*}^{(l)}]} + \frac{\|\mathbf{V}_{\hat{\pi}^*}^{(l)}\|_{\infty} \beta_1}{N} \mathbf{1}. \end{aligned}$$

The above derivation validates the assumption required for Lemma 2. As a result, if $N > \frac{16e^2}{1-\gamma} \beta_1$ and $0 < \omega < 1$, then Lemma 2 leads to the advertised bound

$$\begin{aligned} \|\hat{\mathbf{V}}^{\hat{\pi}^*} - \mathbf{V}^{\hat{\pi}^*}\|_{\infty} &\leq \frac{6}{1-\gamma} \sqrt{\frac{\beta_1}{N(1-\gamma)}} \leq \frac{6}{1-\gamma} \sqrt{\frac{2 \log \left(\frac{32}{(1-\gamma)^2 \omega \delta} |\mathcal{S}| |\mathcal{A}| \log \left(\frac{e}{1-\gamma} \right) \right)}{N(1-\gamma)}} \\ &\leq 6 \sqrt{\frac{2 \log \left(\frac{32 |\mathcal{S}| |\mathcal{A}|}{(1-\gamma)^3 \omega \delta} \right)}{N(1-\gamma)^3}}. \end{aligned} \quad (93)$$

Finally, we move on to the term $\mathbf{V}^* - \hat{\mathbf{V}}^{\pi^*}$. Given that π^* is independent of $\hat{\mathbf{P}}$, invoke Lemma 1 to reach

$$\begin{aligned} \|\hat{\mathbf{V}}^{\pi^*} - \mathbf{V}^*\|_{\infty} &= \|\hat{\mathbf{V}}^{\pi^*} - \mathbf{V}^{\pi^*}\|_{\infty} \leq 6\sqrt{2} \sqrt{\frac{\log \left(\frac{4|\mathcal{S}|}{\delta} \right) + \log \log \left(\frac{e}{1-\gamma} \right)}{N(1-\gamma)^3}} \\ &\leq 6 \sqrt{\frac{2 \log \left(\frac{32 |\mathcal{S}| |\mathcal{A}|}{(1-\gamma)^2 \omega \delta} \right)}{N(1-\gamma)^3}} \end{aligned} \quad (94)$$

with probability at least $1 - \delta$. This together with (34) and (93) immediately establishes Lemma 5.

B.5 Proof of Lemma 6

The proofs for Q_p^* and \hat{Q}_p^* are exactly the same; for the sake of conciseness, we shall only provide the proof for Q_p^* . Here we aim to prove a more general result than Lemma 6, namely, with probability at least $1 - \delta$,

$$\forall s \in \mathcal{S} \text{ and } a_1, a_2 \in \mathcal{A} \text{ with } a_1 \neq a_2 : \quad \left| Q_p^*(s, a_1) - Q_p^*(s, a_2) \right| > \frac{\xi \delta (1-\gamma)}{4|\mathcal{S}| |\mathcal{A}|^2}.$$

Consider any state s and any actions $a_1 \neq a_2$. In what follows, we allow $r_p(s, a_1) = \tau$ to vary, while *freezing* the values of all other rewards $\{r_p(\tilde{s}, a) \mid (\tilde{s}, a) \neq (s, a_1)\}$. To streamline notation, we define

- $\mathbf{r}_\tau = [r_\tau(s, a)]_{(s, a) \in \mathcal{S} \times \mathcal{A}}$: the reward vector obeying

$$r_\tau(s, a_1) = \tau \quad \text{and} \quad r_\tau(\tilde{s}, a) = r_p(\tilde{s}, a) \quad \text{for all } (\tilde{s}, a) \neq (s, a_1);$$

- Q_τ^* : the optimal Q-function when the reward vector is \mathbf{r}_τ ;
- V_τ^* : the optimal value function when the reward vector is \mathbf{r}_τ ;
- π_τ^* : the optimal policy when the reward vector is \mathbf{r}_τ .

Additionally, we claim for the moment that there exists a phase transition boundary τ_{th} such that

$$\pi_\tau^*(s) \neq a_1, \quad \text{for all } \tau < \tau_{\text{th}}; \quad (95a)$$

$$\pi_\tau^*(s) = a_1, \quad \text{for all } \tau > \tau_{\text{th}}. \quad (95b)$$

The proof of this claim is deferred to the end of this section. To establish Lemma 6, the idea is to control the size of the set

$$\mathcal{I}_{0, \omega} := \{\tau \mid |Q_\tau^*(s, a_1) - Q_\tau^*(s, a_2)| < \omega\} \quad (96)$$

for some $\omega > 0$ to be specified shortly. As motivated by (95), we further break down this set into two parts $\mathcal{I}_{0, \omega} = \mathcal{I}_{1, \omega} \cup \mathcal{I}_{2, \omega}$, where

$$\mathcal{I}_{1, \omega} := \{\tau \mid \tau < \tau_{\text{th}}, |Q_\tau^*(s, a_1) - Q_\tau^*(s, a_2)| < \omega\}, \quad (97a)$$

$$\mathcal{I}_{2, \omega} := \{\tau \mid \tau \geq \tau_{\text{th}}, |Q_\tau^*(s, a_1) - Q_\tau^*(s, a_2)| < \omega\}. \quad (97b)$$

In what follows, we first control the size of each set separately, and then demonstrate that the probability of these events happening is very small.

Step 1. We begin with $\mathcal{I}_{1, \omega}$ associated with the range $\tau < \tau_{\text{th}}$. In this case, the value function V_τ^* does not vary with τ , since the reward $r_\tau(s, a_1) = \tau$ is never active when calculating V_τ^* (by virtue of (95a)). Thus, the Bellman equation allows us to write

$$Q_\tau^*(s, a_1) = \tau + B_1 \quad \text{and} \quad Q_\tau^*(s, a_2) = B_2$$

for some quantities B_1 and B_2 , where neither B_1 nor B_2 depends on the value of τ . Armed with this observation, we can easily show that: for any $\omega > 0$, the interval $\mathcal{I}_{1, \omega}$ (cf. (97a)) obeys

$$\mathcal{I}_{1, \omega} \subseteq \{\tau \mid |\tau + B_1 - B_2| < \omega\},$$

and hence has length (or Lebesgue measure) at most 2ω .

Step 2. We then move on to $\mathcal{I}_{2, \omega}$ associated with the range $\tau > \tau_{\text{th}}$ in which case $\pi_\tau^*(s) = a_1$. Towards this, we first make some useful observations.

- To begin with, given the relation $\mathbf{Q}_\tau^* = \mathbf{r}_\tau + \gamma \mathbf{P} \mathbf{V}_\tau^*$, it is easily seen that for any $\tau_2 > \tau_1 > \tau_{\text{th}}$,

$$\mathbf{0} \leq \mathbf{Q}_{\tau_2}^* - \mathbf{Q}_{\tau_1}^* \leq \mathbf{r}_{\tau_2} - \mathbf{r}_{\tau_1} + \gamma \|\mathbf{V}_{\tau_2}^* - \mathbf{V}_{\tau_1}^*\|_\infty. \quad (98)$$

In addition, for any state-action pair $(\tilde{s}, a) \neq (s, a_1)$, by construction we have $r_{\tau_2}(\tilde{s}, a) - r_{\tau_1}(\tilde{s}, a) = 0$, which together with (98) indicates that

$$\forall (\tilde{s}, a) \neq (s, a_1) : \quad 0 \leq Q_{\tau_2}^*(\tilde{s}, a) - Q_{\tau_1}^*(\tilde{s}, a) \leq \gamma \|\mathbf{V}_{\tau_2}^* - \mathbf{V}_{\tau_1}^*\|_\infty. \quad (99)$$

- Next, observe that for any $\tau_2 > \tau_1 > \tau_{\text{th}}$,

$$\forall s' \in \mathcal{S} : \quad 0 \leq V_{\tau_2}^*(s') - V_{\tau_1}^*(s') = \max_a Q_{\tau_2}^*(s', a) - \max_a Q_{\tau_1}^*(s', a) \leq \|Q_{\tau_2}^* - Q_{\tau_1}^*\|_{\infty} \quad (100)$$

and hence $\|Q_{\tau_2}^* - Q_{\tau_1}^*\|_{\infty} \geq \|V_{\tau_2}^* - V_{\tau_1}^*\|_{\infty}$. This combined with (99) and the fact $\gamma < 1$ implies that

$$Q_{\tau_2}^*(s, a_1) - Q_{\tau_1}^*(s, a_1) \geq \|V_{\tau_2}^* - V_{\tau_1}^*\|_{\infty}, \quad (101)$$

which together with the facts $V_{\tau_1}^*(s) = Q_{\tau_1}^*(s, a_1)$ and $V_{\tau_2}^*(s) = Q_{\tau_2}^*(s, a_1)$ (by virtue of (95b)) yields

$$\begin{aligned} V_{\tau_2}^*(s) - V_{\tau_1}^*(s) &= Q_{\tau_2}^*(s, a_1) - Q_{\tau_1}^*(s, a_1) \geq \|V_{\tau_2}^* - V_{\tau_1}^*\|_{\infty} \\ \implies Q_{\tau_2}^*(s, a_1) - Q_{\tau_1}^*(s, a_1) &= \|V_{\tau_2}^* - V_{\tau_1}^*\|_{\infty}. \end{aligned}$$

Invoke the Bellman equation to further derive

$$\begin{aligned} Q_{\tau_2}^*(s, a_1) - Q_{\tau_1}^*(s, a_1) &= \|V_{\tau_2}^* - V_{\tau_1}^*\|_{\infty} = \|\mathbf{r}_{\tau_2} + \gamma \mathbf{P}V_{\tau_2}^* - \mathbf{r}_{\tau_1} - \gamma \mathbf{P}V_{\tau_1}^*\|_{\infty} \\ &\geq \|\mathbf{r}_{\tau_2} - \mathbf{r}_{\tau_1}\|_{\infty} = \tau_2 - \tau_1, \end{aligned} \quad (102)$$

where the last inequality holds since $\mathbf{r}_{\tau_2} + \gamma \mathbf{P}V_{\tau_2}^* - \mathbf{r}_{\tau_1} - \gamma \mathbf{P}V_{\tau_1}^* \geq \mathbf{r}_{\tau_2} - \mathbf{r}_{\tau_1} \geq \mathbf{0}$ (due to the monotonicity properties $\mathbf{r}_{\tau_2} \geq \mathbf{r}_{\tau_1}$ and $V_{\tau_2}^* \geq V_{\tau_1}^*$), and the last identity follows from the definition of \mathbf{r}_{τ} .

With the above two properties (99) and (102) in mind, we are ready to locate $\mathcal{I}_{2,\omega}$ by showing that

$$\mathcal{I}_{2,\omega} \subseteq \left[\tau_{\text{th}}, \tau_{\text{th}} + \frac{\omega}{1-\gamma} \right]. \quad (103)$$

Given that $Q_{\tau_{\text{th}}}^*(s, a_1) \geq Q_{\tau_{\text{th}}}^*(s, a_2)$ (in view of (95)), we have for any $\tau \geq \tau_{\text{th}}$ and any $a_2 \neq a_1$ that

$$\begin{aligned} Q_{\tau}^*(s, a_1) - Q_{\tau}^*(s, a_2) &\geq (Q_{\tau}^*(s, a_1) - Q_{\tau_{\text{th}}}^*(s, a_1)) - (Q_{\tau}^*(s, a_2) - Q_{\tau_{\text{th}}}^*(s, a_2)) \\ &\geq (1-\gamma)(Q_{\tau}^*(s, a_1) - Q_{\tau_{\text{th}}}^*(s, a_1)). \end{aligned} \quad (104)$$

Here, the last inequality holds since

$$Q_{\tau}^*(s, a_2) - Q_{\tau_{\text{th}}}^*(s, a_2) \stackrel{(i)}{\leq} \gamma \|V_{\tau}^* - V_{\tau_{\text{th}}}^*\|_{\infty} \stackrel{(ii)}{\leq} \gamma (Q_{\tau}^*(s, a_1) - Q_{\tau_{\text{th}}}^*(s, a_1)),$$

where (i) follows from (99) and (ii) is due to (101). As a result, for any $\tau > \tau_{\text{th}} + \frac{\omega}{1-\gamma}$, one can invoke (104) and (102) to see that

$$\forall a_2 \neq a_1 : \quad Q_{\tau}^*(s, a_1) - Q_{\tau}^*(s, a_2) \geq (1-\gamma)(Q_{\tau}^*(s, a_1) - Q_{\tau_{\text{th}}}^*(s, a_1)) \geq (1-\gamma)(\tau - \tau_{\text{th}}) > \omega,$$

which necessarily implies that such a τ does not lie within the interval $\mathcal{I}_{2,\omega}$ as defined in (97b). This establishes the claimed relation (103).

Step 3. Putting the results in the above two steps together, we see the set $\mathcal{I}_{0,\omega}$ (cf. (96)) has total length (or Lebesgue measure) at most $\frac{3\omega}{1-\gamma}$. Given that $r_{\text{p}}(s, a) = r(s, a) + \zeta(s, a)$ with $\zeta(s, a) \sim \text{Unif}(0, \xi)$, one has

$$\mathbb{P} \left\{ |Q_{\text{p}}^*(s, a_1) - Q_{\text{p}}^*(s, a_2)| < \omega \right\} \leq \mathbb{P} \{ r(s, a) + \zeta(s, a) \in \mathcal{I}_{0,\omega} \} \leq \frac{3\omega}{\xi(1-\gamma)}.$$

By setting $\omega = \frac{\delta(1-\gamma)\xi}{3|\mathcal{S}||\mathcal{A}|^2}$, we arrive at

$$\mathbb{P} \left\{ |Q_{\text{p}}^*(s, a_1) - Q_{\text{p}}^*(s, a_2)| < \frac{\delta(1-\gamma)\xi}{3|\mathcal{S}||\mathcal{A}|^2} \right\} \leq \frac{\delta}{|\mathcal{S}||\mathcal{A}|^2}.$$

Finally, taking the union bound over all s, a_1, a_2 , we conclude that

$$\mathbb{P} \left\{ \exists s, a_1 \neq a_2 : |Q_{\text{p}}^*(s, a_1) - Q_{\text{p}}^*(s, a_2)| < \frac{\delta(1-\gamma)\xi}{3|\mathcal{S}||\mathcal{A}|^2} \right\} \leq \delta,$$

thus establishing Lemma 6 as long as the claim (95) is valid.

Proof of the claim (95). To establish the claim, it suffices to take

$$\tau_{\text{th}} = \sup \{u \mid \pi_\tau^*(s) \neq a_1 \text{ for all } \tau < u\}. \quad (105)$$

It thus suffices to verify (95b) for our choice (105). Towards this, suppose instead that there exist some $\tau_3 < \tau_{\text{th}} \leq \tau_2 < \tau_1$ such that

$$\pi_{\tau_3}^*(s) \neq a_1, \quad \pi_{\tau_2}^*(s) = a_1, \quad \text{and} \quad \pi_{\tau_1}^*(s) \neq a_1.$$

It is straightforward to see that $V_{\tau_1}^* = V_{\tau_3}^*$, since in both cases, the reward $r_\tau(s, a_1)$ does not enter the calculation of the optimal value function (while the rewards in other state-action pairs are identical in both cases). In view of the monotonicity of the value function w.r.t. the reward vector, we have

$$V_{\tau_1}^* = V_{\tau_2}^* = V_{\tau_3}^*.$$

However, this contradicts our assumption that a_1 is the optimal action for state s at τ_2 but not at τ_3 , since enlarging τ_2 to τ_3 otherwise will enlarge the optimal value function $V_{\tau_3}^*$. We have thus established (95). \square

B.6 Proof of Lemma 7

To begin, we find it helpful to introduce a modified reward function $\tilde{\mathbf{r}} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ as follows

$$\tilde{r}(s, a) := \begin{cases} r(s, a) + \hat{V}^*(s) - \hat{Q}^*(s, a), & \text{if } a = \hat{\pi}_c(s), \\ r(s, a), & \text{otherwise.} \end{cases} \quad (106)$$

Armed with this new reward function, we subsequently define a vector $\tilde{\mathbf{Q}} = [\tilde{Q}(s, a)]$ as follows

$$\tilde{\mathbf{Q}} := \tilde{\mathbf{r}} + \gamma \hat{\mathbf{P}} \hat{\mathbf{V}}^*. \quad (107)$$

In view of the Bellman optimality equation $\hat{\mathbf{Q}}^* = \mathbf{r} + \gamma \hat{\mathbf{P}} \hat{\mathbf{V}}^*$, we see that $\tilde{\mathbf{Q}}$ satisfies $\tilde{\mathbf{Q}} = \tilde{\mathbf{r}} + \hat{\mathbf{Q}}^* - \mathbf{r}$, which combined with the construction (106) gives

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}: \quad \tilde{Q}(s, a) = \begin{cases} \hat{V}^*(s), & \text{if } a = \hat{\pi}_c(s), \\ \hat{Q}^*(s, a), & \text{else.} \end{cases}$$

As a consequence, it is easily seen that

$$\forall s \in \mathcal{S}: \quad \max_a \tilde{Q}(s, a) = \hat{V}^*(s) \quad \text{and} \quad \tilde{Q}(s, \hat{\pi}_c(s)) = \hat{V}^*(s)$$

This taken collectively with (107) demonstrates that $\tilde{\mathbf{Q}}$ and $\hat{\mathbf{V}}^*$ are respectively the optimal Q-function and optimal value function of the MDP $\tilde{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, \tilde{\mathbf{r}}, \hat{\mathbf{P}}, \gamma)$, since they satisfy the Bellman optimality condition w.r.t. $\tilde{\mathcal{M}}$. In addition, if we let $\tilde{\mathbf{V}}^{\hat{\pi}_c}$ represent the value function of the policy $\hat{\pi}_c$ in $\tilde{\mathcal{M}}$, then the preceding relation clearly implies that $\tilde{\mathbf{V}}^{\hat{\pi}_c} = \hat{\mathbf{V}}^*$.

Using the above properties, one can deduce that

$$\begin{aligned} \hat{\mathbf{V}}^* - \hat{\mathbf{V}}^{\hat{\pi}_c} &= \tilde{\mathbf{V}}^{\hat{\pi}_c} - \hat{\mathbf{V}}^{\hat{\pi}_c} \stackrel{(i)}{=} (\mathbf{I} - \gamma \hat{\mathbf{P}}_{\hat{\pi}_c})^{-1} \tilde{\mathbf{r}} - (\mathbf{I} - \gamma \hat{\mathbf{P}}_{\hat{\pi}_c})^{-1} \mathbf{r} = (\mathbf{I} - \gamma \hat{\mathbf{P}}_{\hat{\pi}_c})^{-1} (\tilde{\mathbf{r}} - \mathbf{r}) \\ &\leq \|(\mathbf{I} - \gamma \hat{\mathbf{P}}_{\hat{\pi}_c})^{-1}\|_1 \|\tilde{\mathbf{r}} - \mathbf{r}\|_\infty \mathbf{1} \stackrel{(ii)}{=} \frac{1}{1-\gamma} \|\tilde{\mathbf{r}} - \mathbf{r}\|_\infty \mathbf{1} \stackrel{(iii)}{\leq} \frac{\varsigma}{1-\gamma} \mathbf{1} \stackrel{(iv)}{\leq} \frac{\xi}{1-\gamma} \mathbf{1}. \end{aligned}$$

Here, (i) is due to the Bellman equation, (ii) relies on the fact $\|(\mathbf{I} - \gamma \hat{\mathbf{P}}_{\hat{\pi}_c})^{-1}\|_1 = \frac{1}{1-\gamma}$, (iii) arises since $\hat{V}^*(s) - \hat{Q}^*(s, \hat{\pi}_c(s)) \leq \varsigma$ by construction of $\hat{\pi}_c$, whereas (iv) is valid since $\varsigma \in [0, \xi]$. The lemma then follows by recognizing that $\hat{\mathbf{V}}^{\pi^*} \leq \hat{\mathbf{V}}^*$ due to the optimality of $\hat{\mathbf{V}}^*$.

B.7 Proof of Lemma 8

We first make the key observation that, with probability at least $1 - \delta$, the following event holds true:

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad |\widehat{V}(s) - \widehat{Q}(s, a) - \varsigma| > \frac{\xi\delta}{2|\mathcal{S}||\mathcal{A}|}. \quad (108)$$

To justify this claim, note that its complementary event satisfies

$$\begin{aligned} \mathbb{P}\left(\exists (s, a) \in \mathcal{S} \times \mathcal{A}, |\widehat{V}(s) - \widehat{Q}(s, a) - \varsigma| \leq \frac{\xi\delta}{2|\mathcal{S}||\mathcal{A}|}\right) &\leq \sum_{s,a} \mathbb{P}\left(\varsigma \in \left[\widehat{V}(s) - \widehat{Q}(s, a) \pm \frac{\xi\delta}{2|\mathcal{S}||\mathcal{A}|}\right]\right) \\ &\leq |\mathcal{S}||\mathcal{A}| \cdot \frac{\xi\delta}{\xi|\mathcal{S}||\mathcal{A}|} = \delta, \end{aligned}$$

where the first inequality applies the union bound, and the last line follows since $\varsigma \sim \text{Unif}(0, \xi)$. Here, we abbreviate $[a \pm b] := [a - b, a + b]$. Combining (108) with the assumption (55), we further reach

$$|V(s) - Q(s, a) - \varsigma| \geq |\widehat{V}(s) - \widehat{Q}(s, a) - \varsigma| - 2 \max_{s,a} |Q(s, a) - \widehat{Q}(s, a)| > \frac{\xi\delta}{4|\mathcal{S}||\mathcal{A}|}$$

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

$$V(s) - Q(s, a) < \varsigma - \frac{\xi\delta}{4|\mathcal{S}||\mathcal{A}|} \quad \text{or} \quad V(s) - Q(s, a) > \varsigma + \frac{\xi\delta}{4|\mathcal{S}||\mathcal{A}|} \quad (109)$$

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

We are now prepared to justify the claim of this lemma. To begin with, consider any action $\widehat{a} \in \{a \in \mathcal{A} : Q(s, a) > V(s) - \varsigma\}$. Comparing this with the condition (109), we can easily see that the only possibility is

$$V(s) - Q(s, \widehat{a}) < \varsigma - \frac{\xi\delta}{4|\mathcal{S}||\mathcal{A}|}.$$

Therefore, by invoking a basic decomposition, we ensure that

$$\begin{aligned} \widehat{V}(s) - \widehat{Q}(s, \widehat{a}) &= \widehat{V}(s) - V(s) + V(s) - Q(s, \widehat{a}) + Q(s, \widehat{a}) - \widehat{Q}(s, \widehat{a}) \\ &\leq V(s) - Q(s, \widehat{a}) + |V(s) - \widehat{V}(s)| + |\widehat{Q}(s, \widehat{a}) - Q(s, \widehat{a})| \\ &< \varsigma - \frac{\xi\delta}{4|\mathcal{S}||\mathcal{A}|} + 2 \max_{s,a} |Q(s, a) - \widehat{Q}(s, a)| \\ &\leq \varsigma - \frac{\xi\delta}{4|\mathcal{S}||\mathcal{A}|} + 2 \cdot \frac{\xi\delta}{8|\mathcal{S}||\mathcal{A}|} = \varsigma. \end{aligned}$$

This essentially implies that

$$\{a \in \mathcal{A} : Q(s, a) > V(s) - \varsigma\} \subseteq \{a \in \mathcal{A} : \widehat{Q}(s, a) > \widehat{V}(s) - \varsigma\}. \quad (110)$$

Applying exactly the same argument for any action $\widehat{a} \in \{a \in \mathcal{A} : Q(s, a) \leq V(s) - \varsigma\}$, we can also derive

$$\{a \in \mathcal{A} : \widehat{Q}(s, a) > \widehat{V}(s) - \varsigma\} \subseteq \{a \in \mathcal{A} : Q(s, a) > V(s) - \varsigma\}. \quad (111)$$

These two set inequalities taken collectively establish the lemma.

C Proofs of auxiliary lemmas: finite-horizon MDPs

To begin, we find it helpful to control the entrywise magnitudes of $\{\mathbf{V}_h^{(l)}\}$. This is accomplished via the following lemma, with the proof postponed to Section C.2.

Lemma 12. *The vectors $\{\mathbf{V}_h^{(l)}\}$ defined in (62) obey*

$$\|\mathbf{V}_h^{(l)}\|_\infty \leq (\sqrt{3H})^l H \quad (112)$$

for all $1 \leq h \leq H$ and all $l \geq 0$.

C.1 Proof of Lemma 9

Consider any l obeying $1 \leq l \leq m := \log_2 H$. By construction (cf. (62)), we see that

$$\begin{aligned}\widehat{\mathbf{V}}_h^{(l)} - \mathbf{V}_h^{(l)} &= \widehat{\mathbf{P}}_{h,\pi} \widehat{\mathbf{V}}_{h+1}^{(l)} - \mathbf{P}_{h,\pi} \mathbf{V}_{h+1}^{(l)} \\ &= \widehat{\mathbf{P}}_{h,\pi} (\widehat{\mathbf{V}}_{h+1}^{(l)} - \mathbf{V}_{h+1}^{(l)}) + (\widehat{\mathbf{P}}_{h,\pi} - \mathbf{P}_{h,\pi}) \mathbf{V}_{h+1}^{(l)},\end{aligned}$$

which connects $\mathbf{V}_h^{(l)}$ and $\widehat{\mathbf{V}}_h^{(l)}$ with $\mathbf{V}_{h+1}^{(l)}$ and $\widehat{\mathbf{V}}_{h+1}^{(l)}$. Apply the above relation recursively and make use of the conditions (61) to arrive at

$$\widehat{\mathbf{V}}_h^{(l)} - \mathbf{V}_h^{(l)} = \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \widehat{\mathbf{P}}_{i,\pi} (\widehat{\mathbf{P}}_{j,\pi} - \mathbf{P}_{j,\pi}) \mathbf{V}_{j+1}^{(l)},$$

where we adopt the convenient notation and let $\prod_{i=h}^{h-1} \widehat{\mathbf{P}}_{i,\pi} = \mathbf{I}$.

According to the triangle inequality, we can further deduce that

$$\begin{aligned}|\widehat{\mathbf{V}}_h^{(l)} - \mathbf{V}_h^{(l)}| &\leq \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \widehat{\mathbf{P}}_{i,\pi} \left| (\widehat{\mathbf{P}}_{j,\pi} - \mathbf{P}_{j,\pi}) \mathbf{V}_{j+1}^{(l)} \right| \\ &\leq \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \widehat{\mathbf{P}}_{i,\pi} \left(\sqrt{\frac{\beta_1}{N}} \sqrt{\text{Var}_{\mathbf{P}_{j,\pi}}[\mathbf{V}_{j+1}^{(l)}]} + \frac{\beta_1 \|\mathbf{V}_{j+1}^{(l)}\|_\infty}{N} \mathbf{1} \right) \\ &= \sqrt{\frac{\beta_1}{N}} \left(\sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \widehat{\mathbf{P}}_{i,\pi} \mathbf{r}_j^{(l+1)} \right) + \sum_{j=h}^{H-1} \frac{\beta_1 \|\mathbf{V}_{j+1}^{(l)}\|_\infty}{N} \mathbf{1} \\ &\leq \sqrt{\frac{\beta_1}{N}} \left(\sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \widehat{\mathbf{P}}_{i,\pi} \mathbf{r}_j^{(l+1)} \right) + \frac{\beta_1 H \max_j \|\mathbf{V}_j^{(l)}\|_\infty}{N} \mathbf{1},\end{aligned}\tag{113}$$

where the second line follows from the assumption (63), and the third line makes use of the definition (62) of $\mathbf{r}_h^{(l)}$ and the elementary identity $\prod_{i=h}^{j-1} \widehat{\mathbf{P}}_{i,\pi} \mathbf{1} = \mathbf{1}$ (since each $\widehat{\mathbf{P}}_{i,\pi}$ is a probability transition matrix). In view of the construction (62), we can also derive recursively that

$$\widehat{\mathbf{V}}_h^{(l)} = \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \widehat{\mathbf{P}}_{i,\pi} \mathbf{r}_j^{(l)},\tag{114}$$

which combined with (113) yields

$$|\widehat{\mathbf{V}}_h^{(l)} - \mathbf{V}_h^{(l)}| \leq \sqrt{\frac{\beta_1}{N}} \widehat{\mathbf{V}}_h^{(l+1)} + \frac{\beta_1 H \max_j \|\mathbf{V}_j^{(l)}\|_\infty}{N} \mathbf{1}.\tag{115}$$

Further, the above inequality together with the triangle inequality immediately results in the following recursive relation

$$\begin{aligned}\|\widehat{\mathbf{V}}_h^{(l)} - \mathbf{V}_h^{(l)}\|_\infty &\leq \sqrt{\frac{\beta_1}{N}} \|\widehat{\mathbf{V}}_h^{(l+1)}\|_\infty + \frac{\beta_1 H}{N} \max_j \|\mathbf{V}_j^{(l)}\|_\infty \\ &\leq \sqrt{\frac{\beta_1}{N}} \|\widehat{\mathbf{V}}_h^{(l+1)} - \mathbf{V}_h^{(l+1)}\|_\infty + \sqrt{\frac{\beta_1}{N}} \|\mathbf{V}_h^{(l+1)}\|_\infty + \frac{\beta_1 H}{N} \max_j \|\mathbf{V}_j^{(l)}\|_\infty,\end{aligned}$$

thus revealing a useful connection between $\|\widehat{\mathbf{V}}_h^{(l)} - \mathbf{V}_h^{(l)}\|_\infty$ and $\|\widehat{\mathbf{V}}_h^{(l+1)} - \mathbf{V}_h^{(l+1)}\|_\infty$. Applying this relation recursively with a little algebra leads to

$$\|\widehat{\mathbf{V}}_h^{(0)} - \mathbf{V}_h^{(0)}\|_\infty \leq \left(\sqrt{\frac{\beta_1}{N}} \right)^m \|\widehat{\mathbf{V}}_h^{(m)} - \mathbf{V}_h^{(m)}\|_\infty$$

$$+ \sum_{l=1}^m \left(\sqrt{\frac{\beta_1}{N}} \right)^l \|\mathbf{V}_h^{(l)}\|_\infty + \frac{\beta_1 H}{N} \sum_{l=1}^m \left(\sqrt{\frac{\beta_1}{N}} \right)^{l-1} \max_j \|\mathbf{V}_j^{(l-1)}\|_\infty. \quad (116)$$

Additionally, it is easily seen from the definition (62) that

$$\mathbf{r}_h^{(l+1)} \leq \mathbf{V}_{h+1}^{(l)} \leq \|\mathbf{V}_{h+1}^{(l)}\|_\infty \mathbf{1},$$

which taken together with (113) and the elementary identity $\prod_{i=h}^{j-1} \hat{\mathbf{P}}_{i,\pi} \mathbf{1} = \mathbf{1}$ implies that

$$\begin{aligned} |\hat{\mathbf{V}}_h^{(m)} - \mathbf{V}_h^{(m)}| &\leq \sqrt{\frac{\beta_1}{N}} \max_j \|\mathbf{r}_j^{(m+1)}\|_\infty \left(\sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \hat{\mathbf{P}}_{i,\pi} \mathbf{1} \right) + \frac{\beta_1 H \max_j \|\mathbf{V}_j^{(m)}\|_\infty}{N} \mathbf{1} \\ &\leq \sqrt{\frac{\beta_1}{N}} \max_j \|\mathbf{V}_j^{(m)}\|_\infty \left(\sum_{j=h}^{H-1} \mathbf{1} \right) + \frac{\beta_1 H \max_j \|\mathbf{V}_j^{(m)}\|_\infty}{N} \mathbf{1} \\ &\leq H \left(\sqrt{\frac{\beta_1}{N}} + \frac{\beta_1}{N} \right) \max_j \|\mathbf{V}_j^{(m)}\|_\infty \mathbf{1}. \end{aligned}$$

Substitution into (116) results in

$$\begin{aligned} \|\hat{\mathbf{V}}_h^{(0)} - \mathbf{V}_h^{(0)}\|_\infty &\leq H \left(\sqrt{\frac{\beta_1}{N}} \right)^m \left(\sqrt{\frac{\beta_1}{N}} + \frac{\beta_1}{N} \right) \max_j \|\mathbf{V}_j^{(m)}\|_\infty \\ &\quad + \sum_{l=1}^m \left(\sqrt{\frac{\beta_1}{N}} \right)^l \|\mathbf{V}_h^{(l)}\|_\infty + \frac{\beta_1 H}{N} \sum_{l=1}^m \left(\sqrt{\frac{\beta_1}{N}} \right)^{l-1} \max_j \|\mathbf{V}_j^{(l-1)}\|_\infty. \end{aligned} \quad (117)$$

To finish up, it remains to control the terms $\{\|\mathbf{V}_h^{(l)}\|_\infty\}$. Towards this, combining Lemma 12 with the above inequality (117) yields

$$\begin{aligned} \|\hat{\mathbf{V}}_h^{(0)} - \mathbf{V}_h^{(0)}\|_\infty &\leq H^2 \left(\sqrt{\frac{3\beta_1 H}{N}} \right)^m \left(\sqrt{\frac{\beta_1}{N}} + \frac{\beta_1}{N} \right) + \sum_{l=1}^m H \left(\sqrt{\frac{3\beta_1 H}{N}} \right)^l + \sum_{l=1}^m \frac{\beta_1 H^2}{N} \left(\sqrt{\frac{3\beta_1 H}{N}} \right)^{l-1} \\ &= H^2 \left(\sqrt{\frac{3\beta_1 H}{N}} \right)^m \left(\sqrt{\frac{\beta_1}{N}} + \frac{\beta_1}{N} \right) + \left\{ H \sqrt{\frac{3\beta_1 H}{N}} + \frac{\beta_1 H^2}{N} \right\} \sum_{l=1}^m \left(\sqrt{\frac{3\beta_1 H}{N}} \right)^{l-1} \\ &\leq H^2 \left(\sqrt{\frac{3\beta_1 H}{N}} \right)^m \left(\sqrt{\frac{\beta_1}{N}} + \frac{\beta_1}{N} \right) + 2 \left\{ H \sqrt{\frac{3\beta_1 H}{N}} + \frac{\beta_1 H^2}{N} \right\}, \end{aligned} \quad (118)$$

Here, the last inequality holds true since

$$\sum_{l=1}^m \left(\sqrt{\frac{3\beta_1 H}{N}} \right)^{l-1} \leq \frac{1}{1 - \sqrt{\frac{3\beta_1 H}{N}}} \leq 2,$$

provided that $\sqrt{\frac{3\beta_1 H}{N}} \leq 1/2$. Invoking the assumption $N \geq 12H\beta_1$ again and taking $m = \log_2 H$, we have $\left(\sqrt{\frac{3\beta_1 H}{N}} \right)^m \leq 1/H$. This combined with (119) immediately leads to

$$\|\hat{\mathbf{V}}_h^{(0)} - \mathbf{V}_h^{(0)}\|_\infty \leq 3 \left\{ H \sqrt{\frac{3\beta_1 H}{N}} + \frac{\beta_1 H^2}{N} \right\} \leq 6H \sqrt{\frac{3\beta_1 H}{N}}, \quad (119)$$

with the proviso that $N \geq 3\beta_1 H$. This concludes the proof.

C.2 Proof of Lemma 12

To begin with, it is seen from the notation (28) that for all $1 \leq j \leq H$,

$$\begin{aligned}
\text{Var}_{\mathbf{P}_{j,\pi}}(\mathbf{V}_{j+1}^{(l)}) &= \mathbf{P}_{j,\pi}(\mathbf{V}_{j+1}^{(l)} \circ \mathbf{V}_{j+1}^{(l)}) - (\mathbf{P}_{j,\pi} \mathbf{V}_{j+1}^{(l)}) \circ (\mathbf{P}_{j,\pi} \mathbf{V}_{j+1}^{(l)}) \\
&= \mathbf{P}_{j,\pi}(\mathbf{V}_{j+1}^{(l)} \circ \mathbf{V}_{j+1}^{(l)}) - (\mathbf{V}_j^{(l)} - \mathbf{r}_j^{(l)}) \circ (\mathbf{V}_j^{(l)} - \mathbf{r}_j^{(l)}) \\
&= \mathbf{P}_{j,\pi}(\mathbf{V}_{j+1}^{(l)} \circ \mathbf{V}_{j+1}^{(l)}) - \mathbf{V}_j^{(l)} \circ \mathbf{V}_j^{(l)} + 2\mathbf{V}_j^{(l)} \circ \mathbf{r}_j^{(l)} - \mathbf{r}_j^{(l)} \circ \mathbf{r}_j^{(l)} \\
&\leq \mathbf{P}_{j,\pi}(\mathbf{V}_{j+1}^{(l)} \circ \mathbf{V}_{j+1}^{(l)}) - \mathbf{V}_j^{(l)} \circ \mathbf{V}_j^{(l)} + 2\mathbf{V}_j^{(l)} \circ \mathbf{r}_j^{(l)},
\end{aligned} \tag{120}$$

where the second identity makes use of the fact that $\mathbf{V}_j^{(l)} = \mathbf{r}_j^{(l)} + \mathbf{P}_{j,\pi} \mathbf{V}_{j+1}^{(l)}$ (cf. (62)). Moreover, from the construction (62) we can easily derive

$$\mathbf{V}_h^{(l)} = \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi} \mathbf{r}_j^{(l)} = \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi} \sqrt{\text{Var}_{\mathbf{P}_{j,\pi}}[\mathbf{V}_{j+1}^{(l-1)}]}. \tag{121}$$

The above two results taken collectively give

$$\begin{aligned}
|\mathbf{V}_h^{(l+1)}| &= \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi} \sqrt{\text{Var}_{\mathbf{P}_{j,\pi}}[\mathbf{V}_{j+1}^{(l)}]} \stackrel{(i)}{\leq} \sum_{j=h}^{H-1} \sqrt{\prod_{i=h}^{j-1} \mathbf{P}_{i,\pi} \text{Var}_{\mathbf{P}_{j,\pi}}[\mathbf{V}_{j+1}^{(l)}]} \\
&\stackrel{(ii)}{\leq} \sqrt{H \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi} \text{Var}_{\mathbf{P}_{j,\pi}}[\mathbf{V}_{j+1}^{(l)}]} \\
&\stackrel{(iii)}{\leq} \sqrt{H \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi} \left[\mathbf{P}_{j,\pi}(\mathbf{V}_{j+1}^{(l)} \circ \mathbf{V}_{j+1}^{(l)}) - \mathbf{V}_j^{(l)} \circ \mathbf{V}_j^{(l)} + 2\mathbf{V}_j^{(l)} \circ \mathbf{r}_j^{(l)} \right]} \\
&= \sqrt{H \sum_{j=h}^{H-1} \left(\prod_{i=h}^j \mathbf{P}_{i,\pi}(\mathbf{V}_{j+1}^{(l)} \circ \mathbf{V}_{j+1}^{(l)}) - \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi}(\mathbf{V}_j^{(l)} \circ \mathbf{V}_j^{(l)}) \right) + 2H \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi}(\mathbf{V}_j^{(l)} \circ \mathbf{r}_j^{(l)})}.
\end{aligned}$$

Here, (i) arises from Jensen's inequality; (ii) holds true due to the Cauchy-Schwarz inequality; (iii) follows from (120). By telescoping summation, one further arrives at

$$\begin{aligned}
|\mathbf{V}_h^{(l+1)}| &\leq \sqrt{H \left[\prod_{i=h}^{H-1} \mathbf{P}_{i,\pi}(\mathbf{V}_H^{(l)} \circ \mathbf{V}_H^{(l)}) - \mathbf{V}_h^{(l)} \circ \mathbf{V}_h^{(l)} \right] + 2H \max_j \|\mathbf{V}_j^{(l)}\|_\infty \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi} \mathbf{r}_j^{(l)}} \\
&\stackrel{(iv)}{=} \sqrt{H \left[\prod_{i=h}^{H-1} \mathbf{P}_{i,\pi}(\mathbf{V}_H^{(l)} \circ \mathbf{V}_H^{(l)}) - \mathbf{V}_h^{(l)} \circ \mathbf{V}_h^{(l)} \right] + 2H \max_j \|\mathbf{V}_j^{(l)}\|_\infty \mathbf{V}_h^{(l)}} \\
&\leq \sqrt{H \prod_{i=h}^{H-1} \mathbf{P}_{i,\pi}(\mathbf{V}_H^{(l)} \circ \mathbf{V}_H^{(l)}) + 2H \left(\max_j \|\mathbf{V}_j^{(l)}\|_\infty \right)^2 \mathbf{1}} \\
&\stackrel{(v)}{\leq} \sqrt{3H} \max_j \|\mathbf{V}_j^{(l)}\|_\infty \mathbf{1}.
\end{aligned}$$

Here (iv) invokes the relation $\mathbf{V}_h^{(l)} = \sum_{j=h}^{H-1} \prod_{i=h}^{j-1} \mathbf{P}_{i,\pi} \mathbf{r}_j^{(l)}$ (see (121)); and (v) holds true since

$$\left\| \prod_{i=h}^{H-1} \mathbf{P}_{i,\pi}(\mathbf{V}_H^{(l)} \circ \mathbf{V}_H^{(l)}) \right\|_\infty \leq \left\| \prod_{i=h}^{H-1} \mathbf{P}_{i,\pi} \right\|_1 \left\| (\mathbf{V}_H^{(l)} \circ \mathbf{V}_H^{(l)}) \right\|_\infty \leq \max_j \|\mathbf{V}_j^{(l)}\|_\infty^2.$$

As a consequence, the above inequality allows one to deduce that

$$\max_j \|\mathbf{V}_j^{(l+1)}\|_\infty \leq \sqrt{3H} \max_j \|\mathbf{V}_j^{(l)}\|_\infty,$$

and therefore,

$$\max_j \|\mathbf{V}_j^{(l)}\|_\infty \leq (\sqrt{3H})^l \max_j \|\mathbf{V}_j^{(0)}\|_\infty \leq (\sqrt{3H})^l H,$$

where the last inequality arises from the trivial upper bound $\max_j \|\mathbf{V}_j^{(0)}\|_\infty \leq H$.

References

- Agarwal, A., Kakade, S., and Yang, L. F. (2020). Model-based reinforcement learning with a generative model is minimax optimal. *Conference on Learning Theory (COLT)*.
- Azar, M. G., Munos, R., and Kappen, B. (2012). On the sample complexity of reinforcement learning with a generative model. *arXiv preprint arXiv:1206.6461*.
- 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.
- Azar, M. G., Osband, I., and Munos, R. (2017). Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 263–272. JMLR.org.
- Beck, C. L. and Srikant, R. (2012). Error bounds for constant step-size Q-learning. *Systems & control letters*, 61(12):1203–1208.
- Bellman, R. (1952). On the theory of dynamic programming. *Proceedings of the National Academy of Sciences of the United States of America*, 38(8):716.
- Bertsekas, D. P. (2017). *Dynamic programming and optimal control (4th edition)*. Athena Scientific.
- Bhandari, J., Russo, D., and Singal, R. (2018). A finite time analysis of temporal difference learning with linear function approximation. In *Conference On Learning Theory*, pages 1691–1692.
- Bradtke, S. J. and Barto, A. G. (1996). Linear least-squares algorithms for temporal difference learning. *Machine learning*, 22(1-3):33–57.
- Cai, Q., Yang, Z., Lee, J. D., and Wang, Z. (2019). Neural temporal-difference learning converges to global optima. In *Advances in Neural Information Processing Systems*, pages 11312–11322.
- Chen, Y., Chi, Y., Fan, J., and Ma, C. (2021). Spectral methods for data science: A statistical perspective. *Foundations and Trends® in Machine Learning*, 14(5):566–806.
- Chen, Y., Fan, J., Ma, C., and Wang, K. (2019). Spectral method and regularized mle are both optimal for top-k ranking. *Annals of statistics*, 47(4):2204.
- Chen, Z., Maguluri, S. T., Shakkottai, S., and Shanmugam, K. (2020). Finite-sample analysis of stochastic approximation using smooth convex envelopes. *arXiv preprint arXiv:2002.00874*.
- Dalal, G., Szörényi, B., Thoppe, G., and Mannor, S. (2018). Finite sample analyses for TD(0) with function approximation. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- 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.
- El Karoui, N. (2015). On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, pages 1–81.
- Even-Dar, E. and Mansour, Y. (2003). Learning rates for Q-learning. *Journal of machine learning Research*, 5(Dec):1–25.

- Fan, J., Wang, Z., Xie, Y., and Yang, Z. (2019). A theoretical analysis of deep Q-learning. *arXiv preprint arXiv:1901.00137*.
- Gupta, H., Srikant, R., and Ying, L. (2019). Finite-time performance bounds and adaptive learning rate selection for two time-scale reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 4706–4715.
- Jaakkola, T., Jordan, M. I., and Singh, S. P. (1994). Convergence of stochastic iterative dynamic programming algorithms. In *Advances in neural information processing systems*, pages 703–710.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. (2018). Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873.
- Kakade, S. (2003). *On the sample complexity of reinforcement learning*. PhD thesis, University of London.
- Kaledin, M., Moulines, E., Naumov, A., Tadic, V., and Wai, H.-T. (2020). Finite time analysis of linear two-timescale stochastic approximation with Markovian noise. *arXiv preprint arXiv:2002.01268*.
- Kearns, M., Mansour, Y., and Ng, A. Y. (2002). A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Machine learning*, 49(2-3):193–208.
- Kearns, M. J. and Singh, S. P. (1999). Finite-sample convergence rates for Q-learning and indirect algorithms. In *Advances in neural information processing systems*, pages 996–1002.
- Khamaru, K., Pananjady, A., Ruan, F., Wainwright, M. J., and Jordan, M. I. (2020). Is temporal difference learning optimal? an instance-dependent analysis. *arXiv preprint arXiv:2003.07337*.
- Lakshminarayanan, C. and Szepesvari, C. (2018). Linear stochastic approximation: How far does constant step-size and iterate averaging go? In *International Conference on Artificial Intelligence and Statistics*, pages 1347–1355.
- Lattimore, T. and Hutter, M. (2012). PAC bounds for discounted MDPs. In *International Conference on Algorithmic Learning Theory*, pages 320–334. Springer.
- Li, G., Cai, C., Chen, Y., Gu, Y., Wei, Y., and Chi, Y. (2021a). Is Q-learning minimax optimal? a tight sample complexity analysis. *arXiv preprint arXiv:2102.06548*.
- Li, G., Cai, C., Chen, Y., Gu, Y., Wei, Y., and Chi, Y. (2021b). Tightening the dependence on horizon in the sample complexity of Q-learning. In *International Conference on Machine Learning*, pages 6296–6306.
- Li, G., Shi, L., Chen, Y., Gu, Y., and Chi, Y. (2021c). Breaking the sample complexity barrier to regret-optimal model-free reinforcement learning. *Advances in Neural Information Processing Systems*, 34.
- Li, G., Wei, Y., Chi, Y., Gu, Y., and Chen, Y. (2021d). Sample complexity of asynchronous Q-learning: Sharper analysis and variance reduction. *IEEE Transactions on Information Theory*.
- Ma, C., Wang, K., Chi, Y., and Chen, Y. (2020). Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. *Foundations of Computational Mathematics*, 20(3):451–632.
- Mou, W., Li, C. J., Wainwright, M. J., Bartlett, P. L., and Jordan, M. I. (2020). On linear stochastic approximation: Fine-grained Polyak-Ruppert and non-asymptotic concentration. *arXiv preprint arXiv:2004.04719*.
- Pananjady, A. and Wainwright, M. J. (2019). Value function estimation in Markov reward processes: Instance-dependent ℓ_∞ -bounds for policy evaluation. *arXiv preprint arXiv:1909.08749*.
- Qu, G. and Wierman, A. (2020). Finite-time analysis of asynchronous stochastic approximation and Q-learning. In *Conference on Learning Theory*, pages 3185–3205.

- Shah, D. and Xie, Q. (2018). Q-learning with nearest neighbors. In *Advances in Neural Information Processing Systems*, pages 3111–3121.
- Sidford, A., Wang, M., Wu, X., Yang, L., and Ye, Y. (2018a). Near-optimal time and sample complexities for solving Markov decision processes with a generative model. In *Advances in Neural Information Processing Systems*, pages 5186–5196.
- Sidford, A., Wang, M., Wu, X., and Ye, Y. (2018b). Variance reduced value iteration and faster algorithms for solving Markov decision processes. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 770–787.
- Srikant, R. and Ying, L. (2019). Finite-time error bounds for linear stochastic approximation and TD learning. In *Conference on Learning Theory*, pages 2803–2830.
- Strehl, A. L., Li, L., Wiewiora, E., Langford, J., and Littman, M. L. (2006). PAC model-free reinforcement learning. In *International Conference on Machine Learning*, pages 881–888.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Szepesvári, C. (1998). The asymptotic convergence-rate of Q-learning. In *Advances in Neural Information Processing Systems*, pages 1064–1070.
- Szepesvári, C. (2010). Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103.
- Tsitsiklis, J. and Van Roy, B. (1997). An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control*, 42(5):674–690.
- Tsitsiklis, J. N. (1994). Asynchronous stochastic approximation and Q-learning. *Machine learning*, 16(3):185–202.
- 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.
- Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.
- Wainwright, M. J. (2019a). Stochastic approximation with cone-contractive operators: Sharp ℓ_∞ -bounds for Q-learning. *arXiv preprint arXiv:1905.06265*.
- Wainwright, M. J. (2019b). Variance-reduced Q-learning is minimax optimal. *arXiv preprint arXiv:1906.04697*.
- Wang, B., Yan, Y., and Fan, J. (2021). Sample-efficient reinforcement learning for linearly-parameterized MDPs with a generative model. *Neural Information Processing Systems*.
- Wang, M. (2019). Randomized linear programming solves the Markov decision problem in nearly linear (sometimes sublinear) time. *Mathematics of Operations Research*.
- Xu, P. and Gu, Q. (2020). A finite-time analysis of Q-learning with neural network function approximation. In *International Conference on Machine Learning*, pages 10555–10565.
- Xu, T., Zou, S., and Liang, Y. (2019). Two time-scale off-policy TD learning: Non-asymptotic analysis over Markovian samples. In *Advances in Neural Information Processing Systems*, pages 10633–10643.
- Yang, L. and Wang, M. (2019). Sample-optimal parametric Q-learning using linearly additive features. In *International Conference on Machine Learning*, pages 6995–7004.
- Yin, M., Bai, Y., and Wang, Y.-X. (2021). Near-optimal provable uniform convergence in offline policy evaluation for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1567–1575.