

The Plug-in Approach for Average-Reward and Discounted MDPs: Optimal Sample Complexity Analysis

Matthew Zurek and Yudong Chen

Department of Computer Sciences, University of Wisconsin-Madison
`{matthew.zurek,yudong.chen}@wisc.edu`

Abstract

We study the sample complexity of the plug-in approach for learning ε -optimal policies in average-reward Markov decision processes (MDPs) with a generative model. The plug-in approach constructs a model estimate then computes an average-reward optimal policy in the estimated model. Despite representing arguably the simplest algorithm for this problem, the plug-in approach has never been theoretically analyzed. Unlike the more well-studied discounted MDP reduction method, the plug-in approach requires no prior problem information or parameter tuning. Our results fill this gap and address the limitations of prior approaches, as we show that the plug-in approach is optimal in several well-studied settings without using prior knowledge. Specifically it achieves the optimal diameter- and mixing-based sample complexities of $\tilde{O}(SA\frac{D}{\varepsilon^2})$ and $\tilde{O}(SA\frac{\tau_{\text{unif}}}{\varepsilon^2})$, respectively, without knowledge of the diameter D or uniform mixing time τ_{unif} . We also obtain span-based bounds for the plug-in approach, and complement them with algorithm-specific lower bounds suggesting that they are unimprovable. Our results require novel techniques for analyzing long-horizon problems which may be broadly useful and which also improve results for the discounted plug-in approach, removing effective-horizon-related sample size restrictions and obtaining the first optimal complexity bounds for the full range of sample sizes without reward perturbation.

1 Introduction

Reinforcement learning (RL) has emerged as a powerful framework for sequential decision-making problems, where an agent learns to make decisions by interacting with an environment to maximize cumulative rewards. Average reward RL, in particular, focuses on optimizing the long-term average reward per time step, making it especially relevant in ongoing, infinite-horizon tasks where the goal is to maintain consistent performance over time. In this paper, we study the foundational theoretical problem of determining the sample complexity required to learn a near-optimal policy in a Markov decision process (MDP) with access to a generative model. Although recent research has made significant strides in resolving the optimal sample complexity for this problem, a large amount of prior work (including all sample-optimal methods) relies on methods designed for discounted MDPs, where future rewards are multiplied by a discount factor to prioritize immediate rewards. This approach has several drawbacks: selecting the appropriate discount factor (or sequence of factors) is crucial and often requires prior knowledge about the problem, which may not be available in practice, potentially degrading performance. Even when the discount factor can be suitably tuned, it is still extrinsic to the average-reward problem, making it arguably unnatural to require its introduction. Technical challenges have hindered the analysis of more direct average-reward algorithms.

Our study focuses on analyzing the average-reward *plug-in* approach. This approach estimates the parameters of the MDP model and then uses any method to compute the optimal average-reward policy for the estimated model. In the context of discounted MDPs, this approach has been called model-based planning [Agarwal et al., 2020, Li et al., 2020], although we note that the plug-in approach is a particular “model-based” algorithm. We also note that the plug-in approach is a generic template for constructing estimators for a functional of an unknown distribution/model (by plugging the empirical distribution/model into the functional) which is broadly used beyond RL. This is arguably the most natural model-based approach for

Algorithm	Sample Complexity	Reference	Prior Knowledge
DMDP Reduction	$SA \frac{\tau_{\text{unif}}}{\varepsilon^2}$	Wang et al. [2023b]	Yes
DMDP Reduction	$SA \frac{\ h^*\ _{\text{span}}+1}{\varepsilon^2}$	Zurek and Chen [2024]	Yes
Diameter Estimation + DMDP Reduction	$SA \frac{D}{\varepsilon^2} + S^2 AD^2$	Tuynman et al. [2024]	No
Dynamic Horizon Q-Learning	$SA \frac{\tau_{\text{unif}}^8}{\varepsilon^8}$	Jin et al. [2024]	No
Stochastic Saddle-Point Optimization	$S^2 A^2 \frac{\ h^{\hat{\pi}}\ _{\text{span}}^4}{\varepsilon^2}$	Neu and Okolo [2024]	No
Plug-in Approach	$SA \frac{\ h^*\ _{\text{span}} + \ \hat{h}^*\ _{\text{span}} + 1}{\varepsilon^2} \log(\frac{1}{1-\gamma^*})$	Our Theorem 1	No
Anchored+Perturbed Plug-in Approach	$SA \frac{\ h^*\ _{\text{span}} + \min\{\ \hat{h}^*\ _{\text{span}}, \ \underline{h}^{\hat{\pi}}\ _{\text{span}}\} + 1}{\varepsilon^2}$ $SA \frac{D}{\varepsilon^2}$ $SA \frac{\tau_{\text{unif}}}{\varepsilon^2}$	Our Theorem 8 Our Corollary 5 Our Corollary 7	No
$\sqrt{n}$ -Horizon DMDP Reduction	$SA \frac{\ h^*\ _{\text{span}}^2 + 1}{\varepsilon^2}$	Our Theorem 13	No

Table 1: **Algorithms and sample complexity bounds for average reward MDPs** with S states and A actions, for finding an ε -optimal policy under a generative model (up to log factors). All results assume at least that P is weakly communicating. Furthermore all bounds involving τ_{unif} assume that P is uniformly mixing, and all bounds involving D assume that P is communicating. $\|h^{\hat{\pi}}\|_{\text{span}}$ is the bias span of the policy $\hat{\pi}$ returned by the algorithm. See Section 3 for other definitions. We always have $\|h^*\|_{\text{span}} \leq D$ and $\|h^*\|_{\text{span}}, \|\hat{h}^{\hat{\pi}}\|_{\text{span}}, \|\underline{h}^{\hat{\pi}}\|_{\text{span}} \leq 3\tau_{\text{unif}}$, and we have $\|\hat{h}^*\|_{\text{span}} \leq O(D)$ with high probability when $n \geq \tilde{\Omega}(D)$.

solving average-reward MDPs, yet its finite-sample properties have never been theoretically examined. Not only does our analysis fill a major gap in our understanding of a basic algorithm for this problem, but also we show that this algorithm is optimal in several settings without requiring strong assumptions about prior knowledge of the problem, thus addressing many limitations of previous approaches.

In particular, when combined with the stabilizing *anchoring* technique which has previously appeared in the average-reward literature, we show that this algorithm can simultaneously achieve the optimal diameter- and mixing-based sample complexities of $\tilde{O}(SA \frac{D}{\varepsilon^2})$ and $\tilde{O}(SA \frac{\tau_{\text{unif}}}{\varepsilon^2})$ for learning an ε -optimal policy, respectively, where D is the diameter and τ_{unif} is the uniform mixing time, without needing to have prior knowledge of D or τ_{unif} and without needing to tailor the algorithm to the particular situations. These results are corollaries of our bias-span-based complexity bounds for weakly communicating MDPs, for example $\tilde{O}(SA \frac{\|h^*\|_{\text{span}} + \|\hat{h}^*\|_{\text{span}} + 1}{\varepsilon^2})$, where $\|h^*\|_{\text{span}}$ is the optimal bias span and $\|\hat{h}^*\|_{\text{span}}$ is the optimal bias span in a certain estimated MDP. We further show that the analysis behind this (and related) span-based bound is unimprovable, in the sense that the term $\|\hat{h}^*\|_{\text{span}}$ cannot be removed in general for the performance of the plug-in method.

While the average-reward plug-in approach can be seen as a large-discount-factor limit of the discounted plug-in approach, previous discounted analyses are incapable of being adapted to this problem, requiring the development of novel techniques for analyzing the error of long-horizon problems which may be broadly useful. In particular these techniques lead to improved results for the discounted plug-in approach, including removing effective-horizon-related sample size restrictions of previous results which achieve quadratic dependence on the effective horizon for the fixed MDP setting. We also obtain the first optimal complexity bounds for the full range of sample sizes without the need for reward perturbation.

Reference	Sample Complexity	Sample Size Requirement	Requires Perturbation?
Azar et al. [2013]	$SA \frac{1}{(1-\gamma)^3 \varepsilon^2}$	$\varepsilon \leq \sqrt{\frac{1}{(1-\gamma)S}}$	No
Agarwal et al. [2020]	$SA \frac{1}{(1-\gamma)^3 \varepsilon^2}$	$\varepsilon \leq \sqrt{\frac{1}{1-\gamma}}$	No
Li et al. [2020]	$SA \frac{1}{(1-\gamma)^3 \varepsilon^2}$	None	Yes
Wang et al. [2023a]	$SA \frac{\tau_{\text{unif}}}{(1-\gamma)^2 \varepsilon^2}$	$\varepsilon \leq \sqrt{\frac{\tau_{\text{unif}}}{1-\gamma}}$	Yes
Zurek and Chen [2024]	$SA \frac{\ h^*\ _{\text{span}} + 1}{(1-\gamma)^2 \varepsilon^2}$	$\varepsilon \leq \ h^*\ _{\text{span}}$	Yes
Our Theorem 10	$SA \frac{\ V^*\ _{\text{span}} + \ V^{\hat{\pi}}\ _{\text{span}} + 1}{(1-\gamma)^2 \varepsilon^2}$	None	Yes
Our Theorem 9	$SA \frac{\ V^*\ _{\text{span}} + \ \hat{V}^*\ _{\text{span}} + 1}{(1-\gamma)^2 \varepsilon^2}$	None	No
Our Theorem 9+Lemma 12	$SA \frac{\ h^*\ _{\text{span}} + 1}{(1-\gamma)^2 \varepsilon^2}$	$\varepsilon \leq \ h^*\ _{\text{span}}$	No

Table 2: **Sample complexity bounds for the plug-in approach in γ -discounted MDPs** with S states and A actions, for finding an ε -optimal policy under a generative model (up to log factors). The sample size requirement is the valid range of ε for the respective complexity guarantee, and a result is said to require perturbation if it utilizes a randomly perturbed reward vector. All results containing $\|h^*\|_{\text{span}}$ assume P is weakly communicating, and all results containing τ_{unif} assume uniform mixing. If P is weakly communicating then $\|V^*\|_{\text{span}} \leq 2\|h^*\|_{\text{span}}$, and if P is uniformly mixing then $\|V^*\|_{\text{span}}, \|V^{\hat{\pi}}\|_{\text{span}} \leq 3\tau_{\text{unif}}$. We also always have the naive bounds $\|V^*\|_{\text{span}}, \|V^{\hat{\pi}}\|_{\text{span}} \leq \frac{1}{1-\gamma}$.

1.1 Related Work

We summarize related work on learning optimal policies in average-reward MDPs (AMDPs) in Table 1. There is a long history of work on this problem which we do not fully recount here (e.g. Jin and Sidford [2020, 2021], Li et al. [2022], Wang et al. [2022], Zhang and Xie [2023]), instead starting with the works Wang et al. [2023b] and Zurek and Chen [2024] which were the first to obtain optimal sample complexities in their respective settings (we refer to their references for more history of this problem). Each of these works use the DMDP reduction approach with a carefully chosen effective horizon, $\frac{\tau_{\text{unif}}}{\varepsilon}$ and $\frac{\|h^*\|_{\text{span}}}{\varepsilon}$, respectively, which requires prior knowledge of the values of these complexity parameters. The $\tilde{O}(SA \frac{\tau_{\text{unif}}}{\varepsilon^2})$ complexity result of Zurek and Chen [2024] implies a $\tilde{O}(SA \frac{D}{\varepsilon^2})$ complexity for the finite diameter setting since $\|h^*\|_{\text{span}} \leq D$ [Bartlett and Tewari, 2012, Lattimore and Szepesvári, 2020], and it also implies the $\tilde{O}(SA \frac{\tau_{\text{unif}}}{\varepsilon^2})$ complexity obtained by Wang et al. [2023b] since $\|h^*\|_{\text{span}} \leq 3\tau_{\text{unif}}$ (Lemma 39, also see Wang et al. [2022]). These results match minimax lower bounds of $\tilde{\Omega}(SA \frac{\tau_{\text{unif}}}{\varepsilon^2})$ [Jin and Sidford, 2021] and $\tilde{\Omega}(SA \frac{D}{\varepsilon^2})$ (by the relationships between τ_{unif} , D , and $\|h^*\|_{\text{span}}$, these both imply a $\tilde{\Omega}(SA \frac{\|h^*\|_{\text{span}}}{\varepsilon^2})$ lower bound).

Recently there has been significant interest in removing the need for prior knowledge of complexity parameters. Tuynman et al. [2024] show that an upper bound for the diameter can be estimated and then used within the approach of Zurek and Chen [2024] to circumvent its need for parameter knowledge. However, results from Tuynman et al. [2024] and Zurek and Chen [2024] imply that a similar approach cannot be used to obtain the optimal span-based complexity, both showing that it is not generally possible to obtain a multiplicative approximation of $\|h^*\|_{\text{span}}$ with $\text{poly}(SA\|h^*\|_{\text{span}})$ samples. In the uniformly mixing setting, Jin et al. [2024] use a Q-learning-style algorithm with increasing discount factors to remove the need for knowledge of τ_{unif} . Neu and Okolo [2024] develop an algorithm based on stochastic saddle-point optimization that does not require parameter knowledge in the general weakly communicating setting, but their bounds depend on $\|h^{\hat{\pi}}\|_{\text{span}}$, the bias of the algorithm output policy, which is not generally related to $\|h^*\|_{\text{span}}$.

We present related work on learning optimal policies in discounted MDPs (DMDPs) in Table 2. This problem also has been extensively studied, and we include only results on the plug-in approach which obtain minimax-optimal sample complexities. The optimal complexity $\tilde{O}(SA \frac{1}{(1-\gamma)^3 \varepsilon^2})$ for learning a ε -discounted-

optimal policy was first obtained by Azar et al. [2013] for a restrictive range of ε . This range was enlarged by Agarwal et al. [2020], who introduce the absorbing MDP construction for decoupling statistical dependence which also finds use in our analysis. The matching lower bound is established by Azar et al. [2012], Sidford et al. [2018]. Li et al. [2020] are the first to achieve the optimal $\tilde{O}(SA \frac{1}{(1-\gamma)^3 \varepsilon^2})$ complexity for the full range $\varepsilon \in (0, \frac{1}{1-\gamma}]$, and their results actually yield a stronger instance-dependent bound. This stronger guarantee is used in both Wang et al. [2023b] and Zurek and Chen [2024] to obtain the complexity bounds of $\tilde{O}(SA \frac{\tau_{\text{unif}}}{(1-\gamma)^2 \varepsilon^2})$ and $\tilde{O}(SA \frac{\|h^*\|_{\text{span}}}{(1-\gamma)^2 \varepsilon^2})$, respectively, in the restricted situations that P is uniformly mixing or is weakly communicating. However, the arguments within Li et al. [2020] implicitly require that $n \geq \Omega(\frac{1}{1-\gamma})$. This is without loss of generality if the goal is to show $\tilde{O}(SA \frac{1}{(1-\gamma)^3 \varepsilon^2})$ complexity, since this is equivalent to an error bound of $\tilde{O}(\sqrt{\frac{1}{(1-\gamma)^3 n}})$, which is only nontrivial error if it is below $\frac{1}{1-\gamma}$ which requires $n \geq \Omega(\frac{1}{1-\gamma})$. However, once the target is strengthened to an improved complexity like $\tilde{O}(SA \frac{\tau_{\text{unif}}}{(1-\gamma)^2 \varepsilon^2})$, the condition $n \geq \Omega(\frac{1}{1-\gamma})$ is equivalent to a sample size barrier of $\varepsilon \leq O(\sqrt{\frac{\tau_{\text{unif}}}{1-\gamma}})$.

2 Problem Setup

A Markov decision process (MDP) is a tuple $(\mathcal{S}, \mathcal{A}, P, r)$, where $\mathcal{S}$ is the finite set of states, $\mathcal{A}$ is the finite set of actions, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition kernel with $\Delta(\mathcal{S})$ denoting the probability simplex over $\mathcal{S}$, and $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the reward function. We denote the cardinality of the state and action spaces as $S = |\mathcal{S}|$ and $A = |\mathcal{A}|$, respectively. Unless otherwise noted, all policies considered are Markovian (stationary) policies of the form $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$. For any initial state $s_0 \in \mathcal{S}$ and policy π , we let $\mathbb{E}_{s_0}^\pi$ denote the expectation with respect to the probability distribution over trajectories $(S_0, A_0, S_1, A_1, \dots)$ where $S_0 = s_0$, $A_t \sim \pi(S_t)$, and $S_{t+1} \sim P(\cdot | S_t, A_t)$. We let P_π denote the transition probability matrix of the Markov chain induced by π , that is, $(P_\pi)_{s,s'} := \sum_{a \in \mathcal{A}} \pi(a|s) P(s' | s, a)$. Likewise define $(r_\pi)_s := \sum_{a \in \mathcal{A}} \pi(a|s) r(s, a)$. We also consider P as an $(\mathcal{S} \times \mathcal{A})$ -by- $\mathcal{S}$ matrix where $P_{sa,s'} = P(s' | s, a)$.

We assume access to a generative model [Kearns and Singh, 1998], also known as a simulator, which provides independent samples from $P(\cdot | s, a)$ for any given $s \in \mathcal{S}, a \in \mathcal{A}$. P itself is unknown. We assume the r is deterministic and known, which is standard in generative settings (e.g., Agarwal et al. 2020, Li et al. 2020) since otherwise estimating the mean rewards is relatively easy.

Discounted reward criterion A discounted MDP is a tuple $(\mathcal{S}, \mathcal{A}, P, r, \gamma)$, where $\gamma \in (0, 1)$ is the discount factor. For a stationary policy π , the (discounted) value function $V_\gamma^\pi : \mathcal{S} \rightarrow [0, \infty)$ is defined, for each $s \in \mathcal{S}$, as $V_\gamma^\pi(s) := \mathbb{E}_s^\pi [\sum_{t=0}^\infty \gamma^t R_t]$, where $R_t = r(S_t, A_t)$ is the reward received at time t . There always exists an optimal policy π_γ^* that is deterministic and satisfies $V_\gamma^{\pi_\gamma^*}(s) = V_\gamma^*(s) := \sup_\pi V_\gamma^\pi(s)$ for all $s \in \mathcal{S}$ [Puterman, 1994].

Average-reward criterion In an MDP $(\mathcal{S}, \mathcal{A}, P, r)$, the average reward per stage or the *gain* of a policy π starting from state s is defined as $\rho^\pi(s) := \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_s^\pi [\sum_{t=0}^{T-1} R_t]$. The *bias function* of any stationary policy π is $h^\pi(s) := \text{C-lim}_{T \rightarrow \infty} \mathbb{E}_s^\pi [\sum_{t=0}^{T-1} (R_t - \rho^\pi(S_t))]$, where C-lim denotes the Cesaro limit. When the Markov chain induced by P_π is aperiodic, C-lim can be replaced with the usual limit. For any policy π , ρ^π and h^π satisfy $\rho^\pi = P_\pi \rho^\pi$ and $\rho^\pi + h^\pi = r_\pi + P_\pi h^\pi$. We let $P_\pi^\infty = \text{C-lim}_{T \rightarrow \infty} P_\pi^T$ denote the limiting matrix, and note that $P_\pi^\infty P_\pi = P_\pi P_\pi^\infty = P_\pi^\infty$ and $\rho^\pi = P_\pi^\infty r_\pi$. A policy π^* is *Blackwell-optimal* if there exists some discount factor $\bar{\gamma} \in [0, 1)$ such that for all $\gamma \geq \bar{\gamma}$ we have $V_\gamma^{\pi^*} \geq V_\gamma^\pi$ for all policies π . When $\mathcal{S}$ and $\mathcal{A}$ are finite, there always exists some Blackwell-optimal policy which is Markovian and deterministic, which we denote π^* [Puterman, 1994]. We define the optimal gain $\rho^* \in \mathbb{R}^{\mathcal{S}}$ by $\rho^*(s) = \sup_\pi \rho^\pi(s)$ and note that we have $\rho^* = \rho^{\pi^*}$. We also define $h^* = h^{\pi^*}$ (and we note that this definition does not depend on which Blackwell-optimal π^* is used if there are multiple). A policy π is *gain-optimal* if $\rho^\pi = \rho^*$ and it is *bias-optimal* if it is gain-optimal and $h^\pi = h^*$. For $x \in \mathbb{R}^{\mathcal{S}}$, we define the span semi-norm $\|x\|_{\text{span}} := \max_{s \in \mathcal{S}} x(s) - \min_{s \in \mathcal{S}} x(s)$. An MDP is communicating if for any initial and target states, some Markovian policy can reach the target state from the initial state (with probability 1). The *diameter* is defined as $D := \max_{s_1 \neq s_2} \inf_{\pi \in \Pi} \mathbb{E}_{s_1}^\pi [\eta_{s_2}]$, where η_s denotes the hitting time of a state $s \in \mathcal{S}$. $D < \infty$ if and only if the MDP is communicating. A weakly communicating MDP is such that the states can be partitioned into two disjoint subsets $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$ such that all states in $\mathcal{S}_1$ are transient under all stationary policies and $\mathcal{S}_2$ is communicating. In weakly communicating MDPs ρ^* is

a constant vector (all entries are equal). For each policy π , if the Markov chain induced by P_π has a unique stationary distribution ν_π , we define the mixing time of π as $\tau_\pi := \inf \left\{ t \geq 1 : \max_{s \in \mathcal{S}} \|e_s^\top (P_\pi)^t - \nu_\pi^\top\|_1 \leq \frac{1}{2} \right\}$. If all policies in the set of Markovian *deterministic* policies, denoted Π , satisfy this assumption, we define the *uniform mixing time* $\tau_{\text{unif}} := \sup_{\pi \in \Pi} \tau_\pi$. An MDP is unichain if all Markovian deterministic policies induce a Markov chain P_π with a single recurrent class (and possibly some transient states). We note that this definition of mixing time requires the Markov chain P_π to be unichain but not irreducible.

When using transition kernels besides P , for example denoted $\hat{P}$, we will accordingly write $\hat{V}_\gamma^\pi, \hat{h}^\pi, \hat{\rho}^\pi$ for the associated value, bias, and gain functions respectively. We also occasionally drop the subscript from discounted value functions and write V^π when γ is clear from context.

We use standard Big-Oh notation O, Ω, Θ , and we also use the notation $\tilde{O}, \tilde{\Omega}$ to hide logarithmic factors in $n, S, A, \frac{1}{\delta}$, as well $\frac{1}{\xi}$ where ξ is a perturbation size parameter appearing in some results. We use $C_1, C_2, \dots$ to denote absolute constants.

3 Main Results

Algorithm 1 Plug-in approach for AMDP

input: Sample size per state-action pair n ; optional anchoring state s_0 , optional anchor probability η , optional perturbation level ξ

- 1: **for** each state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ **do**
- 2: Collect n samples $S_{s,a}^1, \dots, S_{s,a}^n$ from $P(\cdot \mid s, a)$
- 3: Form the empirical transition kernel $\hat{P}(s' \mid s, a) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}\{S_{s,a}^i = s'\}$, for all $s' \in \mathcal{S}$
- 4: **end for**
- 5: Form *anchored* empirical transition matrix $\underline{\hat{P}} = (1 - \eta)\hat{P} + \eta \mathbf{1} e_{s_0}^\top$

▷ Set $\eta = 0$ for no anchoring

- 6: Form *perturbed* reward $\tilde{r} = r + \Delta$ where $\Delta(s, a) \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}[0, \xi]$

▷ Set $\xi = 0$ for no perturbation

- 7: $\hat{\pi} = \text{SolveAMDP}(\underline{\hat{P}}, \tilde{r})$
- 8: **return** $\hat{\pi}$

We present a meta-algorithm, Algorithm 1, which encapsulates several variants of the plug-in approach for solving AMDPs. There are three key choices in Algorithm 1, within lines 5, 6, and 7. The first is that instead of solving for an optimal policy in the empirical MDP $\hat{P}$, we may instead choose to use the *anchored* MDP $\underline{\hat{P}} = (1 - \eta)\hat{P} + \eta \mathbf{1} e_{s_0}^\top$, which adds a small probability η of transitioning to an arbitrary *anchor* state s_0 from all states and actions [Fruit et al., 2018]. We discuss the anchoring technique in more detail shortly. This step is optional, and can be skipped by setting $\eta = 0$. Secondly, we may use a slightly perturbed reward vector $\tilde{r} = r + \Delta$ where $\Delta \in \mathbb{R}^{SA}$ has each entry sampled independently from the $\text{Uniform}[0, \xi]$ distribution. This step can also be skipped by setting $\xi = 0$. Finally, any AMDP solver **SolveAMDP** can be used in line 7, but our theorems each require certain conditions on the degree of suboptimality of the output policy $\hat{\pi}$ guaranteed by the solver.

3.1 Standard Plug-in Approach

We first analyze arguably the most natural algorithm for learning optimal policies in AMDPs, the plug-in approach: we form an empirical transition matrix $\hat{P}$ using transition counts from the generative model and then compute a bias-optimal policy $\hat{\pi}$ for the AMDP $(\hat{P}, r)$. This corresponds to Algorithm 1 with no perturbation ($\xi = 0$) and no anchoring ($\eta = 0$). To the best of our knowledge, the following results are the first for this simple algorithm.

Theorem 1. *Suppose P is weakly communicating. Consider Algorithm 1 with $\eta = 0$ and $\xi = 0$. Suppose that the policy $\hat{\pi}$ returned by **SolveAMDP** is guaranteed to be a bias-optimal policy of the AMDP $(\hat{P}, r)$. Let*

$\hat{h}^*$ be the optimal bias of $(\hat{P}, r)$, and let $\hat{\gamma}^*$ be the (random) smallest discount factor such that for all $\gamma \geq \hat{\gamma}^*$, there exists $c \in \mathbb{R}$ (which may depend on γ) such that

$$\left\| \hat{V}_\gamma^* - \hat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n}. \quad (1)$$

Then with probability $1 - \delta$, if $\hat{P}$ is weakly communicating, then

$$\rho^{\hat{\pi}} - \rho^* \leq \sqrt{\frac{C_7 \log^3 \left(\frac{SAn}{\delta(1-\hat{\gamma}^*)} \right)}{n}} \left(\|h^*\|_{\text{span}} + \|\hat{h}^*\|_{\text{span}} + 1 \right) \mathbf{1}.$$

Compared to the minimax optimal rate of $\tilde{O}\left(\sqrt{\frac{\|h^*\|_{\text{span}}}{n}}\right)$ (equivalent to $\tilde{O}\left(SA \frac{\|h^*\|_{\text{span}}}{\varepsilon^2}\right)$ sample complexity), Theorem 1 has the additional term $\|\hat{h}^*\|_{\text{span}}$. We show in Theorem 14 that this additional term is unavoidable for the plug-in approach, in the sense that there exist instances where the plug-in approach satisfies a high-probability bound of $\tilde{O}\left(\sqrt{\frac{\|h^*\|_{\text{span}} + \|\hat{h}^*\|_{\text{span}}}{n}}\right)$ but not $\tilde{O}\left(\sqrt{\frac{\|h^*\|_{\text{span}}}{n}}\right)$. One key feature of the optimal algorithm [Zurek and Chen, 2024], based on DMDP reduction, is that it requires prior knowledge of $\|h^*\|_{\text{span}}$ to set the discount factor, whereas the plug-in method has no need for such information. Lemma 25 establishes basic properties of the quantity $\hat{\gamma}^*$ appearing in Theorem 1. In particular it is well-defined when $\hat{P}$ is weakly communicating.

Within the proof of Theorem 1 we analyze the accuracy of using $(\hat{P}, r)$ to estimate the gain of a fixed policy, leading to the following policy evaluation result of independent interest.

Theorem 2. Fix a policy π such that ρ^π is constant. Let $\hat{\rho}^\pi$ be the gain of π in the empirical AMDP $(\hat{P}, r)$. Then with probability at least $1 - \delta$,

$$\|\hat{\rho}^\pi - \rho^\pi\|_\infty \leq \sqrt{\frac{C_4 \log^3 \left(\frac{SAn}{\delta} \right)}{n}} (\|h^\pi\|_{\text{span}} + 1).$$

3.2 Anchoring-Based Plug-in Approach

Although the standard plug-in approach is arguably the most natural algorithm for learning in AMDPs with a generative model, it and our Theorem 1 have a few limitations. First, the performance bound in Theorem 1 only holds on the event that $\hat{P}$ is weakly communicating, which can be understood as a consequence of the fact that the standard plug-in approach does not provide a way for us to incorporate prior information that P is weakly communicating. Additionally, the $\log\left(\frac{1}{1-\hat{\gamma}^*}\right)$ term is not bounded as $\log \text{poly}(SAn\delta)$ in the worst case, preventing us from applying Theorem 1 to obtain optimal D or τ_{unif} -based rates. Finally, since arbitrary weakly communicating MDPs do not possess sufficient stability properties for a definition of policy near-optimality which suffices for our purposes, our Theorem 1 requires finding an exactly bias-optimal policy in $(\hat{P}, r)$.

Fortunately, we can overcome all of these limitations with a simple technique which has been used many times (for various purposes) in the literature on average-reward reinforcement learning, which we term *anchoring*. For a small probability $\eta \in [0, 1]$ and any arbitrarily chosen state s_0 , we can form the *anchored* transition matrix $\hat{\underline{P}} = (1 - \eta)\hat{P} + \eta \mathbf{1} e_{s_0}^\top$ (where $\mathbf{1} \in \mathbb{R}^{SA}$ is all-1 and $e_{s_0} \in \mathbb{R}^S$ is all-0 except for a 1 in entry s_0). In words, $\hat{\underline{P}}$ follows $\hat{P}$ a $(1 - \eta)$ fraction of the time, but all state-action pairs have a small chance η to return to s_0 . This technique has been used in average-reward and related settings (e.g., Fruit et al. 2018, Yin et al. 2022) for essentially *computational* reasons, since it ensures that the associated Bellman operator is a $1 - \eta$ (span-)contraction [Puterman, 1994, Theorem 6.6.6], whereas without anchoring there is no guarantee of contractivity and thus standard average-reward value iteration has no finite-time convergence guarantee. Such works often set $\eta = \frac{1}{n}$, in which case these computational benefits are essentially without loss of statistical efficiency since an order $O(1/n)$ perturbation contributes a lower-order term relative to the statistical error. An arguably more standard perspective would be to consider anchoring + value iteration as a particular solver for the empirical AMDP $(\hat{P}, r)$, but we can incorporate anchoring within the plug-in

framework and thus allow arbitrary AMDP solvers by having them solve the anchored AMDP $(\hat{P}, r)$. $\hat{P}$ is always weakly communicating (in fact unichain), thus providing a simple way to enforce our model estimate to be weakly communicating.

Lemma 24 summarizes all these (and other) properties of the anchoring technique, in particular showing that anchoring with $\eta = \frac{1}{n}$ is essentially equivalent to DMDP reduction with an effective horizon of $\frac{1}{1-\gamma} = n$. This is a much larger effective horizon than those used in prior work on DMDP reduction for solving AMDP [Jin and Sidford, 2021, Wang et al., 2023b, 2022, Zurek and Chen, 2024] and unlike prior work, does not require knowledge of complexity parameters such as τ_{unif} or $\|h^*\|_{\text{span}}$. Prior analysis of DMDP (e.g. Agarwal et al. [2020], Li et al. [2020]) does not allow or gives vacuous guarantees for $\frac{1}{1-\gamma} = n$, but our novel analysis for the AMDP plug-in method (which heuristically is a DMDP reduction with arbitrarily large effective horizon) can be repurposed to handle this situation. We thus believe anchoring is better understood as a stabilized method for directly solving the AMDP $(\hat{P}, r)$ rather than as a discounted reduction. Still, the DMDP reduction method (with horizon n) obtains nearly identical guarantees to those in this section, which we provide in Appendix F.

Now we present our first result on the anchored AMDP plug-in approach. We define $\hat{\rho}^\pi$ and $\hat{h}^\pi$ as the gain and bias of a policy π in the anchored AMDP $(\hat{P}, r)$, and likewise define $\hat{\rho}^*$ and $\hat{h}^*$ as the optimal gain and bias in the anchored AMDP.

Theorem 3. *Suppose P is weakly communicating. Let s_0 be an arbitrary state, let $\eta = \frac{1}{n}$, and set $\xi = 0$ in Algorithm 1. Also suppose that SolveAMDP is guaranteed to return a policy $\hat{\pi}$ satisfying*

$$\hat{\rho}^{\hat{\pi}} \geq \hat{\rho}^* - \frac{1}{3n^2} \quad \text{and} \quad \|\hat{h}^{\hat{\pi}} - \hat{h}^*\|_\infty \leq \frac{1}{3n^2}. \quad (2)$$

Then with probability at least $1 - \delta$,

$$\rho^{\hat{\pi}} - \rho^* \leq \sqrt{\frac{C_5 \log^3\left(\frac{SAn}{\delta}\right)}{n}} \left(\|h^*\|_{\text{span}} + \|\hat{h}^*\|_{\text{span}} + 1 \right) \mathbf{1}.$$

Regarding the difference between the terms $\|\hat{h}^*\|_{\text{span}}$ and $\|\hat{h}^*\|_{\text{span}}$ appearing in Theorems 3 and 1 respectively, by Lemma 24, whenever $\hat{P}$ is weakly communicating (which is required for the bound within Theorem 1 to hold), we have that $\|\hat{h}^*\|_{\text{span}} \leq O(\|\hat{h}^*\|_{\text{span}})$.

We now apply Theorem 3 to the diameter-based complexity setting, where we assume P is communicating with diameter D and derive a complexity bound depending on D . Theorem 3 will yield an optimal sample complexity $\tilde{O}(SA\frac{D}{\epsilon^2})$ (matching the lower bound in Wang et al. 2022), and this optimal complexity follows directly from upper-bounding the guarantee of Theorem 3 in terms of D without any algorithmic modifications. In particular, no prior knowledge of D is required. The optimal bias span is always bounded by the diameter [Bartlett and Tewari, 2012, Lattimore and Szepesvári, 2020], so we have $\|h^*\|_{\text{span}} \leq D$, and similarly it is possible to show $\|\hat{h}^*\|_{\text{span}} \leq O(\hat{D})$ where $\hat{D}$ is the diameter of the empirical MDP $\hat{P}$. The key fact is that whenever $n \geq \tilde{\Omega}(D)$, we additionally have $\hat{D} \leq O(D)$, that is, $\hat{P}$ will be communicating and have diameter order D .

Lemma 4. *Suppose that the MDP P is communicating and has diameter D . Then there exists a constant C_7 such that if $n \geq C_8 D \log^3\left(\frac{SADn}{\delta}\right)$, then with probability at least $1 - \delta$, $\hat{D} \leq 14D$. In particular, in this same event, $\hat{P}$ is communicating.*

Lemma 4 follows from our later results on DMDPs, using the fact that the maximum travel time between states in $\hat{P}$ can be bounded by analyzing certain discounted value functions associated with auxiliary problems each measuring travel time to a certain state.

Corollary 5. *Suppose P is communicating and has diameter D . Let s_0 be an arbitrary state, let $\eta = \frac{1}{n}$, and set $\xi = 0$ in Algorithm 1. Also suppose that SolveAMDP is guaranteed to return a policy $\hat{\pi}$ which satisfies condition (2). Then with probability at least $1 - \delta$,*

$$\rho^{\hat{\pi}} - \rho^* \leq \sqrt{\frac{C_9 \log^3\left(\frac{SADn}{\delta}\right)}{n}} D \mathbf{1}.$$

The only other method which does not require prior knowledge of D and which yields a diameter-based complexity guarantee is that of [Tuytman et al. \[2024\]](#). In contrast to the explicit diameter estimation employed by [Tuytman et al. \[2024\]](#) which leads to a worse error bound (since for large values of ε , the diameter estimation subroutine dominates the complexity), [Theorem 5](#) yields the optimal diameter-based complexity for the full range of ε and does so with a simpler algorithm.

Another important and heavily studied sub-setting is the uniformly mixing setting, wherein all deterministic Markovian policies are assumed to have bounded mixing time τ_{unif} in P , and the goal is to obtain a complexity bound in terms of τ_{unif} . We conjecture that [Theorem 3](#) should also imply an optimal complexity of $\tilde{O}(SA\frac{\tau_{\text{unif}}}{\varepsilon^2})$ for this setting by an analogous argument. We always have $\|h^*\|_{\text{span}} \leq O(\tau_{\text{unif}})$ ([Lemma 39; Wang et al. 2022](#)), however we are unaware of how to bound the uniform mixing time of $\hat{P}$ by τ_{unif} for small values of n (which would imply that $\|\hat{h}^*\|_{\text{span}} \leq O(\tau_{\text{unif}})$). For this reason we present a slightly different algorithm and guarantee which replaces the $\|\hat{h}^*\|_{\text{span}}$ term with $\|\underline{h}^{\hat{\pi}}\|_{\text{span}}$, the bias span of the returned policy $\hat{\pi}$ in the *true* anchored MDP $\underline{P} = (1 - \eta)P + (1 - \eta)\mathbf{1}e_{s_0}^\top$, which can be straightforwardly bounded as $O(\tau_{\text{unif}})$.

Theorem 6. *Suppose P is weakly communicating. Let s_0 be an arbitrary state, let $\eta = \frac{1}{n}$, and set $\xi \in (0, \frac{1}{n}]$ in [Algorithm 1](#). Also suppose that the policy $\hat{\pi}$ returned by *SolveAMDP* is guaranteed to be the exact Blackwell-optimal policy of the AMDP $(\hat{P}, \tilde{r})$. Then with probability at least $1 - \delta$,*

$$\rho^* - \rho^{\hat{\pi}} \leq \sqrt{\frac{C_3 \log^3\left(\frac{SA n}{\delta \xi}\right)}{n}} \left(\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right) \mathbf{1}.$$

While finding a Blackwell-optimal policy may generally be computationally expensive, similarly to the discounted setting [Li et al. \[2020\]](#), with high probability the perturbation ensures a small separation between the bias of the Blackwell optimal policy of $(\hat{P}, \tilde{r})$ and all other (Markovian, deterministic) gain-optimal policies, thus ensuring that Blackwell optimality for $(\hat{P}, \tilde{r})$ reduces to bias optimality and also that $\tilde{O}(n)$ steps of value iteration suffice to find an exactly Blackwell-optimal policy. See [Lemma 28](#) for the formal statement.

Now, as promised, we can show that the plug-in approach with anchoring and perturbation obtains the optimal τ_{unif} -based sample complexity.

Corollary 7. *Suppose P has a finite uniform mixing time τ_{unif} . Let s_0 be an arbitrary state, let $\eta = \frac{1}{n}$, and set $\xi \in (0, \frac{1}{n}]$ in [Algorithm 1](#). Also suppose that the policy $\hat{\pi}$ returned by *SolveAMDP* is the exact Blackwell-optimal policy of the AMDP $(\hat{P}, \tilde{r})$. Then with probability at least $1 - \delta$,*

$$\rho^* - \rho^{\hat{\pi}} \leq \sqrt{\frac{7C_3 \log^3\left(\frac{SA n}{\delta \xi}\right)}{n}} \tau_{\text{unif}} \mathbf{1}.$$

Compared to extensive prior work on this setting, this is the first algorithm which achieves the optimal complexity without requiring prior knowledge of τ_{unif} , and additionally we believe the algorithm is much simpler than previous approaches.

Rather than having two different (albeit highly similar, differing only in whether or not the reward vector r is perturbed) algorithms which are optimal for different settings, one might prefer to have one algorithm which achieves the best of all the aforementioned guarantees. In fact, we can view the exact solution of the *perturbed* empirical AMDP $(\hat{P}, \tilde{r})$ as an approximate solution of the *unperturbed* empirical AMDP $(\hat{P}, r)$, with the degree of suboptimality depending on the perturbation magnitude ξ , and thus for sufficiently small ξ , we can also apply the performance guarantees for the unperturbed anchored plug-in approach from [Theorem 3](#) to the perturbed anchored plug-in approach.

Theorem 8. *Suppose P is weakly communicating. Let s_0 be an arbitrary state, let $\eta = \frac{1}{n}$, and set $\xi = \frac{1}{2n^2}$ in [Algorithm 1](#). Also suppose that the policy $\hat{\pi}$ returned by *SolveAMDP* is guaranteed to be the exact Blackwell-optimal policy of the AMDP $(\hat{P}, \tilde{r})$. Then with probability at least $1 - \delta$,*

$$\rho^* - \rho^{\hat{\pi}} \leq \sqrt{\frac{C_6 \log^3\left(\frac{SA n}{\delta}\right)}{n}} \left(\|h^*\|_{\text{span}} + \min\{\|\hat{h}^{\hat{\pi}}\|_{\text{span}}, \|\underline{h}^{\hat{\pi}}\|_{\text{span}}\} + 1 \right).$$

Algorithm 2 Plug-in approach for DMDP

input: Sample size per state-action pair n , discount factor γ ; optional perturbation level ξ

1: **for** each state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ **do**

2: Collect n samples $S_{s,a}^1, \dots, S_{s,a}^n$ from $P(\cdot \mid s, a)$

3: Form the empirical transition kernel $\hat{P}(s' \mid s, a) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}\{S_{s,a}^i = s'\}$, for all $s' \in \mathcal{S}$

4: **end for**

5: Form *perturbed* reward $\tilde{r} = r + \Delta$ where $\Delta(s, a) \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}[0, \xi]$

▷ Set $\xi = 0$ for no perturbation

6: $\hat{\pi} = \text{SolveDMDP}(\hat{P}, \tilde{r}, \gamma)$

7: **return** $\hat{\pi}$

Following identical steps as in the Corollaries 5 and 7, we can thus show that the anchored plug-in approach with $\xi = \frac{1}{2n^2}$ automatically satisfies the optimal error bounds for both the diameter and uniform-mixing-based settings, without any required prior knowledge. Since some prior DMDP-reduction-based methods for these settings require the discount factor to be set in terms of D or τ_{unif} , prior algorithms which are optimal for the diameter-based setting may not be optimal for the uniformly-mixing setting and vice versa, unlike our result.

3.3 Plug-in Approach for Discounted MDPs

The new analysis techniques developed for the AMDP plug-in approach also lead to improvements for the DMDP plug-in method. Similarly to the previous section, we present a meta-algorithm, Algorithm 2, and allow different choices of perturbation ξ and solver `SolveDMDP` within the theorems.

Theorem 9. *Let $\xi = 0$ in Algorithm 2. Suppose that `SolveDMDP` returns a policy $\hat{\pi}$ satisfying*

$$\hat{V}^{\hat{\pi}} \geq \hat{V}^* - \frac{1}{n} \mathbf{1}. \quad (3)$$

Then with probability at least $1 - \delta$,

$$\|V^{\hat{\pi}} - V^*\|_{\infty} \leq \frac{1}{1-\gamma} \sqrt{\frac{C_1 \log^3 \left(\frac{SA n}{(1-\gamma)\delta} \right)}{n} (\|V^*\|_{\text{span}} + \|\hat{V}^{\hat{\pi}}\|_{\text{span}} + 1)}.$$

Using the bounds $\|V^*\|_{\text{span}}, \|\hat{V}^{\hat{\pi}}\|_{\text{span}} \leq \frac{1}{1-\gamma}$, Theorem 9 is the first to imply that the discounted plug-in approach attains the minimax optimal sample complexity of $\tilde{O} \left(\frac{SA}{(1-\gamma)^3 \varepsilon^2} \right)$ *without perturbation*.

Analogously to the situation for the anchored AMDP plug-in approach, by adding perturbation we can replace the $\|\hat{V}^{\hat{\pi}}\|_{\text{span}}$ term with $\|V^{\hat{\pi}}\|_{\text{span}}$.

Theorem 10. *Set $\xi \in (0, \frac{1}{n}]$ in Algorithm 2. Suppose that the policy $\hat{\pi}$ returned by `SolveDMDP` is guaranteed to be exactly optimal for the DMDP $(\hat{P}, \tilde{r}, \gamma)$. Then with probability at least $1 - \delta$,*

$$\|V^{\hat{\pi}} - V^*\|_{\infty} \leq \frac{1}{1-\gamma} \sqrt{\frac{C_2 \log^3 \left(\frac{SA n}{(1-\gamma)\delta\xi} \right)}{n} (\|V^*\|_{\text{span}} + \|V^{\hat{\pi}}\|_{\text{span}} + 1)}.$$

As shown in Li et al. [2020], the perturbation ensures that the exact optimal policy of the DMDP $(\hat{P}, \tilde{r}, \gamma)$ can be computed in finite time, for example with $\tilde{O} \left(\frac{1}{1-\gamma} \right)$ value iteration steps.

Also analogously to the situation for the AMDP plug-in approach, for sufficiently small perturbation ξ we can combine both guarantees for one algorithm.

Theorem 11. Set $\xi = \frac{1-\gamma}{2n}$ in Algorithm 2. Suppose that the policy $\hat{\pi}$ returned by *SolveDMDP* is the exact optimal policy of the DMDP $(\hat{P}, \tilde{r}, \gamma)$. Then with probability at least $1 - \delta$,

$$\|V^{\hat{\pi}} - V^{\star}\|_{\infty} \leq \frac{1}{1-\gamma} \sqrt{\frac{C_{10} \log^3\left(\frac{SA n}{(1-\gamma)\delta}\right)}{n}} \left(\|V^{\star}\|_{\text{span}} + \min\{\|\hat{V}^{\hat{\pi}}\|_{\text{span}}, \|V^{\hat{\pi}}\|_{\text{span}}\} + 1 \right).$$

Now we discuss the relationship between the terms $\|V^{\star}\|_{\text{span}}$, $\|\hat{V}^{\hat{\pi}}\|_{\text{span}}$, and $\|V^{\hat{\pi}}\|_{\text{span}}$ appearing in our theorems and the usual complexity parameters τ_{unif} , D , and $\|h^{\star}\|_{\text{span}}$. If P is weakly communicating, we have $\|V^{\star}\|_{\text{span}} \leq 2\|h^{\star}\|_{\text{span}}$ [Wei et al., 2020, Lemma 2], and as previously mentioned we have $\|h^{\star}\|_{\text{span}} \leq 3\tau_{\text{unif}}$ (Lemma 39, Wang et al. [2022]) and $\|h^{\star}\|_{\text{span}} \leq D$ [Bartlett and Tewari, 2012, Lattimore and Szepesvári, 2020]. Therefore, under the event described in Lemma 4, since $\hat{P}$ is communicating with diameter $\hat{D} \leq O(D)$, we can apply these same statements to $\|\hat{V}^{\star}\|_{\text{span}}$ to obtain that $\|\hat{V}^{\star}\|_{\text{span}} \leq 2\|\hat{h}^{\star}\|_{\text{span}}$ (since $\hat{P}$ is weakly communicating) and that $\|\hat{h}^{\star}\|_{\text{span}} \leq \hat{D}$, and thus that $\|\hat{V}^{\star}\|_{\text{span}} \leq O(D)$, and finally the optimality condition (3) implies that $\|\hat{V}^{\hat{\pi}}\|_{\text{span}} \leq \|\hat{V}^{\star}\|_{\text{span}} + \frac{1}{n} \leq O(D)$. Therefore, Theorem 9 (combined with Lemma 4) implies a $\tilde{O}\left(SA \frac{D}{(1-\gamma)^2 \varepsilon^2}\right)$ sample complexity bound, for the full nontrivial range of $\varepsilon \in (0, \frac{1}{1-\gamma}]$. Also, we have the bound $\|V^{\hat{\pi}}\|_{\text{span}} \leq 3\tau_{\text{unif}}$ (see Lemma 39), which when combined with Theorem 10 implies a $\tilde{O}\left(SA \frac{\tau_{\text{unif}}}{(1-\gamma)^2 \varepsilon^2}\right)$ sample complexity bound, also for the entire nontrivial range of $\varepsilon \in (0, \frac{1}{1-\gamma}]$. This improves on Wang et al. [2023b] which only obtains this complexity for $\varepsilon \leq \sqrt{\frac{\tau_{\text{unif}}}{1-\gamma}}$. (See Wang et al. [2023a] for the matching lower bound.)

When n is sufficiently large relative to other problem parameters we can show that $\|\hat{V}^{\star}\|_{\text{span}}$ is bounded in terms of $\|h^{\star}\|_{\text{span}}$ akin to $\|V^{\star}\|_{\text{span}}$.

Lemma 12. If P is weakly communicating and $n \geq \frac{C_1 \log^3\left(\frac{SA n}{(1-\gamma)\delta}\right)}{(1-\gamma)^2(\|h^{\star}\|_{\text{span}}+1)}$, then with probability at least $1 - \delta$, $\|\hat{V}^{\star}\|_{\text{span}} \leq 4(\|h^{\star}\|_{\text{span}} + 1)$.

When P is weakly communicating, the algorithm of Zurek and Chen [2024] achieves the span-based bound $\tilde{O}\left(SA \frac{\|h^{\star}\|_{\text{span}}+1}{(1-\gamma)^2 \varepsilon^2}\right)$ under the restriction that $\varepsilon \leq \|h^{\star}\|_{\text{span}}$, or equivalently that $n \geq \tilde{\Omega}\left(\frac{1}{(1-\gamma)^2(\|h^{\star}\|_{\text{span}}+1)}\right)$. Under this condition the requirement of Lemma 12 will be met, and by combining it with our Theorem 9, we recover the result of Zurek and Chen [2024]. An analogous version of Lemma 12 could also be shown to bound $\|V^{\hat{\pi}}\|_{\text{span}}$. However, similarly to the situation for the average-reward plug-in method, generally the terms $\|V^{\hat{\pi}}\|_{\text{span}}$ and $\|\hat{V}^{\hat{\pi}}\|_{\text{span}}$ cannot be removed from the analysis of the DMDP plug-in approach, as is shown in Theorem 14.

AMDP-to-DMDP Reduction Approach While our focus is not on analyzing the well-studied AMDP-to-DMDP reduction approach for solving AMDPs, we briefly mention some corollaries of our DMDP results for the complexity of this method. First, for target AMDP error ε , if $\|h^{\star}\|_{\text{span}}$ is known, then we can use an effective horizon of $\frac{1}{1-\gamma} = C \frac{\|h^{\star}\|_{\text{span}}}{\varepsilon}$ (as do Wang et al. [2022], Zurek and Chen [2024]) and the condition in Lemma 12 will be satisfied as long as we have $n \geq \tilde{\Omega}\left(\frac{\|h^{\star}\|_{\text{span}}+1}{\varepsilon^2}\right)$. Combining the resulting error bound with Theorem 9 and with standard DMDP reduction results [Wang et al., 2022], this recovers the result of Zurek and Chen [2024] which obtains the optimal $\tilde{O}\left(SA \frac{\|h^{\star}\|_{\text{span}}+1}{\varepsilon^2}\right)$ sample complexity, but we remove the need for reward perturbation.

More interestingly, we can satisfy the conditions of Lemma 12 with a smaller effective horizon of approximately $\sqrt{n}$. This is too small to yield the optimal complexity, since a DMDP reduction with discount γ incurs error of order $(1-\gamma)\|h^{\star}\|_{\text{span}} \approx \frac{\|h^{\star}\|_{\text{span}}}{\sqrt{n}}$ (even with infinite samples) [Wang et al., 2022]. However, since Lemma 12 holds, we can obtain the first complexity bound depending only on S, A, ε , and $\|h^{\star}\|_{\text{span}}$ without requiring prior knowledge of $\|h^{\star}\|_{\text{span}}$.

Theorem 13. Suppose P is weakly communicating. Let $\xi = 0$ and $\gamma = 1 - \sqrt{\frac{C_1 \log^3\left(\frac{SA n^2}{\delta}\right)}{n}}$ in Algorithm 9

and suppose that *SolveMDP* guarantees (3). Then with probability $1 - \delta$,

$$\rho^* - \rho^{\hat{\pi}} \leq \sqrt{\frac{C_{11} \log^3(\frac{SAn}{\delta})}{n}} (\|h^*\|_{\text{span}}^2 + 1) \mathbf{1}.$$

3.4 Hard Instance for Plug-in Approach

Now we provide a concrete MDP where both the average-reward and discounted plug-in methods fail to achieve an $\|h^*\|_{\text{span}}$ -based complexity, implying that the terms $\|\hat{h}^*\|_{\text{span}}$, $\|\hat{V}^{\hat{\pi}}\|_{\text{span}}$, and $\|V^{\hat{\pi}}\|_{\text{span}}$ cannot be generally removed from the Theorems 1, 9, and 10, respectively.

Theorem 14. *For any fixed $n \geq 10$, there exists an MDP P (depending on n) with $S = A = 2$ such that P has $\|h^*\|_{\text{span}} = 1$, diameter $D = n$, and a uniform mixing time $\tau_{\text{unif}} = \Theta(n)$. Also, with probability at least $\frac{1}{25}$,*

1. $\hat{P}$ is communicating.
2. Letting $\hat{\pi}^*$ be the Blackwell-optimal policy of the AMDP $(\hat{P}, r)$, $\|\rho^* - \rho^{\hat{\pi}^*}\|_{\infty} \geq \frac{1}{5}$.
3. Letting $\hat{\pi}_{\gamma}^*$ be the optimal policy for the DMDP $(\hat{P}, \gamma, r)$ with effective horizon $\frac{1}{1-\gamma} = n^2$, $\|V_{\gamma}^* - V_{\gamma}^{\hat{\pi}^*}\|_{\infty} \geq \frac{n^2}{5} = \frac{1}{1-\gamma} \frac{1}{5}$.

Consequently, for any constant $C > 0$, there exists n, P, r such that the statement

$$\mathbb{P}\left(\|\rho^{\hat{\pi}^*} - \rho^*\|_{\infty} \leq C \sqrt{\frac{\|h^*\|_{\text{span}} \log(\|h^*\|_{\text{span}} n)}{n}}\right) > 1 - \frac{1}{25}$$

is false. Additionally, for any constant $C > 0$, there exists n, P, r, γ such that the statement

$$\mathbb{P}\left(\|V_{\gamma}^{\hat{\pi}^*} - V_{\gamma}^*\|_{\infty} \leq C \frac{1}{1-\gamma} \sqrt{\frac{\|h^*\|_{\text{span}} \log(\|h^*\|_{\text{span}} n)}{n}}\right) > 1 - \frac{1}{25}$$

is false.

We show the construction for P in Appendix G, along with the proof of Theorem 14. At a high level, P causes constant probability of sampling $\hat{P}$ which has optimal bias span $\|\hat{h}^*\|_{\text{span}} \gg \|h^*\|_{\text{span}}$. If we had knowledge of the true $\|h^*\|_{\text{span}}$ we could use it to find a near-optimal policy with controlled complexity, which is accomplished by DMDP reduction using a $\|h^*\|_{\text{span}}$ -based effective horizon [Zurek and Chen, 2024]. In contrast the AMDP plug-in method (and the DMDP plug-in method with sufficiently large horizon) has no way of controlling the span of the empirical optimal policy, leading to a potentially greater span than $\|h^*\|_{\text{span}}$ and correspondingly larger error.

4 Proof Techniques

At the heart of all our main results is a novel decomposition of the difference between limiting distributions associated with $\hat{P}$ and P . This technique bears some resemblance to the “higher-order” simulation lemma expansion introduced by Li et al. [2020] to show the DMDP plug-in method achieves $\tilde{O}\left(SA \frac{1}{(1-\gamma)^3 \varepsilon^2}\right)$ sample complexity for the full range $\varepsilon \in (0, \frac{1}{1-\gamma}]$. Both techniques decompose the error with the simulation lemma, use law-of-total-variance-style arguments to bound some leading terms, and obtain lower-order error terms which can be inductively bounded again with the simulation lemma. However, there are many subtle differences. Most importantly for the average-reward setting, the arguments of Li et al. [2020] require $n \geq \Omega(\frac{1}{1-\gamma})$, which is without loss of generality for the minimax rate of $\tilde{O}\left(\frac{1}{\sqrt{(1-\gamma)^3 n}}\right)$ (it is necessary for nontrivial accuracy), but breaks for the arbitrarily large effective horizons needed in average-reward problems. Also, the arguments of Li et al. [2020] are designed to use concentration inequalities involving variance parameters of certain auxiliary MDPs, which requires a more delicate leave-one-out analysis (hence their use of reward perturbation), whereas our argument requires concentration bounds on terms which

are simpler functions of the original/empirical MDPs, enabling the flexibility to utilize the absorbing MDP arguments of Agarwal et al. [2020] or those of Li et al. [2020].

We briefly illustrate our techniques as applied to the proof of the policy evaluation bound Theorem 2. Hence we fix a policy π with constant gain ρ^π and attempt to bound $\|\hat{\rho}^\pi - \rho^\pi\|_\infty$. In this sketch we use the $\lesssim$ notation to ignore constants and log factors. By an average-reward version of the simulation lemma (see Lemma 23), since ρ^π is a constant vector,

$$\hat{\rho}^\pi - \rho^\pi = \hat{P}_\pi^\infty r_\pi - P_\pi^\infty r_\pi = \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) h^\pi \quad (4)$$

Using Bernstein's inequality (S times), we can obtain an elementwise inequality

$$\left| (\hat{P}_\pi - P_\pi) h^\pi \right| \lesssim \sqrt{\frac{\mathbb{V}_{P_\pi}[h^\pi]}{n}} + \frac{\|h^\pi\|_\infty}{n} \mathbf{1} \quad (5)$$

where $\mathbb{V}_{P_\pi}[h^\pi] = P_\pi(h^\pi)^{\circ 2} - (P_\pi h^\pi)^{\circ 2}$ is a (vector) variance of the next-state bias function. ($x^{\circ k}$ denotes the elementwise k th-power operation.) Since all entries of $\hat{P}_\pi^\infty$ are nonnegative, we can combine this with (4) and obtain

$$|\hat{\rho}^\pi - \rho^\pi| \leq \hat{P}_\pi^\infty \left| (\hat{P}_\pi - P_\pi) h^\pi \right| \lesssim \hat{P}_\pi^\infty \sqrt{\frac{\mathbb{V}_{P_\pi}[h^\pi]}{n}} + \frac{\|h^\pi\|_\infty}{n} \hat{P}_\pi^\infty \mathbf{1}. \quad (6)$$

The second term of (6) is $\frac{\|h^\pi\|_\infty}{n} \hat{P}_\pi^\infty \mathbf{1} = \frac{\|h^\pi\|_\infty}{n} \mathbf{1}$, which is smaller than the desired bound $\sqrt{\frac{\|h^\pi\|_\infty}{n}}$ (in the nontrivial accuracy regime where both of these terms are ≤ 1) so we focus on bounding the first term in the RHS of (6). Using Jensen's inequality to move $\hat{P}_\pi^\infty$ inside the square root and the Poisson/Bellman equation $\rho^\pi + h^\pi = r_\pi + P_\pi h^\pi$, we have (elementwise)

$$\begin{aligned} \hat{P}_\pi^\infty \sqrt{\mathbb{V}_{P_\pi}[h^\pi]} &\leq \sqrt{\hat{P}_\pi^\infty \mathbb{V}_{P_\pi}[h^\pi]} = \sqrt{\hat{P}_\pi^\infty (P_\pi(h^\pi)^{\circ 2} - (P_\pi h^\pi)^{\circ 2})} \\ &= \sqrt{\hat{P}_\pi^\infty \left(P_\pi(h^\pi)^{\circ 2} - ((\rho^\pi - r_\pi) + h^\pi)^{\circ 2} \right)} \\ &= \sqrt{\hat{P}_\pi^\infty (P_\pi(h^\pi)^{\circ 2} - (h^\pi)^{\circ 2} + 2(\rho^\pi - r_\pi) \circ h^\pi - (\rho^\pi - r_\pi)^{\circ 2})} \\ &\leq \sqrt{\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 2} + 2\hat{P}_\pi^\infty \|\rho^\pi - r_\pi\|_\infty \|h^\pi\|_\infty \mathbf{1}} \\ &\lesssim \sqrt{\left| \hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 2} \right|} + \sqrt{\|h^\pi\|_\infty} \mathbf{1}. \end{aligned} \quad (7)$$

Combining all these steps, we have shown

$$|\hat{\rho}^\pi - \rho^\pi| \lesssim \frac{\|h^\pi\|_\infty}{n} \mathbf{1} + \sqrt{\frac{\|h^\pi\|_\infty}{n}} \mathbf{1} + \frac{1}{\sqrt{n}} \sqrt{\left| \hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 2} \right|}. \quad (8)$$

Using that $\hat{P}_\pi^\infty \hat{P}_\pi = \hat{P}_\pi^\infty$, we can recognize that $\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 2} = \hat{P}_\pi^\infty (P_\pi - \hat{P}_\pi)(h^\pi)^{\circ 2}$, a term of a very similar form to the RHS of the average-reward simulation lemma step (4). This suggests that we can apply analogous steps to bound this term and thus replace the final term in the RHS of (8) with lower-order quantities. Using Bernstein's inequality again,

$$\left| (\hat{P}_\pi - P_\pi)(h^\pi)^{\circ 2} \right| \lesssim \sqrt{\frac{\mathbb{V}_{P_\pi}[(h^\pi)^{\circ 2}]}{n}} + \frac{\|(h^\pi)^{\circ 2}\|_\infty}{n} \mathbf{1} \leq \sqrt{\frac{P_\pi(h^\pi)^{\circ 4} - (P_\pi h^\pi)^{\circ 4}}{n}} + \frac{\|h^\pi\|_\infty^2}{n} \mathbf{1}$$

where we used $\mathbb{V}_{P_\pi}[(h^\pi)^{\circ 2}] = P_\pi(h^\pi)^{\circ 4} - (P_\pi(h^\pi)^{\circ 2})^{\circ 2} \leq P_\pi(h^\pi)^{\circ 4} - (P_\pi h^\pi)^{\circ 4}$ by Jensen's inequality. Thus

similarly to steps (6) and (7) we can bound

$$\begin{aligned}
|\hat{P}_\pi^\infty(\hat{P}_\pi - P_\pi)(h^\pi)^{\circ 2}| &\leq \hat{P}_\pi^\infty |(\hat{P}_\pi - P_\pi)(h^\pi)^{\circ 2}| \\
&\lesssim \frac{1}{\sqrt{n}} \sqrt{\hat{P}_\pi^\infty (P_\pi(h^\pi)^{\circ 4} - (P_\pi h^\pi)^{\circ 4})} + \frac{\|h^\pi\|_\infty^2}{n} \mathbf{1} \\
&= \frac{1}{\sqrt{n}} \sqrt{\hat{P}_\pi^\infty (P_\pi(h^\pi)^{\circ 4} - ((\rho^\pi - r_\pi) + h^\pi)^{\circ 4})} + \frac{\|h^\pi\|_\infty^2}{n} \mathbf{1} \\
&\leq \frac{1}{\sqrt{n}} \sqrt{\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 4} + (2^4 - 1)(\|h^\pi\|_\infty + 1)^3 \mathbf{1}} + \frac{\|h^\pi\|_\infty^2}{n} \mathbf{1} \\
&\lesssim \frac{1}{\sqrt{n}} \sqrt{|\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 4}|} + \sqrt{\frac{2^4(\|h^\pi\|_\infty + 1)^3}{n}} \mathbf{1} + \frac{\|h^\pi\|_\infty^2}{n} \mathbf{1}
\end{aligned} \tag{9}$$

where we obtain inequality (9) by noticing $((\rho^\pi - r_\pi) + h^\pi)^{\circ 4}$ expands to 2^4 terms, one of which is $(h^\pi)^{\circ 4}$ and the rest of which have $\|\cdot\|_\infty$ bounded by $\|h^\pi\|_\infty^{4-k} \|\rho^\pi - r_\pi\|_\infty^k$ for some $k \in \{1, 2, 3, 4\}$, and also $\|\rho^\pi - r_\pi\|_\infty \leq 1$. We have shown

$$\begin{aligned}
\frac{1}{\sqrt{n}} \sqrt{|\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 2}|} &= \frac{1}{\sqrt{n}} \sqrt{|\hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi)(h^\pi)^{\circ 2}|} \\
&\lesssim \frac{1}{\sqrt{n}} \sqrt{\frac{1}{\sqrt{n}} \sqrt{|\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 4}|}} + \sqrt{\frac{2^4(\|h^\pi\|_\infty + 1)^3}{n}} \mathbf{1} + \frac{\|h^\pi\|_\infty^2}{n} \mathbf{1} \\
&\leq (2^4)^{1/4} \left(\frac{\|h^\pi\|_\infty + 1}{n} \right)^{3/4} \mathbf{1} + \frac{\|h^\pi\|_\infty}{n} \mathbf{1} + \frac{1}{n^{3/4}} |\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 4}|^{1/4}
\end{aligned}$$

and plugging back into (8) and simplifying, we have

$$|\hat{\rho}^\pi - \rho^\pi| \lesssim 2 \frac{\|h^\pi\|_\infty}{n} \mathbf{1} + \left(\frac{\|h^\pi\|_\infty + 1}{n} \right)^{1/2} \mathbf{1} + \left(\frac{\|h^\pi\|_\infty + 1}{n} \right)^{3/4} \mathbf{1} + \frac{1}{n^{3/4}} |\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 4}|^{1/4}.$$

As this argument suggests, we can continue bounding terms of the form $|\hat{P}_\pi^\infty (P_\pi - I)(h^\pi)^{\circ 2k}|^{2^{-k}}$, picking up additional terms which are lower-order relative to $\left(\frac{\|h^\pi\|_\infty + 1}{n} \right)^{1/2}$ and increasing the powers of 2. After roughly $\log_2 \log_2 \frac{\|h^\pi\|_\infty + 1}{n}$ steps all terms will be small enough to end the argument, yielding the desired bound $\|\hat{\rho}^\pi - \rho^\pi\|_\infty \lesssim \sqrt{\frac{\|h^\pi\|_\infty + 1}{n}}$. See Lemmas 30 and 31.

Now we briefly outline the additional steps required for our additional results. First, we note that a basically analogous argument, but with $(I - \gamma \hat{P}_\pi)^{-1}$ replacing $\hat{P}_\pi^\infty$, V_γ^π replacing h^π , and other straightforward adaptations, can be used in the DMDP setting, leading to our DMDP results. One important difference is that, while $\|h^\pi\|_{\text{span}}$ and $\|h^\pi\|_\infty$ are equivalent up to a factor of 2, we generally have $\|V_\gamma^\pi\|_\infty \gg \|V_\gamma^\pi\|_{\text{span}}$. However, all steps of the argument still go through if we replace V_γ^π by $\bar{V} = V_\gamma^\pi - (\min_s V_\gamma^\pi(s)) \mathbf{1}$, and thus the resulting bound will be in terms of $\|\bar{V}\|_\infty = \|V_\gamma^\pi\|_{\text{span}}$. See Lemmas 15 and 16 for details.

One final point is that unlike the sketched Theorem 2, our other results show the optimality of an empirical (near-)optimal policy $\hat{\pi}$. In the average-reward case (e.g. for Theorem 1) this requires bounding the two terms $\|\hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}}\|_\infty$ and $\|\hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}}\|_\infty$. The same technique as sketched above is still used, but the Bernstein inequality steps (e.g. (5)) require more care in order to decouple statistical dependency between $\hat{P}$ and $h^{\hat{\pi}}$. Agarwal et al. [2020] and Li et al. [2020] have developed different leave-one-out techniques for this purpose in DMDPs, either of which can be used to establish the ‘‘Bernstein-like’’ inequalities required in our argument. Anchoring plays a key role in facilitating the use of their DMDP-based bounds for AMDPs, since by Lemma 24, the bias functions in anchored AMDPs are equivalent (up to a constant shift) to certain DMDP value functions. See Lemmas 20, 22, 34, and 35 where we establish the Bernstein-like inequalities needed for our different results.

5 Conclusion

In this paper we performed the first analysis of the plug-in approach for average-reward MDPs, showing that this simple method obtains optimal rates for the diameter- and mixing-based settings without requiring prior knowledge. Our techniques also lead to improved results for DMDPs. While Theorem 14 suggests our span-based results cannot be improved for the plug-in method, it remains an interesting open question as to whether an improved algorithm can achieve the optimal $\tilde{O}\left(\frac{\|h^*\|_{\text{span}}}{\varepsilon^2}\right)$ sample complexity without knowledge of $\|h^*\|_{\text{span}}$. In conclusion, we believe this work fills a gap in our understanding of average-reward RL algorithms, and we hope that our techniques can be more broadly useful for the analysis of natural average-reward algorithms.

Acknowledgments

Y. Chen and M. Zurek were supported in part by National Science Foundation grants CCF-2233152 and DMS-2023239.

References

- Alekh Agarwal, Sham Kakade, and Lin F. Yang. Model-Based Reinforcement Learning with a Generative Model is Minimax Optimal, April 2020. URL <http://arxiv.org/abs/1906.03804>. arXiv:1906.03804 [cs, math, stat] version: 3.
- Mohammad Gheshlaghi Azar, Remi Munos, and Bert Kappen. On the Sample Complexity of Reinforcement Learning with a Generative Model, June 2012. URL <http://arxiv.org/abs/1206.6461>. arXiv:1206.6461 [cs, stat].
- Mohammad Gheshlaghi Azar, Rémi Munos, and Hilbert J. Kappen. Minimax PAC bounds on the sample complexity of reinforcement learning with a generative model. *Machine Learning*, 91(3):325–349, June 2013. ISSN 1573-0565. doi: 10.1007/s10994-013-5368-1. URL <https://doi.org/10.1007/s10994-013-5368-1>.
- Peter L. Bartlett and Ambuj Tewari. REGAL: A Regularization based Algorithm for Reinforcement Learning in Weakly Communicating MDPs, May 2012. URL <https://arxiv.org/abs/1205.2661v1>.
- X. R. Cao. Single Sample Path-Based Optimization of Markov Chains. *Journal of Optimization Theory and Applications*, 100(3):527–548, March 1999. ISSN 1573-2878. doi: 10.1023/A:1022634422482. URL <https://doi.org/10.1023/A:1022634422482>.
- Richard Durrett. *Probability: theory and examples*. Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, Cambridge, fifth edition edition, 2019. ISBN 978-1-108-47368-2.
- Ronan Fruit, Matteo Pirota, Alessandro Lazaric, and Ronald Ortner. Efficient Bias-Span-Constrained Exploration-Exploitation in Reinforcement Learning, July 2018. URL <http://arxiv.org/abs/1802.04020>. arXiv:1802.04020 [cs, stat].
- Ying Jin, Ramki Gummadi, Zhengyuan Zhou, and Jose Blanchet. Feasible Q^* -Learning for Average Reward Reinforcement Learning. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, pages 1630–1638. PMLR, April 2024. URL <https://proceedings.mlr.press/v238/jin24b.html>. ISSN: 2640-3498.
- Yujia Jin and Aaron Sidford. Efficiently Solving MDPs with Stochastic Mirror Descent, August 2020. URL <https://arxiv.org/abs/2008.12776v1>.
- Yujia Jin and Aaron Sidford. Towards Tight Bounds on the Sample Complexity of Average-reward MDPs, June 2021. URL <http://arxiv.org/abs/2106.07046>. arXiv:2106.07046 [cs, math].

- Michael Kearns and Satinder Singh. Finite-Sample Convergence Rates for Q-Learning and Indirect Algorithms. In *Advances in Neural Information Processing Systems*, volume 11. MIT Press, 1998. URL <https://proceedings.neurips.cc/paper/1998/hash/99adff456950dd9629a5260c4de21858-Abstract.html>.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, Cambridge ; New York, NY, 2020. ISBN 978-1-108-57140-1.
- David A. Levin and Yuval Peres. *Markov Chains and Mixing Times*. American Mathematical Soc., October 2017. ISBN 978-1-4704-2962-1.
- Gen Li, Yuting Wei, Yuejie Chi, Yuantao Gu, and Yuxin Chen. Breaking the Sample Size Barrier in Model-Based Reinforcement Learning with a Generative Model. In *Advances in Neural Information Processing Systems*, volume 33, pages 12861–12872. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/96ea64f3a1aa2fd00c72faacf0cb8ac9-Abstract.html>.
- Tianjiao Li, Feiyang Wu, and Guanghui Lan. Stochastic first-order methods for average-reward Markov decision processes, May 2022. URL <https://arxiv.org/abs/2205.05800v5>.
- Andreas Maurer and Massimiliano Pontil. Empirical Bernstein Bounds and Sample Variance Penalization, July 2009. URL <http://arxiv.org/abs/0907.3740>. arXiv:0907.3740 [stat] version: 1.
- Carl D. Meyer, Jr. The Condition of a Finite Markov Chain and Perturbation Bounds for the Limiting Probabilities. *SIAM Journal on Algebraic Discrete Methods*, 1(3):273–283, September 1980. ISSN 0196-5212. doi: 10.1137/0601031. URL <https://epubs.siam.org/doi/abs/10.1137/0601031>. Publisher: Society for Industrial and Applied Mathematics.
- Gergely Neu and Nneka Okolo. Dealing with unbounded gradients in stochastic saddle-point optimization, June 2024. URL <http://arxiv.org/abs/2402.13903>. arXiv:2402.13903 [cs, math, stat] version: 2.
- Charles Chapman Pugh. *Real mathematical analysis*. Undergraduate texts in mathematics. Springer, Cham Heidelberg, 2. ed edition, 2015. ISBN 978-3-319-17770-0.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley, 1 edition, April 1994. ISBN 978-0-471-61977-2 978-0-470-31688-7. doi: 10.1002/9780470316887. URL <https://onlinelibrary.wiley.com/doi/book/10.1002/9780470316887>.
- P. J. Schweitzer and A. Federgruen. The Functional Equations of Undiscounted Markov Renewal Programming. *Mathematics of Operations Research*, 3(4):308–321, November 1978. ISSN 0364-765X. doi: 10.1287/moor.3.4.308. URL <https://pubsonline.informs.org/doi/10.1287/moor.3.4.308>. Publisher: INFORMS.
- Aaron Sidford, Mengdi Wang, Xian Wu, Lin Yang, and Yinyu Ye. Near-Optimal Time and Sample Complexities for Solving Markov Decision Processes with a Generative Model. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/bb03e43ffe34eeb242a2ee4a4f125e56-Abstract.html>.
- Satinder P. Singh and Richard C. Yee. An Upper Bound on the Loss from Approximate Optimal-Value Functions. *Machine Learning*, 16(3):227–233, September 1994. ISSN 1573-0565. doi: 10.1023/A:1022693225949. URL <https://doi.org/10.1023/A:1022693225949>.
- Adrienne Tuynman, Rémy Degenne, and Emilie Kaufmann. Finding good policies in average-reward Markov Decision Processes without prior knowledge, May 2024. URL <http://arxiv.org/abs/2405.17108>. arXiv:2405.17108 [cs].
- Jinghan Wang, Mengdi Wang, and Lin F. Yang. Near Sample-Optimal Reduction-based Policy Learning for Average Reward MDP, December 2022. URL <http://arxiv.org/abs/2212.00603>. arXiv:2212.00603 [cs].

- Shengbo Wang, Jose Blanchet, and Peter Glynn. Optimal Sample Complexity of Reinforcement Learning for Mixing Discounted Markov Decision Processes, February 2023a. URL <https://arxiv.org/abs/2302.07477v3>.
- Shengbo Wang, Jose Blanchet, and Peter Glynn. Optimal Sample Complexity for Average Reward Markov Decision Processes, October 2023b. URL <https://arxiv.org/abs/2310.08833v1>.
- Chen-Yu Wei, Mehdi Jafarnia-Jahromi, Haipeng Luo, Hiteshi Sharma, and Rahul Jain. Model-free Reinforcement Learning in Infinite-horizon Average-reward Markov Decision Processes, February 2020. URL <http://arxiv.org/abs/1910.07072>. arXiv:1910.07072 [cs, stat].
- Ming Yin, Wenjing Chen, Mengdi Wang, and Yu-Xiang Wang. Offline Stochastic Shortest Path: Learning, Evaluation and Towards Optimality. In *The 38th Conference on Uncertainty in Artificial Intelligence*, May 2022. URL <https://openreview.net/forum?id=HtW4EdIsqIq>.
- Zihan Zhang and Qiaomin Xie. Sharper Model-free Reinforcement Learning for Average-reward Markov Decision Processes, June 2023. URL <http://arxiv.org/abs/2306.16394>. arXiv:2306.16394 [cs].
- Matthew Zurek and Yudong Chen. Span-Based Optimal Sample Complexity for Weakly Communicating and General Average Reward MDPs, March 2024. URL <https://arxiv.org/abs/2403.11477v2>.

A Additional Notation and Guide to Appendices

A.1 Additional Notation

Here we provide definitions for some miscellaneous additional notation used within the appendices. Letting $x \in \mathbb{R}^S$, we define a next-state transition variance vector with respect to P , $\mathbb{V}_P[x] \in \mathbb{R}^{SA}$, as $(\mathbb{V}_P[x])_{sa} := \sum_{s' \in \mathcal{S}} (P)_{sa,s'} [x(s') - \sum_{s''} (P_{sa,s''} x(s''))]^2$. For a policy π , we also define a policy-specific version $\mathbb{V}_{P_\pi}[x] \in \mathbb{R}^S$ as $(\mathbb{V}_{P_\pi}[x])_s := \sum_{s' \in \mathcal{S}} (P_\pi)_{s,s'} [x(s') - \sum_{s''} (P_\pi)_{s,s''} x(s'')]^2$. For any policy π , we define the policy matrix $M^\pi \in \mathbb{R}^{S \times SA}$ by $M^\pi_{s,sa} = \pi(a | s)$, and $M^\pi_{s,s'a} = 0$ if $s \neq s'$. We also define the maximization operator $M : \mathbb{R}^{SA} \rightarrow \mathbb{R}^S$ by $M(x)_s = \max_a x_{sa}$. We also note that for any $x \in \mathbb{R}^{SA}$ and any policy π , $M(x) \geq M^\pi(x)$. We say a policy π is greedy with respect to a vector $x \in \mathbb{R}^{SA}$ if $M(x) = M^\pi(x)$. For any transition matrix P and policy π , we define the deviation matrix H_{P_π} as the Drazin inverse of $I - P_\pi$ (see [Puterman, 1994, Appendix A] for its basic properties). Given a DMDP (P, r, γ) , we define the Q -function of policy π as $Q_\gamma^\pi = r + \gamma P V_\gamma^\pi$, and we define the optimal Q -function as $Q_\gamma^* = r + \gamma P V_\gamma^*$. Note that for any policy π , we have that $V_\gamma^\pi = (I - \gamma P_\pi)^{-1} r_\pi$. We let $\|B\|_{\infty \rightarrow \infty}$ denote the $\|\cdot\|_\infty$ to $\|\cdot\|_\infty$ operator norm of a matrix B , and we note that this is equal to the maximum of the ℓ^1 -norms of the rows of B . In particular, if B is a stochastic matrix (all rows are probability distributions) then $\|B\|_{\infty \rightarrow \infty} = 1$.

A.2 Guide to Appendices

Now we provide an outline of the appendices. In Appendix B we prove our main results for the DMDP plug-in method, Theorems 9 and 10. This section is further split into Subsection B.1, where we show a deterministic error decomposition (which can be understood as a DMDP version of the arguments sketched in Section 4), Subsection B.2, where we check the concentration inequalities that are required for this error decomposition, and Subsections B.3 and B.4, where we complete the proofs of Theorems 9 and 10, respectively.

In Appendix C we prove our main results for the AMDP plug-in method, Theorems 1, 2, 3, 6, and 8. This section is likewise split into further subsections. Subsection C.1 contains useful lemmas, including many properties of the anchoring technique. Subsection C.2 contains a proof of Theorem 6, which can be shown as a consequence of the DMDP result Theorem 10. Analogous to the proofs of the DMDP results, Subsection C.3 contains a deterministic error decomposition (the formal version of the sketched proof techniques) and Subsection C.4 checks the required concentration inequalities. We then complete the proofs of Theorems 2, 3, 8, and 1 in Subsections C.5, C.6, C.7, and C.8, respectively.

In Appendix D we provide proofs which lead to our corollaries for the diameter- and mixing-based settings, proving Lemma 4 in Subsection D.1, proving the diameter-based Corollary 5 in Subsection D.2, and proving the mixing-based Corollary 7 in Subsection D.3.

In Appendix E we provide additional DMDP-related results, Theorem 11, Lemma 12, and Theorem 13.

In Appendix F we provide guarantees for the DMDP reduction approach with effective horizon n , using the close connection to the anchored AMDP plugin approach.

Finally, in Appendix G we provide the proof of Theorem 14 on the impossibility of obtaining a purely $\|h^*\|_{\text{span}}$ -based complexity with the plug-in approach.

B Proofs of DMDP Theorems

In this section we prove Theorems 9 and 10.

B.1 Higher-order variance bounds

Lemma 15. Fix $k \geq 0$ and let $\bar{V} = V^\pi - (\min_s V^\pi(s)) \mathbf{1}$. If the inequality

$$\left| \left(\hat{P}_\pi - P_\pi \right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_\pi} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_\infty + 1)^{2^k} \mathbf{1} \quad (10)$$

holds for some $\alpha \in \mathbb{R}$, then

$$\begin{aligned}
& \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-k}} \left\| \gamma^{2^k} (I - \gamma^{2^k} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^k} \right\|_\infty^{2^{-k}} \\
& \leq \frac{4\alpha}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1) + \frac{2}{1-\gamma} \left(\frac{2\alpha}{n} (\|\bar{V}\|_\infty + 1) \right)^{1-2^{-(k+1)}} \\
& \quad + \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-(k+1)}} \left\| \gamma^{2^{k+1}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^{k+1}} \right\|_\infty^{2^{-(k+1)}}.
\end{aligned}$$

We note that this lemma is purely algebraic, and the same statement holds if we swap all appearances of P and V^π with $\hat{P}$ and $\hat{V}^\pi$ (respectively), both within the assumption and conclusion of the lemma.

Proof. Since $(I - \gamma^{2^k} \hat{P}_\pi)^{-1}$ is elementwise non-negative and using the assumption of the lemma,

$$\begin{aligned}
\left| (I - \gamma^{2^k} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^k} \right| & \leq (I - \gamma^{2^k} \hat{P}_\pi)^{-1} \left| (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^k} \right| \\
& \leq \sqrt{\frac{\alpha}{n}} (I - \gamma^{2^k} \hat{P}_\pi)^{-1} \sqrt{\mathbb{V}_{P_\pi} [(\bar{V})^{\circ 2^k}]} + \frac{\alpha \cdot 2^k}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1)^{2^k} \mathbf{1} \quad (11)
\end{aligned}$$

using the fact that $(I - \gamma^{2^k} \hat{P}_\pi)^{-1} \mathbf{1} = \frac{1}{1-\gamma^{2^k}} \mathbf{1} \leq \frac{1}{1-\gamma} \mathbf{1}$. Therefore using Jensen's inequality (since each row of $(1 - \gamma^{2^k})(I - \gamma^{2^k} \hat{P}_\pi)^{-1}$ is a probability distribution), and then using the elementary inequality $\|(I - \beta P)^{-1} x\|_\infty \leq 2 \|(I - \beta^2 P)^{-1} x\|_\infty$ for $\beta \in (0, 1)$ [Agarwal et al., 2020], we have

$$\begin{aligned}
\left\| (I - \gamma^{2^k} \hat{P}_\pi)^{-1} \sqrt{\mathbb{V}_{P_\pi} [(\bar{V})^{\circ 2^k}]} \right\|_\infty & \leq \frac{1}{\sqrt{1-\gamma^{2^k}}} \sqrt{\left\| (I - \gamma^{2^k} \hat{P}_\pi)^{-1} \mathbb{V}_{P_\pi} [(\bar{V})^{\circ 2^k}] \right\|_\infty} \\
& \leq \frac{\sqrt{2}}{\sqrt{1-\gamma^{2^k}}} \sqrt{\left\| (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \mathbb{V}_{P_\pi} [(\bar{V})^{\circ 2^k}] \right\|_\infty} \\
& \leq \frac{\sqrt{2}}{\sqrt{1-\gamma}} \sqrt{\left\| (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \mathbb{V}_{P_\pi} [(\bar{V})^{\circ 2^k}] \right\|_\infty}. \quad (12)
\end{aligned}$$

Note that $(I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \mathbb{V}_{P_\pi} [(\bar{V})^{\circ 2^k}] \geq 0$ so it suffices to upper bound $(I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \mathbb{V}_{P_\pi} [(\bar{V})^{\circ 2^k}]$ elementwise.

Abbreviating $\nu = (\min_s V^\pi(s))(s)$, by the Bellman equation for V^π we have

$$\begin{aligned}
P_\pi \bar{V} &= P_\pi (V^\pi - \nu \mathbf{1}) \\
&= P_\pi V^\pi - \nu \mathbf{1} \\
&= \frac{1}{\gamma} (V^\pi - r_\pi) - \nu \mathbf{1} \\
&= \frac{1}{\gamma} (V^\pi - r_\pi) + \frac{(1-\gamma) - 1}{\gamma} \nu \mathbf{1} \\
&= \frac{1}{\gamma} (V^\pi - \nu \mathbf{1} - r_\pi) + \frac{1-\gamma}{\gamma} \nu \mathbf{1} \\
&= \frac{1}{\gamma} (\bar{V} - r_\pi + (1-\gamma) \nu \mathbf{1}).
\end{aligned}$$

Using this, we can calculate that

$$\begin{aligned}
\mathbb{V}_{P_\pi} \left[(\bar{V})^{\circ 2^k} \right] &= P_\pi (\bar{V})^{\circ 2^{k+1}} - \left(P_\pi (\bar{V})^{\circ 2^k} \right)^{\circ 2} \\
&\leq P_\pi (\bar{V})^{\circ 2^{k+1}} - \left((P_\pi \bar{V})^{\circ 2^k} \right)^{\circ 2} \\
&= P_\pi (\bar{V})^{\circ 2^{k+1}} - (P_\pi \bar{V})^{\circ 2^{k+1}} \\
&= P_\pi (\bar{V})^{\circ 2^{k+1}} - \frac{1}{\gamma^{2^{k+1}}} (\bar{V} - r_\pi + (1 - \gamma)\nu \mathbf{1})^{\circ 2^{k+1}} \\
&= \frac{1}{\gamma^{2^{k+1}}} \left(\gamma^{2^{k+1}} P_\pi - I \right) (\bar{V})^{\circ 2^{k+1}} - \frac{1}{\gamma^{2^{k+1}}} \left((\bar{V} - r_\pi + (1 - \gamma)\nu \mathbf{1})^{\circ 2^{k+1}} - (\bar{V})^{\circ 2^{k+1}} \right) \\
&\leq \frac{1}{\gamma^{2^{k+1}}} \left(\gamma^{2^{k+1}} P_\pi - I \right) (\bar{V})^{\circ 2^{k+1}} + \frac{2^{2^{k+1}}}{\gamma^{2^{k+1}}} \max\{\|\bar{V}\|_\infty, 1\}^{2^{k+1}-1} \mathbf{1} \\
&\leq \frac{1}{\gamma^{2^{k+1}}} \left(\gamma^{2^{k+1}} P_\pi - I \right) (\bar{V})^{\circ 2^{k+1}} + \frac{2^{2^{k+1}}}{\gamma^{2^{k+1}}} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1} \mathbf{1}
\end{aligned}$$

where we used the fact that $P_\pi (\bar{V})^{\circ 2^k} \geq (P_\pi \bar{V})^{\circ 2^k}$ by Jensen's inequality (applied to each row), and then the fact that

$$(\bar{V} - (r_\pi - (1 - \gamma)\nu \mathbf{1}))^{\circ 2^{k+1}} - (\bar{V})^{\circ 2^{k+1}}$$

contains $< 2^{2^{k+1}}$ terms, all of which have magnitude bounded by $\max\{\|\bar{V}\|_\infty, 1\}^{2^{k+1}-1}$ since $\|r_\pi - (1 - \gamma)\nu\|_\infty \leq 1$ (because $r_\pi \in [0, 1]$ elementwise and also $(1 - \gamma)\nu \in [0, 1]$ since $0 \leq \nu \leq \frac{1}{1-\gamma}$). Plugging this into the RHS of (12), we obtain

$$\begin{aligned}
&(I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \mathbb{V}_{P_\pi} \left[(\bar{V})^{\circ 2^k} \right] \\
&\leq \frac{1}{\gamma^{2^{k+1}}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \left(\gamma^{2^{k+1}} P_\pi - I \right) (\bar{V})^{\circ 2^{k+1}} + \frac{2^{2^{k+1}}}{\gamma^{2^{k+1}}} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \mathbf{1} \\
&= \frac{1}{\gamma^{2^{k+1}}} (I - \gamma^{2^{k+1}} P_\pi)^{-1} \left(\gamma^{2^{k+1}} P_\pi - I \right) (\bar{V})^{\circ 2^{k+1}} \\
&\quad + \frac{1}{\gamma^{2^{k+1}}} \left((I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} - (I - \gamma^{2^{k+1}} P_\pi)^{-1} \right) \left(\gamma^{2^{k+1}} P_\pi - I \right) (\bar{V})^{\circ 2^{k+1}} \\
&\quad + \frac{2^{2^{k+1}}}{(1 - \gamma^{2^{k+1}}) \gamma^{2^{k+1}}} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1} \mathbf{1} \\
&= -\frac{1}{\gamma^{2^{k+1}}} (\bar{V})^{\circ 2^{k+1}} + \frac{2^{2^{k+1}}}{(1 - \gamma^{2^{k+1}}) \gamma^{2^{k+1}}} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1} \mathbf{1} \\
&\quad + \frac{1}{\gamma^{2^{k+1}} \gamma^{2^{k+1}}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \left(\hat{P}_\pi - P_\pi \right) (I - \gamma^{2^{k+1}} P_\pi)^{-1} \left(\gamma^{2^{k+1}} P_\pi - I \right) (\bar{V})^{\circ 2^{k+1}} \\
&= -\frac{1}{\gamma^{2^{k+1}}} (\bar{V})^{\circ 2^{k+1}} + \frac{2^{2^{k+1}}}{(1 - \gamma^{2^{k+1}}) \gamma^{2^{k+1}}} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1} \mathbf{1} \\
&\quad - \frac{1}{\gamma^{2^{k+1}} \gamma^{2^{k+1}}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \left(\hat{P}_\pi - P_\pi \right) (\bar{V})^{\circ 2^{k+1}} \\
&\leq \frac{1}{\gamma^{2^{k+1}}} \left(\frac{2^{2^{k+1}}}{1 - \gamma} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1} + \left\| \gamma^{2^{k+1}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} \left(\hat{P}_\pi - P_\pi \right) (\bar{V})^{\circ 2^{k+1}} \right\|_\infty \right) \mathbf{1}.
\end{aligned}$$

Thus inequality (12), combined with the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, yields that

$$\left\| (I - \gamma^{2^k} \hat{P}_\pi)^{-1} \sqrt{\mathbb{V}_{P_\pi}[(\bar{V})^{\circ 2^k}]} \right\|_\infty \quad (13)$$

$$\begin{aligned} &\leq \frac{\sqrt{2}}{\sqrt{1-\gamma}} \sqrt{\frac{1}{\gamma^{2^{k+1}}}} \sqrt{\frac{2^{2^{k+1}}}{1-\gamma} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1}} \\ &\quad + \frac{\sqrt{2}}{\sqrt{1-\gamma}} \sqrt{\frac{1}{\gamma^{2^{k+1}}}} \sqrt{\left\| \gamma^{2^{k+1}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^{k+1}} \right\|_\infty} \end{aligned} \quad (14)$$

$$\begin{aligned} &= \frac{\sqrt{2}}{\sqrt{1-\gamma}} \frac{1}{\gamma^{2^k}} \sqrt{\frac{2^{2^{k+1}}}{1-\gamma} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1}} \\ &\quad + \frac{\sqrt{2}}{\sqrt{1-\gamma}} \frac{1}{\gamma^{2^k}} \sqrt{\left\| \gamma^{2^{k+1}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^{k+1}} \right\|_\infty}. \end{aligned} \quad (15)$$

Combining this with inequality (11) gives

$$\begin{aligned} &\left\| (I - \gamma^{2^k} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^k} \right\|_\infty \\ &\leq \sqrt{\frac{\alpha}{n}} \left\| (I - \gamma^{2^k} \hat{P}_\pi)^{-1} \sqrt{\mathbb{V}_{P_\pi}[(\bar{V})^{\circ 2^k}]} \right\|_\infty + \frac{\alpha \cdot 2^k}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1)^{2^k} \\ &\leq \frac{\alpha \cdot 2^k}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1)^{2^k} + \sqrt{\frac{\alpha}{n}} \frac{\sqrt{2}}{\sqrt{1-\gamma}} \frac{1}{\gamma^{2^k}} \sqrt{\frac{2^{2^{k+1}}}{1-\gamma} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1}} \\ &\quad + \sqrt{\frac{\alpha}{n}} \frac{\sqrt{2}}{\sqrt{1-\gamma}} \frac{1}{\gamma^{2^k}} \sqrt{\left\| \gamma^{2^{k+1}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^{k+1}} \right\|_\infty} \end{aligned}$$

so using the fact that $(a+b)^{2^{-k}} \leq a^{2^{-k}} + b^{2^{-k}}$ and that $\gamma^{2^k} < 1$,

$$\begin{aligned} &\left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-k}} \left\| \gamma^{2^k} (I - \gamma^{2^k} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^k} \right\|_\infty^{2^{-k}} \\ &\leq \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-k}} \left(\frac{\alpha \cdot 2^k}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1)^{2^k} \right)^{2^{-k}} \\ &\quad + \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-k}} \left(\sqrt{\frac{\alpha}{n}} \frac{\sqrt{2}}{\sqrt{1-\gamma}} \sqrt{\frac{2^{2^{k+1}}}{1-\gamma} (\|\bar{V}\|_\infty + 1)^{2^{k+1}-1}} \right)^{2^{-k}} \\ &\quad + \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-k}} \left(\sqrt{\frac{\alpha}{n}} \frac{\sqrt{2}}{\sqrt{1-\gamma}} \sqrt{\left\| \gamma^{2^{k+1}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^{k+1}} \right\|_\infty} \right)^{2^{-k}} \\ &= \frac{2^{1-2^{-k}+k2^{-k}} \alpha}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1) + \frac{2}{1-\gamma} \left(\frac{2\alpha}{n} (\|\bar{V}\|_\infty + 1) \right)^{1-2^{-(k+1)}} \\ &\quad + \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-(k+1)}} \left\| \gamma^{2^{k+1}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^{k+1}} \right\|_\infty^{2^{-(k+1)}} \\ &\leq \frac{4\alpha}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1) + \frac{2}{1-\gamma} \left(\frac{2\alpha}{n} (\|\bar{V}\|_\infty + 1) \right)^{1-2^{-(k+1)}} \\ &\quad + \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-(k+1)}} \left\| \gamma^{2^{k+1}} (I - \gamma^{2^{k+1}} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^{k+1}} \right\|_\infty^{2^{-(k+1)}} \end{aligned}$$

as desired. $\square$

Lemma 16. Let $\bar{V} = V^\pi - (\min_s V^\pi(s)) \mathbf{1}$ and $\ell = \lceil \log_2 \log_2 (\|\bar{V}\|_\infty + 4) \rceil$. Suppose that for some $\alpha \in \mathbb{R}$, the inequalities

$$\left| \left(\hat{P}_\pi - P_\pi \right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_\pi} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_\infty + 1)^{2^k} \mathbf{1}$$

hold for all $k = 0, \dots, \ell$. Then

$$\|\hat{V}^\pi - V^\pi\|_\infty \leq \frac{4(\ell+1)\alpha}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1) + \frac{2(\ell+1)}{1-\gamma} \sqrt{\frac{2\alpha (\|\bar{V}\|_\infty + 1)}{n}}.$$

Proof. First we give a weaker but non-recursive bound which can be used on the final term. Note that

$$\begin{aligned} \left| \left(\hat{P}_\pi - P_\pi \right) (\bar{V})^{\circ 2^\ell} \right| &\leq \sqrt{\frac{\alpha \mathbb{V}_{P_\pi} \left[(\bar{V})^{\circ 2^\ell} \right]}{n}} + \frac{\alpha \cdot 2^\ell}{n} (\|\bar{V}\|_\infty + 1)^{2^\ell} \mathbf{1} \\ &\leq \sqrt{\frac{\alpha}{n}} \|\bar{V}\|_\infty^{2^\ell} \mathbf{1} + \frac{\alpha \cdot 2^\ell}{n} (\|\bar{V}\|_\infty + 1)^{2^\ell} \mathbf{1} \end{aligned}$$

so

$$\begin{aligned} \left\| \gamma^{2^\ell} (I - \gamma^{2^\ell} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^\ell} \right\|_\infty &\leq \gamma^{2^\ell} \left\| (I - \gamma^{2^\ell} \hat{P}_\pi)^{-1} \left| (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^\ell} \right| \right\|_\infty \\ &\leq \frac{1}{1-\gamma} \sqrt{\frac{\alpha}{n}} \|\bar{V}\|_\infty^{2^\ell} + \frac{1}{1-\gamma} \frac{\alpha \cdot 2^\ell}{n} (\|\bar{V}\|_\infty + 1)^{2^\ell} \\ &\leq \frac{1}{1-\gamma} \sqrt{\frac{\alpha}{n}} (\|\bar{V}\|_\infty + 1)^{2^\ell} + \frac{1}{1-\gamma} \frac{\alpha \cdot 2^\ell}{n} (\|\bar{V}\|_\infty + 1)^{2^\ell} \end{aligned}$$

since $(I - \gamma^{2^\ell} \hat{P}_\pi)^{-1} \mathbf{1} = \frac{1}{1-\gamma^{2^\ell}} \mathbf{1} \leq \frac{1}{1-\gamma} \mathbf{1}$ and $\gamma < 1$. Therefore

$$\begin{aligned} &\left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-\ell}} \left\| \gamma^{2^\ell} (I - \gamma^{2^\ell} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^\ell} \right\|_\infty^{2^{-\ell}} \\ &\leq \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-\ell}} \left(\frac{1}{1-\gamma} \sqrt{\frac{\alpha}{n}} (\|\bar{V}\|_\infty + 1)^{2^\ell} \right)^{2^{-\ell}} + \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-\ell}} \left(\frac{1}{1-\gamma} \frac{\alpha \cdot 2^\ell}{n} (\|\bar{V}\|_\infty + 1)^{2^\ell} \right)^{2^{-\ell}} \\ &\leq \frac{2}{1-\gamma} \left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} (\|\bar{V}\|_\infty + 1) + \frac{1}{1-\gamma} \frac{4\alpha}{n} (\|\bar{V}\|_\infty + 1). \end{aligned} \tag{16}$$

Note that

$$\begin{aligned} \hat{V}^\pi - V^\pi &= \gamma (I - \gamma \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) V^\pi \\ &= \gamma (I - \gamma \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) \left(V^\pi - \left(\min_s V^\pi(s) \right) \mathbf{1} \right) \end{aligned}$$

since $(\hat{P}_\pi - P_\pi) \mathbf{1} = 0$. Using Lemma 15 ℓ times and using the above bound (16) for the final term, we

obtain

$$\begin{aligned}
\|\hat{V}^\pi - V^\pi\|_\infty &= \left\| \gamma(I - \gamma\hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) \left(V^\pi - \left(\min_s V^\pi(s) \right) \mathbf{1} \right) \right\|_\infty \\
&\leq \sum_{k=0}^{\ell-1} \left(\frac{4\alpha}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1) + \frac{2}{1-\gamma} \left(\frac{2\alpha}{n} (\|\bar{V}\|_\infty + 1) \right)^{1-2^{-(k+1)}} \right) \\
&\quad + \left(\frac{2\alpha}{(1-\gamma)n} \right)^{1-2^{-\ell}} \left\| \gamma^{2^\ell} (I - \gamma^{2^\ell} \hat{P}_\pi)^{-1} (\hat{P}_\pi - P_\pi) (\bar{V})^{\circ 2^\ell} \right\|_\infty^{2^{-\ell}} \\
&\leq \sum_{k=0}^{\ell-1} \left(\frac{4\alpha}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1) + \frac{2}{1-\gamma} \left(\frac{2\alpha}{n} (\|\bar{V}\|_\infty + 1) \right)^{1-2^{-(k+1)}} \right) \\
&\quad + \frac{2}{1-\gamma} \left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} (\|\bar{V}\|_\infty + 1) + \frac{1}{1-\gamma} \frac{4\alpha}{n} (\|\bar{V}\|_\infty + 1) \\
&\leq \frac{4(\ell+1)\alpha}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1) + \frac{2\ell}{1-\gamma} \sqrt{\frac{2\alpha (\|\bar{V}\|_\infty + 1)}{n}} + \frac{2}{1-\gamma} \left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} (\|\bar{V}\|_\infty + 1) \quad (17)
\end{aligned}$$

where we assume that $\frac{2\alpha}{n} (\|\bar{V}\|_\infty + 1) < 1$ in the final inequality step, so that the $k=0$ term is the largest term in $\sum_{k=0}^{\ell-1} \left(\frac{2\alpha}{n} (\|\bar{V}\|_\infty + 1) \right)^{1-2^{-(k+1)}}$. Now we check that $\ell = \lceil \log_2 \log_2 (\|\bar{V}\|_\infty + 4) \rceil$ is sufficiently large so that the rightmost term in the RHS of (17) is smaller than $\frac{2}{1-\gamma} \sqrt{\frac{2\alpha (\|\bar{V}\|_\infty + 1)}{n}}$. We have

$$\begin{aligned}
\left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} (\|\bar{V}\|_\infty + 1) &= \left(\frac{\alpha}{n} (\|\bar{V}\|_\infty + 1) \right)^{1-2^{-(\ell+1)}} (\|\bar{V}\|_\infty + 1)^{2^{-(\ell+1)}} \\
&\leq \left(\frac{\alpha}{n} (\|\bar{V}\|_\infty + 1) \right)^{1/2} (\|\bar{V}\|_\infty + 1)^{2^{-(\ell+1)}}
\end{aligned}$$

since $\ell \geq 0$. Furthermore we have the equivalences

$$\begin{aligned}
(\|\bar{V}\|_\infty + 1)^{2^{-(\ell+1)}} &\leq \sqrt{2} \iff 2^{-(\ell+1)} \log_2 (\|\bar{V}\|_\infty + 1) \leq 1/2 \\
&\iff \log_2 (\|\bar{V}\|_\infty + 1) \leq 2^\ell \\
&\iff \ell \geq \log_2 \log_2 (\|\bar{V}\|_\infty + 1)
\end{aligned}$$

and this RHS of the final inequality is smaller than our choice of ℓ . Thus our choice of ℓ indeed guarantees that

$$\begin{aligned}
\frac{2}{1-\gamma} \left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} (\|\bar{V}\|_\infty + 1) &\leq \frac{2}{1-\gamma} \sqrt{\frac{\alpha (\|\bar{V}\|_\infty + 1)}{n}} (\|\bar{V}\|_\infty + 1)^{2^{-(\ell+1)}} \\
&\leq \frac{2}{1-\gamma} \sqrt{\frac{2\alpha (\|\bar{V}\|_\infty + 1)}{n}}.
\end{aligned}$$

Plugging this into (17), we conclude that

$$\|\hat{V}^\pi - V^\pi\|_\infty \leq \frac{4(\ell+1)\alpha}{(1-\gamma)n} (\|\bar{V}\|_\infty + 1) + \frac{2(\ell+1)}{1-\gamma} \sqrt{\frac{2\alpha (\|\bar{V}\|_\infty + 1)}{n}}$$

as desired. $\square$

B.2 Bernstein-like inequalities

Now we show that each of the different versions of the Bernstein-like inequality (10) which are needed for our different results hold with high probability.

For an optimal γ -discounted policy π_γ^* , which is fixed and independent of $\hat{P}$, showing the Bernstein-like inequality (10) follows almost immediately from Bernstein's inequality.

Lemma 17. Let $\bar{V} = V^* - (\min_s V^*(s)) \mathbf{1}$. With probability at least $1 - \delta$, we have that for all $k = 0, \dots, \lceil \log_2 \log_2 (\|V^*\|_{\text{span}} + 4) \rceil$,

$$\left| \left(\hat{P}_{\pi_\gamma^*} - P_{\pi_\gamma^*} \right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_{\pi_\gamma^*}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha}{n} \|\bar{V}\|_\infty^{2^k} \mathbf{1}$$

where $\alpha = 2 \log \left(\frac{6S \log_2 \log_2 (\|V^*\|_{\text{span}} + 4)}{\delta} \right)$.

Proof. First we fix $k \in \{0, \dots, \lceil \log_2 \log_2 (\|V^*\|_{\text{span}} + 4) \rceil\}$. Fix a state $s \in \mathcal{S}$ and note that we defined π_γ^* to be a deterministic policy, so we may treat $\pi_\gamma^*(s)$ as an element of $\mathcal{A}$. By applying Bernstein's inequality (e.g. [Maurer and Pontil 2009](#), Theorem 3), we have that with probability at least $1 - 2\delta'$,

$$\begin{aligned} \left| \left(\hat{P}_{s, \pi_\gamma^*(s)} - P_{s, \pi_\gamma^*(s)} \right) (\bar{V})^{\circ 2^k} \right| &\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{s, \pi_\gamma^*(s)}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\log \left(\frac{1}{\delta'} \right)}{3n} \left\| (\bar{V})^{\circ 2^k} \right\|_\infty \\ &= \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{s, \pi_\gamma^*(s)}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\log \left(\frac{1}{\delta'} \right)}{3n} \|\bar{V}\|_\infty^{2^k} \\ &\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{s, \pi_\gamma^*(s)}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right)}{n} \|\bar{V}\|_\infty^{2^k}. \end{aligned}$$

Now taking a union bound over $s \in \mathcal{S}$, we have that the above inequality holds for all s simultaneously with probability at least $1 - 2S\delta'$, in which case we have (elementwise)

$$\left| \left(\hat{P}_{\pi_\gamma^*} - P_{\pi_\gamma^*} \right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{\pi_\gamma^*}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right)}{n} \|\bar{V}\|_\infty^{2^k} \mathbf{1}.$$

Finally, taking another union bound over all possible values of k , of which there are

$$1 + \lceil \log_2 \log_2 (\|V^*\|_{\text{span}} + 4) \rceil \leq 2 + \log_2 \log_2 (\|V^*\|_{\text{span}} + 4) \leq 3 \log_2 \log_2 (\|V^*\|_{\text{span}} + 4),$$

(since $\log_2 \log_2 (\|V^*\|_{\text{span}} + 4) \geq \log_2 \log_2 4 \geq 1$) and setting $\delta' = \frac{\delta}{6S \log_2 \log_2 (\|V^*\|_{\text{span}} + 4)}$, we obtain the desired conclusion. $\square$

For the empirical (near)-optimal policy $\hat{\pi}$, which is statistically dependent on $\hat{P}$, this requires more effort, in particular the use of the absorbing MDP construction pioneered by [Agarwal et al. \[2020\]](#) to decouple the statistical dependency. We first present their construction.

Theorem 18 ([Agarwal et al. \[2020\]](#)). *There exists a collection of random variables $\hat{V}_{s,u}^*$ for $s \in \mathcal{S}$ and $u \in [0, 1]$ such that*

1. *For all $s \in \mathcal{S}, u \in [0, 1]$, $\hat{V}_{s,u}^*$ is independent from all of the random variables $S_{s,a}^i$ for all $i = 1, \dots, n$ and all $a \in \mathcal{A}$ (recall these are the n observed transitions from state-action pair (s, a) which are used to form $\hat{P}$).*
2. *For all $s \in \mathcal{S}, u, u' \in [0, 1]$, $\left\| \hat{V}_{s,u}^* - \hat{V}_{s,u'}^* \right\|_\infty \leq \frac{|u - u'|}{1 - \gamma}$.*
3. *Letting $u^*(s) = (1 - \gamma) \hat{V}^*(s)$, $\hat{V}_{s, u^*(s)}^* = \hat{V}^*$.*

As a consequence, there exists a finite set U of $\left\lceil \frac{1}{2(1-\gamma)\varepsilon} \right\rceil$ equally-spaced points in $[0, 1]$ such that for all $s \in \mathcal{S}$, there exists $u \in U$ such that $\left\| \hat{V}_{s,u}^* - \hat{V}^* \right\|_\infty \leq \varepsilon$.

We also note a useful elementary inequality.

Lemma 19. *For any natural number n ,*

$$|a^n - b^n| \leq n|a - b| (\max\{|a|, |b|\})^{n-1}.$$

Proof. This follows immediately from the algebraic identity

$$a^n - b^n = (a - b)(a^{n-1} + a^{n-2}b + a^{n-3}b^2 + \dots + b^{n-1}).$$

□

Now we can use the leave-one-out construction of Agarwal et al. [2020] to check a version of the Bernstein-like inequalities.

Lemma 20. *If $n \geq 4$, then with probability at least $1 - \delta$, for all $\hat{\pi}$ which satisfy $\|\hat{V}^{\hat{\pi}} - \hat{V}^{\star}\|_{\infty} \leq \frac{1}{n}$, letting $\bar{V} = \hat{V}^{\hat{\pi}} - \left(\min_s \hat{V}^{\hat{\pi}}(s)\right) \mathbf{1}$, for all $k = 0, \dots, \left\lceil \log_2 \log_2 \left(\|\hat{V}^{\hat{\pi}}\|_{\text{span}} + 4\right) \right\rceil$, we have*

$$\left| \left(\hat{P}_{\hat{\pi}} - P_{\hat{\pi}}\right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{\hat{P}_{\hat{\pi}}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_{\infty} + 1)^{2^k} \mathbf{1}$$

where $\alpha = 16 \log \left(12 \frac{SAn}{(1-\gamma)^2 \delta} \right)$.

Proof. We use the leave-one-out construction from Agarwal et al. [2020]: by Theorem 18, there exists a finite set U with $|U| = \left\lceil \frac{n}{2(1-\gamma)} \right\rceil$ such that for all $s \in \mathcal{S}$, there exists $u \in U$ such that $\|\hat{V}_{s,u}^{\star} - \hat{V}^{\star}\|_{\infty} \leq \frac{1}{n}$. Also note that since $n \geq 2$ by assumption, $1 \leq \frac{n}{2(1-\gamma)}$, so we can bound $|U| \leq \frac{n}{2(1-\gamma)} + 1 \leq \frac{n}{1-\gamma}$. We define, for all $s \in \mathcal{S}$ and $u \in U$, $\bar{V}_{s,u} = \hat{V}_{s,u}^{\star} - \left(\min_{s'} \hat{V}_{s,u}^{\star}(s')\right) \mathbf{1}$.

We note that the quantity $\left\lceil \log_2 \log_2 \left(\|\hat{V}^{\hat{\pi}}\|_{\text{span}} + 4\right) \right\rceil$ appearing in the lemma statement is random, but since $\|\hat{V}^{\hat{\pi}}\|_{\text{span}} \leq \frac{1}{1-\gamma}$, we can simply bound

$$\begin{aligned} \left\lceil \log_2 \log_2 \left(\|\hat{V}^{\hat{\pi}}\|_{\text{span}} + 4\right) \right\rceil &\leq \left\lceil \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4\right) \right\rceil \\ &\leq 1 + \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4\right) \\ &\leq 1 + 2 \frac{1}{1-\gamma} \\ &\leq 3 \frac{1}{1-\gamma} \end{aligned} \tag{18}$$

(using that $\log_2 \log_2(x+4) \leq 2x$ for $x \geq 1$), so we can check the inequality for all values of k up to the upper bound $3 \frac{1}{1-\gamma}$, which is at most $1 + 3 \frac{1}{1-\gamma} \leq 4 \frac{1}{1-\gamma}$ values of k . Therefore for the rest of the proof we will focus on showing the desired conclusion only for some fixed k , from which we can immediately obtain the desired conclusion by taking a union bound and adjusting the failure probability.

For each state s , action a , and $u \in U$, we will use [Maurer and Pontil, 2009, Theorem 10] to show that, with probability at least $1 - \delta'$,

$$\sqrt{\mathbb{V}_{P_{sa}} \left[(\bar{V}_{s,u})^{\circ 2^k} \right]} \leq \sqrt{2 \mathbb{V}_{\hat{P}_{sa}} \left[(\bar{V}_{s,u})^{\circ 2^k} \right]} + 2 \left\| (\bar{V}_{s,u})^{\circ 2^k} \right\|_{\infty} \sqrt{\frac{\log(1/\delta')}{n}}. \tag{19}$$

To match the notation of [Maurer and Pontil, 2009, Theorem 10], let $x_i := (\bar{V}_{s,u})^{\circ 2^k}(S_{s,a}^i)$, that is the value of $(\bar{V}_{s,u})^{\circ 2^k}$ at the state $S_{s,a}^i$, which is the i th sample from the state-action distribution $P(\cdot | s, a)$. [Maurer

and Pontil, 2009, Theorem 10] assumes that $x_i \in [0, 1]$, but we can apply their theorem to the quantities $x'_i := \frac{x_i}{\|(\bar{V}_{s,u})^{\circ 2^k}\|_\infty}$ to obtain that if $n \geq 2$, with probability at least $1 - \delta'$,

$$\sqrt{\mathbb{V}_{P_{sa}} \left[(\bar{V}_{s,u})^{\circ 2^k} \right]} \leq \sqrt{\frac{n}{n-1} \mathbb{V}_{\hat{P}_{sa}} \left[(\bar{V}_{s,u})^{\circ 2^k} \right]} + \|(\bar{V}_{s,u})^{\circ 2^k}\|_{\text{span}} \sqrt{\frac{2 \log(1/\delta')}{n-1}}.$$

This is true because we have

$$\begin{aligned} \frac{n}{n-1} \mathbb{V}_{\hat{P}_{sa}} \left[(\bar{V}_{s,u})^{\circ 2^k} \right] &= \frac{n}{n-1} \left(\hat{P}_{sa} \left((\bar{V}_{s,u})^{\circ 2^k} \right)^{\circ 2} - \left(\hat{P}_{sa} (\bar{V}_{s,u})^{\circ 2^k} \right)^{\circ 2} \right) \\ &= \frac{n}{n-1} \left(\frac{1}{n} \sum_{i=1}^n x_i^2 - \left(\frac{1}{n} \sum_{i=1}^n x_i \right)^2 \right) \\ &= \frac{1}{n-1} \sum_{i=1}^n x_i^2 + \frac{1}{n(n-1)} \left(\sum_{i=1}^n x_i \right) \left(\sum_{j=1}^n x_j \right) \\ &= \frac{1}{n(n-1)} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{x_i^2}{2} + \sum_{i=1}^n \sum_{j=1}^n \frac{x_j^2}{2} \right) + \frac{1}{n(n-1)} \left(\sum_{i=1}^n x_i \right) \left(\sum_{j=1}^n x_j \right) \\ &= \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^n \frac{x_i^2 - 2x_i x_j + x_j^2}{2} \\ &= \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^n \frac{(x_i - x_j)^2}{2} \end{aligned}$$

which is the quantity $V_n(X)$ appearing in [Maurer and Pontil, 2009, Theorem 10]. Also since $n \geq 2$ we have $\frac{n}{n-1} \leq 2$ and $\frac{1}{n-1} \leq 2\frac{1}{n}$, yielding (19). Using Bernstein's inequality (e.g. [Maurer and Pontil, 2009, Theorem 3]), with an additional failure probability of at most $2\delta'$, we have

$$\left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V}_{s,u})^{\circ 2^k} \right| \leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{sa}} \left[(\bar{V}_{s,u})^{\circ 2^k} \right]}{n}} + \frac{\log \left(\frac{1}{\delta'} \right) \left\| (\bar{V}_{s,u})^{\circ 2^k} \right\|_\infty}{3n}. \quad (20)$$

Combining this with (19), we have

$$\left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V}_{s,u})^{\circ 2^k} \right| \leq \sqrt{\frac{4 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{\hat{P}_{sa}} \left[(\bar{V}_{s,u})^{\circ 2^k} \right]}{n}} + \frac{4 \log \left(\frac{1}{\delta'} \right) \left\| (\bar{V}_{s,u})^{\circ 2^k} \right\|_\infty}{n} \quad (21)$$

(using that $\frac{1}{3} + 2\sqrt{2} \leq 4$ for the second term.)

Also $\left\| (\bar{V}_{s,u})^{\circ 2^k} \right\|_\infty = \left\| \bar{V}_{s,u} \right\|_\infty^{2^k}$. Combining this with (21) yields

$$\left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V}_{s,u})^{\circ 2^k} \right| \leq \sqrt{\frac{4 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{\hat{P}_{sa}} \left[(\bar{V}_{s,u})^{\circ 2^k} \right]}{n}} + \frac{4 \log \left(\frac{1}{\delta'} \right)}{n} \left\| \bar{V}_{s,u} \right\|_\infty^{2^k}. \quad (22)$$

From this point, we will take $\delta' = \frac{\delta}{3SA|U|}$ and operate under the event that the above inequality holds for all $s \in S, a \in A, u \in U$, which by the union bound and our choice of δ' has probability at least $1 - \delta$. U is chosen so that there exists some $u \in U$ such that $\left\| \hat{V}_{s,u}^* - \hat{V}^* \right\|_\infty \leq \frac{1}{n}$. Additionally, by assumption $\left\| \hat{V}^{\hat{\pi}} - \hat{V}^* \right\|_\infty \leq \frac{1}{n}$, so by triangle inequality $\left\| \hat{V}^{\hat{\pi}} - \hat{V}_{s,u}^* \right\|_\infty \leq \frac{2}{n}$. Therefore

$$\left\| \bar{V} - \bar{V}_{s,u} \right\|_\infty = \left\| \hat{V}^{\hat{\pi}} - \hat{V}_{s,u}^* + \min_{s'} \hat{V}_{s,u}^*(s') - \min_{s'} \hat{V}^{\hat{\pi}}(s') \right\|_\infty \leq 2 \left\| \hat{V}^{\hat{\pi}} - \hat{V}_{s,u}^* \right\|_\infty \leq \frac{4}{n}.$$

We now use this to obtain a version of (22) but with $\bar{V}$ in place of $\bar{V}_{s,u}$.

First we note that by Lemma 19 (applied elementwise) we have

$$\left\| (\bar{V})^{\circ 2^k} - (\bar{V}_{s,u})^{\circ 2^k} \right\|_{\infty} \leq 2^k \|\bar{V} - \bar{V}_{s,u}\|_{\infty} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k-1} \leq 2^k \frac{4}{n} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k-1} \quad (23)$$

since $\max\{\|\bar{V}\|_{\infty}, \|\bar{V}_{s,u}\|_{\infty}\} \leq \|\bar{V}\|_{\infty} + \frac{4}{n}$ because $\|\bar{V} - \bar{V}_{s,u}\|_{\infty} \leq \frac{4}{n}$. For any s, a , we thus have

$$\begin{aligned} & \left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V})^{\circ 2^k} \right| \\ & \leq \left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V}_{s,u})^{\circ 2^k} \right| + \left| \left(\hat{P}_{sa} - P_{sa} \right) \left((\bar{V})^{\circ 2^k} - (\bar{V}_{s,u})^{\circ 2^k} \right) \right| \\ & \leq \left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V}_{s,u})^{\circ 2^k} \right| + \left\| \hat{P}_{sa} - P_{sa} \right\|_1 \left\| (\bar{V})^{\circ 2^k} - (\bar{V}_{s,u})^{\circ 2^k} \right\|_{\infty} \\ & \leq \left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V}_{s,u})^{\circ 2^k} \right| + 2^k \frac{8}{n} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k-1} \\ & \leq \sqrt{\frac{4 \log\left(\frac{1}{\delta'}\right) \mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V}_{s,u})^{\circ 2^k}\right]}{n}} + \frac{4 \log\left(\frac{1}{\delta'}\right)}{n} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k} + 2^k \frac{8}{n} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k-1} \end{aligned} \quad (24)$$

using the triangle inequality, Holder's inequality, $\left\| \hat{P}_{sa} - P_{sa} \right\|_1 \leq 2$ and (23), and (22) combined with the fact that $\|\bar{V}_{s,u}\|_{\infty} \leq \|\bar{V}\|_{\infty} + \frac{4}{n}$.

We also have that

$$\begin{aligned} \sqrt{\mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V}_{s,u})^{\circ 2^k}\right]} &= \sqrt{\mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V})^{\circ 2^k} + (\bar{V}_{s,u})^{\circ 2^k} - (\bar{V})^{\circ 2^k}\right]} \\ &\leq \sqrt{\mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V})^{\circ 2^k}\right]} + \sqrt{\mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V}_{s,u})^{\circ 2^k} - (\bar{V})^{\circ 2^k}\right]} \\ &\leq \sqrt{\mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V})^{\circ 2^k}\right]} + \left\| (\bar{V})^{\circ 2^k} - (\bar{V}_{s,u})^{\circ 2^k} \right\|_{\infty} \\ &\leq \sqrt{\mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V})^{\circ 2^k}\right]} + 2^k \frac{4}{n} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k-1} \end{aligned}$$

using the triangle inequality for the norm $\sqrt{\mathbb{E}X^2}$, $\mathbb{V}_{\hat{P}_{sa}}[X] \leq \|X\|_{\infty}^2$, and using (23). Plugging this into (24) and simplifying, we have

$$\begin{aligned} \left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V})^{\circ 2^k} \right| &\leq \sqrt{\frac{4 \log\left(\frac{1}{\delta'}\right) \mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V})^{\circ 2^k}\right]}{n}} + \frac{2 \sqrt{\log\left(\frac{1}{\delta'}\right)}}{\sqrt{n}} 2^k \frac{4}{n} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k-1} \\ &\quad + \frac{4 \log\left(\frac{1}{\delta'}\right)}{n} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k} + 2^k \frac{8}{n} \left(\|\bar{V}\|_{\infty} + \frac{4}{n} \right)^{2^k-1} \\ &\leq \sqrt{\frac{4 \log\left(\frac{1}{\delta'}\right) \mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V})^{\circ 2^k}\right]}{n}} + \frac{16 \cdot 2^k \log\left(\frac{1}{\delta'}\right)}{n} (\|\bar{V}\|_{\infty} + 1)^{2^k} \\ &\leq \sqrt{\frac{\alpha \mathbb{V}_{\hat{P}_{sa}}\left[(\bar{V})^{\circ 2^k}\right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_{\infty} + 1)^{2^k} \end{aligned}$$

for $\alpha = 16 \log\left(\frac{1}{\delta'}\right)$ in the final inequality. (For the simplification steps, we have $\log\left(\frac{1}{\delta'}\right) \geq 1$ so $\sqrt{\log\left(\frac{1}{\delta'}\right)} \leq \log\left(\frac{1}{\delta'}\right)$, and, since $n \geq 4$ we have $\|\bar{V}\|_{\infty} + \frac{4}{n} \leq \|\bar{V}\|_{\infty} + 1$, and also $\|\bar{V}\|_{\infty} + 1 \geq 1$. Also since $n \geq 4$,

$\frac{2}{\sqrt{n}} \leq 1$.) Since this holds for all s, a , by definition of the vector $\mathbb{V}_{\hat{P}} \left[(\bar{V})^{\circ 2^k} \right]$ we thus have (elementwise)

$$\left| (\hat{P} - P) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{\hat{P}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_{\infty} + 1)^{2^k} \mathbf{1} \quad (25)$$

and so

$$\begin{aligned} \left| (\hat{P}_{\hat{\pi}} - P_{\hat{\pi}}) (\bar{V})^{\circ 2^k} \right| &= \left| M^{\hat{\pi}} (\hat{P} - P) (\bar{V})^{\circ 2^k} \right| \\ &\leq M^{\hat{\pi}} \left| (\hat{P} - P) (\bar{V})^{\circ 2^k} \right| \\ &\leq M^{\hat{\pi}} \sqrt{\frac{\alpha \mathbb{V}_{\hat{P}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_{\infty} + 1)^{2^k} M^{\hat{\pi}} \mathbf{1} \\ &\leq \sqrt{\frac{\alpha M^{\hat{\pi}} \mathbb{V}_{\hat{P}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_{\infty} + 1)^{2^k} \mathbf{1} \\ &= \sqrt{\frac{\alpha \mathbb{V}_{\hat{P}_{\hat{\pi}}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_{\infty} + 1)^{2^k} \mathbf{1} \end{aligned} \quad (26)$$

where the first inequality is because all entries of $M^{\hat{\pi}}$ are ≥ 0 , and the third inequality is using Jensen's inequality (since the rows of $M^{\hat{\pi}}$ are probability distributions) and using the fact that $M^{\hat{\pi}} \mathbf{1} = \mathbf{1}$ (note $\mathbf{1}$ has different dimensions on each side of this equation). (We note that this step, which replaces all appearances of $P, \hat{P}$ in (25) with $P_{\hat{\pi}}, \hat{P}_{\hat{\pi}}$, could be done for any arbitrary policy π .)

Finally, taking a union bound over all $\leq 4 \frac{1}{1-\gamma}$ values of k and adjusting the failure probabilities so that the overall failure probability is $\leq \delta$, we can upper bound the resulting value of α by

$$16 \log \left(\frac{3SA|U|}{\delta} 4 \frac{1}{1-\gamma} \right) \leq 16 \log \left(\frac{12SAn}{(1-\gamma)^2 \delta} \right). \quad (27)$$

□

The following result summarizes the leave-one-out construction from Li et al. [2020] (specifically it is a direct combination of [Li et al., 2020, Lemma 6] and [Li et al., 2020, Lemma 4]).

Theorem 21 (Li et al. [2020]). *Let $\tilde{r} = r + \Delta$ where $\Delta(s, a) \sim \text{Uniform}[0, \xi]$ independently for each $s \in \mathcal{S}, a \in \mathcal{A}$. Let ${}_p \hat{V}^{\pi}$ denote the value of policy π the discounted MDP with transition matrix $\hat{P}$ and reward $\tilde{r}$, and likewise let ${}_p \hat{V}^{\star}$ denote the optimal value function in this MDP. We also let ${}_p \hat{Q}^{\pi}$ and ${}_p \hat{Q}^{\star}$ be the Q -functions of policy π and of an optimal policy in this MDP. Then*

1. *There exists a family of MDPs $\hat{P}^{(s,a,u)}$ for $s \in \mathcal{S}, a \in \mathcal{A}, u \in U$ (for some finite set U) such that*

$$(a) \quad |U| \leq \frac{24SA^2}{(1-\gamma)^3 \xi \delta}.$$

(b) *For each $s \in \mathcal{S}, a \in \mathcal{A}, u \in U$, the MDP $\hat{P}^{(s,a,u)}$ is independent of all of the random variables $(S_{s,a}^i)_{i=1}^n$ (the n observed transitions from state-action pair s, a).*

2. *With probability at least $1 - \delta$,*

(a) *The optimal policy $\hat{\pi}_p^{\star}$ in the DMDP $(\hat{P}, \tilde{r}, \gamma)$ is unique and is a deterministic policy.*

(b) *For all $s \in \mathcal{S}, a \in \mathcal{A}$ such that $a \neq \hat{\pi}_p^{\star}(s)$,*

$${}_p \hat{Q}^{\star}(s, \hat{\pi}_p^{\star}(s)) - {}_p \hat{Q}^{\star}(s, a) \geq \frac{\xi \delta (1-\gamma)}{3SA^2}.$$

(c) For each $s \in \mathcal{S}, a \in \mathcal{A}$, there exists $u^* \in U$ such that the unique optimal policy in the DMDP $(\hat{P}^{(s,a,u^*)}, \tilde{r}, \gamma)$, which we label $\hat{\pi}_{s,a,u}^*$, is equal to $\hat{\pi}_p^*$.

This construction allows us to check the Bernstein-like inequality (10) but with $V^{\hat{\pi}_p^*}$, the true value function (in the true DMDP (P, r, γ)) of a policy $\hat{\pi}_p^*$ which is optimal in the perturbed empirical DMDP $(\hat{P}, \tilde{r}, \gamma)$.

Lemma 22. Let $\tilde{r} = r + \Delta$ where $\Delta(s, a) \sim \text{Uniform}[0, \xi]$ independently for each $s \in \mathcal{S}, a \in \mathcal{A}$. Let $\hat{\pi}_p^*$ be a policy which is optimal for the DMDP $(\hat{P}, \tilde{r}, \gamma)$. With probability at least $1 - \delta$, letting $\bar{V} = V^{\hat{\pi}_p^*} - (\min_s V^{\hat{\pi}_p^*}(s)) \mathbf{1}$, for all $k = 0, \dots, \left\lceil \log_2 \log_2 \left(\|V^{\hat{\pi}_p^*}\|_{\text{span}} + 4 \right) \right\rceil$, we have

$$\left| \left(\hat{P}_{\hat{\pi}_p^*} - P_{\hat{\pi}_p^*} \right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_{\hat{\pi}_p^*}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_\infty + 1)^{2^k} \mathbf{1}$$

where $\alpha = 2 \log \left(\frac{768 S^2 A^3}{(1-\gamma)^4 \xi \delta^2} \right)$.

Proof. From Theorem 21, the policy $\hat{\pi}_{s,a,u}^*$ is independent of the observed transitions $(S_{s,a}^i)_{i=1}^n$ from state-action pair (s, a) , so in particular the random variable $V^{\hat{\pi}_{s,a,u}^*}$ is independent from $(S_{s,a}^i)_{i=1}^n$ (and so is $(V^{\hat{\pi}_{s,a,u}^*} - (\min_{s'} V^{\hat{\pi}_{s,a,u}^*}(s')) \mathbf{1})^{\circ 2^k}$ for any k). Furthermore, Theorem 21 guarantees that with probability at least $1 - \delta/2$, for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$ there exists some $u^* \in U$ such that $\hat{\pi}_{s,a,u^*}^* = \hat{\pi}_p^*$, which implies that $V^{\hat{\pi}_{s,a,u^*}^*} = V^{\hat{\pi}_p^*}$. Also we have $|U| \leq \frac{48 S A^2}{(1-\gamma)^3 \xi \delta}$. Therefore, letting $\bar{V}_{s,a,u} = V^{\hat{\pi}_{s,a,u}^*} - (\min_{s'} V^{\hat{\pi}_{s,a,u}^*}(s')) \mathbf{1}$, if we check that

$$\left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V}_{s,a,u})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_{sa}} \left[(\bar{V}_{s,a,u})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}_{s,a,u}\|_\infty + 1)^{2^k} \quad (28)$$

for all combinations of $s \in \mathcal{S}, a \in \mathcal{A}, u \in U$, and $k = 0, \dots, \left\lceil \log_2 \log_2 \left(\|V^{\hat{\pi}_p^*}\|_{\text{span}} + 4 \right) \right\rceil$ with probability at least $1 - \delta/2$, then we can combine with Theorem 21 to obtain that

$$\left| \left(\hat{P}_{sa} - P_{sa} \right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_{sa}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_\infty + 1)^{2^k} \quad (29)$$

(for all a, s, k), and consequently that

$$\left| \left(\hat{P}_{\hat{\pi}_p^*} - P_{\hat{\pi}_p^*} \right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_{\hat{\pi}_p^*}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{V}\|_\infty + 1)^{2^k} \mathbf{1}$$

(for all k) as desired (since the scalar inequality (29) applies in particular to all (s, a) of the form $(s, \hat{\pi}_p^*(a))$).

Identically to the proof of Lemma 20, we can bound the number of values of k to be checked as $\leq 4 \frac{1}{1-\gamma}$. Fixing $s \in \mathcal{S}, a \in \mathcal{A}, u \in U$, and a value of k , by Bernstein's inequality (e.g. [Maurer and Pontil, 2009, Theorem 3]), since $(S_{sa}^i)_{i=1}^n$ (which determine $\hat{P}_{sa}$) are independent of $(\bar{V}_{s,a,u})^{\circ 2^k}$, we have that with

probability at least $1 - 2\delta'$,

$$\begin{aligned}
\left| \left(\hat{P}_{sa} - P_{sa} \right) \left(\bar{V}_{s,a,u} \right)^{\circ 2^k} \right| &\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{sa}} \left[\left(\bar{V}_{s,a,u} \right)^{\circ 2^k} \right]}{n}} + \frac{\log \left(\frac{1}{\delta'} \right)}{3n} \left\| \left(\bar{V}_{s,a,u} \right)^{\circ 2^k} \right\|_{\infty} \\
&= \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{sa}} \left[\left(\bar{V}_{s,a,u} \right)^{\circ 2^k} \right]}{n}} + \frac{\log \left(\frac{1}{\delta'} \right)}{3n} \left\| \bar{V}_{s,a,u} \right\|_{\infty}^{2^k} \\
&\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{sa}} \left[\left(\bar{V}_{s,a,u} \right)^{\circ 2^k} \right]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right)}{n} \left\| \bar{V}_{s,a,u} \right\|_{\infty}^{2^k} \\
&\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{sa}} \left[\left(\bar{V}_{s,a,u} \right)^{\circ 2^k} \right]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right) \cdot 2^k}{n} \left(\left\| \bar{V}_{s,a,u} \right\|_{\infty} + 1 \right)^{2^k}.
\end{aligned}$$

Taking a union bound over all $\leq S \cdot A \cdot \frac{48SA^2}{(1-\gamma)^3\xi\delta} \cdot \frac{4}{1-\gamma} = \frac{192S^2A^3}{(1-\gamma)^4\xi\delta}$ combinations of s, a, u, k and choosing $\delta' = \frac{\delta}{4} \frac{(1-\gamma)^4\xi\delta}{192S^2A^3}$, we thus obtain (28) with $\alpha = 2 \log \left(\frac{768S^2A^3}{(1-\gamma)^4\xi\delta^2} \right)$ as desired. $\square$

B.3 Completing the proof of Theorem 9

Proof of Theorem 9. Let $\alpha = 16 \log \left(\frac{12SA n}{(1-\gamma)^2\delta} \right)$, which is larger than $2 \log \left(\frac{6S \log_2 \log_2 (\|V^*\|_{\text{span}} + 4)}{\delta} \right)$ since $\log_2 \log_2 (\|V^*\|_{\text{span}} + 4) \leq 2 \frac{1}{1-\gamma}$ by the same arguments as those within the proof of Lemma 20. Combining Lemma 17 with Lemma 16, we obtain that with probability at least $1 - \delta$,

$$\begin{aligned}
\left\| \hat{V}^{\pi_\gamma^*} - V^{\pi_\gamma^*} \right\|_{\infty} &\leq \frac{4 \cdot 3 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right) \alpha}{(1-\gamma)n} \left(\|V^{\pi_\gamma^*}\|_{\text{span}} + 1 \right) \\
&\quad + \frac{2 \cdot 3 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{2\alpha \left(\|V^{\pi_\gamma^*}\|_{\text{span}} + 1 \right)}{n}} \\
&\leq \frac{12 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right) \alpha}{(1-\gamma)n} \left(\|V^{\pi_\gamma^*}\|_{\text{span}} + 1 \right) \\
&\quad + \frac{12 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha \left(\|V^{\pi_\gamma^*}\|_{\text{span}} + 1 \right)}{n}} \tag{30}
\end{aligned}$$

$$\leq \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha \left(\|V^{\pi_\gamma^*}\|_{\text{span}} + 1 \right)}{n}} \tag{31}$$

where we used that $1 + \lceil \log_2 \log_2 (\|V^*\|_{\text{span}} + 4) \rceil \leq 2 + \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right) \leq 3 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)$. We also use that the second term on the RHS of (30) is always the larger of the two when $\alpha \frac{\|V^{\pi_\gamma^*}\|_{\text{span}} + 1}{n} \leq 1$, and if this quantity is not ≤ 1 then by the trivial bound that $\left\| \hat{V}^{\pi_\gamma^*} - V^{\pi_\gamma^*} \right\|_{\infty} \leq \frac{1}{1-\gamma}$, inequality (31) still holds. Similarly, combining Lemma 20 with Lemma 16 and assuming $n \geq 4$ (so that the bound in Lemma 20 satisfies the assumptions of Lemma 16, noting also that the assumptions are satisfied since we assume $\hat{\pi}$ satisfies $\hat{V}^{\hat{\pi}} \geq \hat{V}^* - \frac{1}{n} \mathbf{1}$, which implies $\left\| \hat{V}^{\hat{\pi}} - \hat{V}^* \right\|_{\infty} \leq \frac{1}{n}$ since $\hat{V}^* \geq \hat{V}^{\hat{\pi}}$), we can perform analogous calculations to obtain that with probability at least $1 - \delta$,

$$\left\| \hat{V}^{\hat{\pi}} - V^{\hat{\pi}} \right\|_{\infty} \leq \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha \left(\left\| \hat{V}^{\hat{\pi}} \right\|_{\text{span}} + 1 \right)}{n}}. \tag{32}$$

Since $\widehat{V}^{\widehat{\pi}} \geq \widehat{V}^{\star} - \frac{1}{n}\mathbf{1}$, we have that elementwise

$$\begin{aligned}
V^{\star} &\geq V^{\widehat{\pi}} \geq \widehat{V}^{\widehat{\pi}} - \left\| \widehat{V}^{\widehat{\pi}} - V^{\widehat{\pi}} \right\|_{\infty} \mathbf{1} \\
&\geq \widehat{V}^{\star} - \frac{1}{n}\mathbf{1} - \left\| \widehat{V}^{\widehat{\pi}} - V^{\widehat{\pi}} \right\|_{\infty} \mathbf{1} \\
&\geq \widehat{V}^{\pi_{\gamma}^{\star}} - \frac{1}{n}\mathbf{1} - \left\| \widehat{V}^{\widehat{\pi}} - V^{\widehat{\pi}} \right\|_{\infty} \mathbf{1} \\
&\geq V^{\star} - \left\| \widehat{V}^{\pi_{\gamma}^{\star}} - V^{\pi_{\gamma}^{\star}} \right\|_{\infty} \mathbf{1} - \frac{1}{n}\mathbf{1} - \left\| \widehat{V}^{\widehat{\pi}} - V^{\widehat{\pi}} \right\|_{\infty} \mathbf{1}
\end{aligned} \tag{33}$$

(since $\widehat{V}^{\star} \geq \widehat{V}^{\pi_{\gamma}^{\star}}$ and $V^{\star} = V^{\pi_{\gamma}^{\star}}$). Combining this inequality with the bounds (31) and (32) (and taking the union bound, giving an overall failure probability bounded by 2δ) yields

$$\begin{aligned}
\left\| V^{\widehat{\pi}} - V^{\star} \right\|_{\infty} &\leq \frac{1}{n} + \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \left(\sqrt{\frac{\alpha \left(\left\| \widehat{V}^{\widehat{\pi}} \right\|_{\text{span}} + 1 \right)}{n}} + \sqrt{\frac{\alpha \left(\left\| V^{\pi_{\gamma}^{\star}} \right\|_{\text{span}} + 1 \right)}{n}} \right) \\
&\leq \frac{1}{n} + \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{2\alpha \left(\left\| \widehat{V}^{\widehat{\pi}} \right\|_{\text{span}} + \left\| V^{\pi_{\gamma}^{\star}} \right\|_{\text{span}} + 2 \right)}{n}} \\
&\leq \frac{25 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{2\alpha \left(\left\| \widehat{V}^{\widehat{\pi}} \right\|_{\text{span}} + \left\| V^{\pi_{\gamma}^{\star}} \right\|_{\text{span}} + 2 \right)}{n}} \\
&\leq \frac{25 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{4\alpha \left(\left\| \widehat{V}^{\widehat{\pi}} \right\|_{\text{span}} + \left\| V^{\pi_{\gamma}^{\star}} \right\|_{\text{span}} + 1 \right)}{n}} \\
&= \frac{1}{1-\gamma} \sqrt{4 \left(25 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right) \right)^2 16 \log \left(\frac{12SA n}{(1-\gamma)^2 \delta} \right) \frac{\left(\left\| \widehat{V}^{\widehat{\pi}} \right\|_{\text{span}} + \left\| V^{\pi_{\gamma}^{\star}} \right\|_{\text{span}} + 1 \right)}{n}} \\
&\leq \frac{1}{1-\gamma} \sqrt{4 \left(25 \log \left(\frac{4}{1-\gamma} \right) \right)^2 16 \log \left(\frac{12SA n}{(1-\gamma)^2 \delta} \right) \frac{\left(\left\| \widehat{V}^{\widehat{\pi}} \right\|_{\text{span}} + \left\| V^{\pi_{\gamma}^{\star}} \right\|_{\text{span}} + 1 \right)}{n}} \\
&\leq \frac{1}{1-\gamma} \sqrt{C_1 \log^3 \left(\frac{SA n}{(1-\gamma)\delta} \right) \frac{\left(\left\| \widehat{V}^{\widehat{\pi}} \right\|_{\text{span}} + \left\| V^{\pi_{\gamma}^{\star}} \right\|_{\text{span}} + 1 \right)}{n}}
\end{aligned}$$

where we used the fact that $\sqrt{a} + \sqrt{b} \leq \sqrt{2(a+b)}$, the definition of α , that $\log_2 \log_2(x+4) \leq \log 4x$ for $x \geq 1$, and finally chose C_1 sufficiently large (in particular large enough to ensure the above bound is vacuous when the $n \geq 4$ assumption is not satisfied; $C_1 \geq 4$ suffices). $\square$

B.4 Completing the proof of Theorem 10

First we briefly outline the proof. The proof of Theorem 9 already bounds the quantity $\left\| \widehat{V}^{\pi_{\gamma}^{\star}} - V^{\pi_{\gamma}^{\star}} \right\|_{\infty}$, which can be reused here. The more difficult step is bounding the “evaluation error” of the empirical optimal policy of the perturbed MDP, $\widehat{\pi}_p^{\star}$. (While in the statement of Theorem 10 we referred to this policy as $\widehat{\pi}$, here we rename it to $\widehat{\pi}_p^{\star}$ to emphasize that it is optimal for the perturbed MDP.) There are several possible choices for which pair of value functions to bound (for example, $\left\| \widehat{V}^{\widehat{\pi}_p^{\star}} - V^{\widehat{\pi}_p^{\star}} \right\|_{\infty}$ and $\left\| {}_p\widehat{V}^{\widehat{\pi}_p^{\star}} - V^{\widehat{\pi}_p^{\star}} \right\|_{\infty}$ are two reasonable choices). Since Lemma 15 assumes that the reward is bounded by 1 while we only have $\widetilde{r} \leq 1 + \xi$, it is most convenient to bound the term $\left\| \widehat{V}^{\widehat{\pi}_p^{\star}} - V^{\widehat{\pi}_p^{\star}} \right\|_{\infty}$ since it does not involve the perturbed reward function $\widetilde{r}$. (Lemma 15 could be trivially modified to handle differently scaled rewards, but by bounding

$\left\| \widehat{V}^{\widehat{\pi}_p^*} - V^{\widehat{\pi}_p^*} \right\|_\infty$ (we avoid this extra bookkeeping.) This consideration motivated our choice of the particular Bernstein-like-condition to check within Lemma 22. Thus, combining Lemma 22 with Lemma 15, we obtain a bound on $\left\| \widehat{V}^{\widehat{\pi}_p^*} - V^{\widehat{\pi}_p^*} \right\|_\infty$. We can use the fact that the perturbation is small to bound $\left\| {}_p\widehat{V}^{\widehat{\pi}_p^*} - \widehat{V}^{\widehat{\pi}_p^*} \right\|_\infty$ and $\left\| {}_p\widehat{V}^{\pi_\gamma^*} - \widehat{V}^{\pi_\gamma^*} \right\|_\infty$, which by triangle inequality gives us bounds on the quantities $\left\| {}_p\widehat{V}^{\widehat{\pi}_p^*} - V^{\widehat{\pi}_p^*} \right\|_\infty$ and $\left\| {}_p\widehat{V}^{\pi_\gamma^*} - V^{\pi_\gamma^*} \right\|_\infty$. We conclude by using the fact that ${}_p\widehat{V}^{\widehat{\pi}_p^*} \geq {}_p\widehat{V}^{\pi_\gamma^*}$, since $\widehat{\pi}_p^*$ is optimal for the perturbed empirical MDP.

Proof of Theorem 10. We follow the above sketch. First we bound $\left\| \widehat{V}^{\pi_\gamma^*} - V^{\pi_\gamma^*} \right\|_\infty$. Combining Lemma 17 with Lemma 16, similarly to the proof of Theorem 9 we obtain that with probability at least $1 - \delta$, with $\alpha_1 = 2 \log \left(\frac{6S \log_2 \log_2 (\|V^{\pi_\gamma^*}\|_{\text{span}} + 4)}{\delta} \right)$,

$$\begin{aligned} \left\| \widehat{V}^{\pi_\gamma^*} - V^{\pi_\gamma^*} \right\|_\infty &\leq \frac{4 \cdot 3 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right) \alpha_1}{(1-\gamma)n} \left(\|V^{\pi_\gamma^*}\|_{\text{span}} + 1 \right) \\ &\quad + \frac{2 \cdot 3 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{2\alpha_1 (\|V^{\pi_\gamma^*}\|_{\text{span}} + 1)}{n}} \\ &\leq \frac{12 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right) \alpha_1}{(1-\gamma)n} \left(\|V^{\pi_\gamma^*}\|_{\text{span}} + 1 \right) \\ &\quad + \frac{12 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_1 (\|V^{\pi_\gamma^*}\|_{\text{span}} + 1)}{n}} \\ &\leq \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_1 (\|V^{\pi_\gamma^*}\|_{\text{span}} + 1)}{n}} \end{aligned} \quad (34)$$

where again we used that $1 + \lceil \log_2 \log_2 (\|V^*\|_{\text{span}} + 4) \rceil \leq 2 + \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right) \leq 3 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)$ and the fact that the second term is always larger whenever the bound on $\left\| \widehat{V}^{\pi_\gamma^*} - V^{\pi_\gamma^*} \right\|_\infty$ is non-vacuous. To bound $\left\| \widehat{V}^{\widehat{\pi}_p^*} - V^{\widehat{\pi}_p^*} \right\|_\infty$, we can combine Lemma 22 with Lemma 16 and perform analogous calculations to obtain that with additional failure probability at most δ ,

$$\left\| \widehat{V}^{\widehat{\pi}_p^*} - V^{\widehat{\pi}_p^*} \right\|_\infty \leq \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \quad (35)$$

where $\alpha_2 = 2 \log \left(\frac{768S^2A^3}{(1-\gamma)^4\xi\delta^2} \right)$.

Next we bound the terms $\left\| {}_p\widehat{V}^{\widehat{\pi}_p^*} - \widehat{V}^{\widehat{\pi}_p^*} \right\|_\infty$ and $\left\| {}_p\widehat{V}^{\pi_\gamma^*} - \widehat{V}^{\pi_\gamma^*} \right\|_\infty$. We have

$$\begin{aligned} \left\| {}_p\widehat{V}^{\widehat{\pi}_p^*} - \widehat{V}^{\widehat{\pi}_p^*} \right\|_\infty &= \left\| (I - \gamma \widehat{P}_{\widehat{\pi}_p^*})^{-1} \widetilde{r}_{\widehat{\pi}_p^*} - (I - \gamma \widehat{P}_{\widehat{\pi}_p^*})^{-1} r_{\widehat{\pi}_p^*} \right\|_\infty \\ &\leq \left\| (I - \gamma \widehat{P}_{\widehat{\pi}_p^*})^{-1} \right\|_{\infty \rightarrow \infty} \left\| \widetilde{r}_{\widehat{\pi}_p^*} - r_{\widehat{\pi}_p^*} \right\|_\infty \\ &\leq \frac{1}{1-\gamma} \|\widetilde{r} - r\|_\infty \\ &\leq \frac{\xi}{1-\gamma} \end{aligned} \quad (36)$$

and likewise

$$\left\| {}_p\widehat{V}^{\pi_\gamma^*} - \widehat{V}^{\pi_\gamma^*} \right\|_\infty \leq \frac{\xi}{1-\gamma}. \quad (37)$$

Finally, we can combine all of these bounds with the fact that ${}_p\widehat{V}^{\widehat{\pi}_p^*} \geq {}_p\widehat{V}^{\pi_\gamma^*}$ (since $\widehat{\pi}_p^*$ is optimal for the perturbed empirical MDP) to obtain that with probability at least $1 - 2\delta$,

$$\begin{aligned} V^{\widehat{\pi}_p^*} &\geq \widehat{V}^{\widehat{\pi}_p^*} - \left\| \widehat{V}^{\widehat{\pi}_p^*} - V^{\widehat{\pi}_p^*} \right\|_\infty \mathbf{1} \\ &\geq \widehat{V}^{\widehat{\pi}_p^*} - \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \mathbf{1} \end{aligned} \quad (35)$$

$$\begin{aligned} &\geq {}_p\widehat{V}^{\widehat{\pi}_p^*} - \left\| {}_p\widehat{V}^{\widehat{\pi}_p^*} - \widehat{V}^{\widehat{\pi}_p^*} \right\|_\infty \mathbf{1} - \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \mathbf{1} \\ &\geq {}_p\widehat{V}^{\widehat{\pi}_p^*} - \frac{\xi}{1-\gamma} \mathbf{1} - \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \mathbf{1} \end{aligned} \quad (36)$$

$$\begin{aligned} &\geq {}_p\widehat{V}^{\pi_\gamma^*} - \frac{\xi}{1-\gamma} \mathbf{1} - \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \mathbf{1} \\ &\geq {}_p\widehat{V}^{\pi_\gamma^*} - \left\| {}_p\widehat{V}^{\pi_\gamma^*} - \widehat{V}^{\pi_\gamma^*} \right\|_\infty \mathbf{1} - \frac{\xi}{1-\gamma} \mathbf{1} - \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \mathbf{1} \\ &\geq {}_p\widehat{V}^{\pi_\gamma^*} - \frac{2\xi}{1-\gamma} \mathbf{1} - \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \mathbf{1} \end{aligned} \quad (37)$$

$$\begin{aligned} &\geq V^{\pi_\gamma^*} - \left\| \widehat{V}^{\pi_\gamma^*} - V^{\pi_\gamma^*} \right\|_\infty \mathbf{1} - \frac{2\xi}{1-\gamma} \mathbf{1} - \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \mathbf{1} \\ &\geq V^{\pi_\gamma^*} - \frac{2\xi}{1-\gamma} \mathbf{1} - \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \left(\sqrt{\frac{\alpha_2 (\|V^{\widehat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \right. \\ &\quad \left. + \sqrt{\frac{\alpha_1 (\|V^{\pi_\gamma^*}\|_{\text{span}} + 1)}{n}} \right) \mathbf{1} \end{aligned} \quad (34).$$

Now using that $\xi \leq \frac{1}{n}$, the fact that $\sqrt{a} + \sqrt{b} \leq 2\sqrt{a+b}$, and that

$$\alpha_1 = 2 \log \left(\frac{6S \log_2 \log_2 \left(\|V^{\pi_\gamma^*}\|_{\text{span}} + 4 \right)}{\delta} \right) \leq 2 \log \left(\frac{12S}{(1-\gamma)\delta} \right) \leq \alpha_2$$

(since $\log_2 \log_2 (\|V^{\pi_\gamma^*}\|_{\text{span}} + 4) \leq 2\frac{1}{1-\gamma}$ as shown within the proof of Lemma 20), we can simplify

$$\begin{aligned}
& \frac{2\xi}{1-\gamma} + \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4\right)}{1-\gamma} \left(\sqrt{\frac{\alpha_2 (\|V^{\hat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} + \sqrt{\frac{\alpha_1 (\|V^{\pi_\gamma^*}\|_{\text{span}} + 1)}{n}} \right) \\
& \leq \frac{2}{n(1-\gamma)} + \frac{48 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4\right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\pi_\gamma^*}\|_{\text{span}} + \|V^{\hat{\pi}_p^*}\|_{\text{span}} + 2)}{n}} \\
& \leq \frac{96 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4\right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\pi_\gamma^*}\|_{\text{span}} + \|V^{\hat{\pi}_p^*}\|_{\text{span}} + 2)}{n}} \\
& \leq \frac{96\sqrt{2} \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4\right)}{1-\gamma} \sqrt{\frac{\alpha_2 (\|V^{\pi_\gamma^*}\|_{\text{span}} + \|V^{\hat{\pi}_p^*}\|_{\text{span}} + 1)}{n}} \\
& \leq \frac{C_2}{1-\gamma} \sqrt{\frac{\log^3 \left(\frac{SAn}{(1-\gamma)\delta\xi}\right) (\|V^{\pi_\gamma^*}\|_{\text{span}} + \|V^{\hat{\pi}_p^*}\|_{\text{span}} + 1)}{n}}
\end{aligned}$$

where in the final inequality we use that $\log_2 \log_2 \left(\frac{1}{1-\gamma} + 4\right) \leq \alpha_2$ and chosen sufficiently large constant C_1 (including absorbing an additional constant due to adjusting the failure probability to be at most δ rather than 2δ). $\square$

C Proofs of AMDP Theorems

C.1 Useful Lemmas

The following is an average-reward version of the simulation lemma. Such techniques are well-known [Cao, 1999, Meyer, 1980].

Lemma 23. *Fix a policy π , and let $P, \hat{P}$ be any two MDP transition matrices. Let $\Delta = \hat{P}_\pi - P_\pi$. Then*

$$\hat{P}_\pi^\infty - P_\pi^\infty = \hat{P}_\pi^\infty \Delta H_{P_\pi} - (\hat{P}_\pi^\infty - P_\pi^\infty) P_\pi^\infty.$$

Consequently,

1. If the Markov chain P_π satisfies $P_\pi^\infty = \mathbf{1}\mu^\top$ for some probability distribution $\mu^\top$, then

$$\hat{P}_\pi^\infty - P_\pi^\infty = \hat{P}_\pi^\infty \Delta H_{P_\pi}.$$

2. If the quantity $\rho^\pi = P_\pi^\infty r_\pi$ is constant (has the form $\alpha \mathbf{1}$ for some $\alpha \in \mathbb{R}$), then

$$\hat{\rho}^\pi - \rho^\pi = \hat{P}_\pi^\infty r_\pi - P_\pi^\infty r_\pi = \hat{P}_\pi^\infty \Delta H_{P_\pi} r_\pi = \hat{P}_\pi^\infty \Delta h^\pi.$$

Proof. By the properties of limiting matrices,

$$\begin{aligned}
\hat{P}_\pi^\infty - P_\pi^\infty &= \hat{P}_\pi^\infty \hat{P}_\pi - P_\pi^\infty P_\pi \\
&= \hat{P}_\pi^\infty (P_\pi + \Delta) - P_\pi^\infty P_\pi \\
&= (\hat{P}_\pi^\infty - P_\pi^\infty) P_\pi + \hat{P}_\pi^\infty \Delta.
\end{aligned}$$

Therefore $(\hat{P}_\pi^\infty - P_\pi^\infty)(I - P_\pi) = \hat{P}_\pi^\infty \Delta$. Now post-multiplying both sides by the deviation matrix H_{P_π} , which satisfies $(I - P_\pi)H_{P_\pi} = I - P_\pi^\infty$, we obtain

$$\hat{P}_\pi^\infty \Delta H_{P_\pi} = (\hat{P}_\pi^\infty - P_\pi^\infty)(I - P_\pi)H_{P_\pi} = (\hat{P}_\pi^\infty - P_\pi^\infty)(I - P_\pi^\infty) = \hat{P}_\pi^\infty - P_\pi^\infty + (\hat{P}_\pi^\infty - P_\pi^\infty)P_\pi^\infty.$$

For the first consequence, note that if $P_\pi^\infty = \mathbf{1}\mu^\top$, then since $\hat{P}_\pi^\infty$ and P_π^∞ are stochastic matrices and have $\mathbf{1}$ as a right eigenvector with eigenvalue 1, we have

$$(\hat{P}_\pi^\infty - P_\pi^\infty)P_\pi^\infty = (\hat{P}_\pi^\infty - P_\pi^\infty)\mathbf{1}\mu^\top = \mathbf{1}\mu^\top - \mathbf{1}\mu^\top = 0.$$

Similarly, for the second consequence, we have

$$(\hat{P}_\pi^\infty - P_\pi^\infty)P_\pi^\infty r_\pi = (\hat{P}_\pi^\infty - P_\pi^\infty)\alpha\mathbf{1} = \alpha\mathbf{1} - \alpha\mathbf{1} = 0.$$

□

Lemma 24. *Let P be any transition matrix, let $\eta \in (0, 1)$, and let $s_0 \in \mathcal{S}$ be an arbitrary state. Form the anchored transition matrix $\tilde{P} = (1 - \eta)P + \eta\mathbf{1}e_{s_0}^\top \in \mathbb{R}^{\mathcal{S} \times \mathcal{S}}$. Also we use $\tilde{h}^\pi$, $\tilde{P}_\pi^\infty$ to denote the bias function of policy π in $\tilde{P}$ and the limiting matrix of $\tilde{P}_\pi$. Then*

1. *For all policies π , the state s_0 is recurrent in the Markov chain $\tilde{P}_\pi$. Consequently $\tilde{P}$ is unichain.*
2. *Fix a policy π . Then*
 - (a) $\tilde{P}_\pi^\infty = \eta\mathbf{1}e_{s_0}^\top (I - (1 - \eta)P_\pi)^{-1}$.
 - (b) $\tilde{\rho}^\pi = \eta\mathbf{1}V_{1-\eta}^\pi(s_0)$, where $V_{1-\eta}^\pi$ is the discounted value function for policy π with discount factor $1 - \eta$ (or equivalently effective horizon $\frac{1}{\eta}$).
 - (c) $\tilde{h}^\pi = V_{1-\eta}^\pi + c\mathbf{1}$ for some scalar c .
 - (d) $\|\tilde{\rho}^\pi - \rho^\pi\|_\infty \leq 2\eta\|\tilde{h}^\pi\|_\infty$.
 - (e) If ρ^π is a state-independent constant, then $\|\tilde{h}^\pi\|_{\text{span}} \leq 2\|h^\pi\|_{\text{span}}$.
3. *Letting $\tilde{\rho}^*$ and $\tilde{h}^*$ be the optimal gain and optimal bias of $\tilde{P}$, respectively, we have*
 - (a) $\tilde{\rho}^* = \eta\mathbf{1}V_{1-\eta}^*(s_0)$, where $V_{1-\eta}^*$ is the optimal discounted value function with discount factor $1 - \eta$.
 - (b) $\tilde{h}^* = V_{1-\eta}^* + c\mathbf{1}$ for some scalar c .
 - (c) If ρ^* is a state-independent constant, then $\|\tilde{h}^*\|_{\text{span}} \leq 2\|h^*\|_{\text{span}}$.
 - (d) The average-reward Bellman optimality operator for $\tilde{P}$, $\tilde{\mathcal{T}}(h) := M(r + \tilde{P}h)$, is a $(1 - \eta)$ -span contraction: $\|\tilde{\mathcal{T}}(h) - \tilde{\mathcal{T}}(h')\|_{\text{span}} \leq (1 - \eta)\|h - h'\|_{\text{span}}$.

Proof of Lemma 24. We start with 1. Fix a policy π and consider the Markov chain $\tilde{P}_\pi$. Since this is a finite Markov chain, there must exist some recurrent state s , and since in $\tilde{P}_\pi$ there is probability at least $\eta > 0$ of transitioning to s_0 from s , the state s_0 is also recurrent [Durrett, 2019, Chapter 5.3]. Furthermore this shows that for any state s which is recurrent, it is in the same recurrent class as s_0 , and therefore there is only one recurrent class in $\tilde{P}_\pi$. Since this holds for arbitrary π (and in particular for all deterministic π), the MDP $\tilde{P}$ is unichain.

Now we show all the properties in statement 2. Again fix a policy π . Since $\tilde{P}_\pi^\infty$ must satisfy $\tilde{P}_\pi^\infty \tilde{P}_\pi = \tilde{P}_\pi^\infty$, expanding the definition of $\tilde{P}_\pi^\infty$ we have

$$\tilde{P}_\pi^\infty = \tilde{P}_\pi \tilde{P}_\pi^\infty = (1 - \eta)\tilde{P}_\pi^\infty P_\pi + \eta\tilde{P}_\pi^\infty \mathbf{1}e_{s_0}^\top = (1 - \eta)\tilde{P}_\pi^\infty P_\pi + \eta\mathbf{1}e_{s_0}^\top$$

using that $\tilde{P}_\pi^\infty \mathbf{1} = \mathbf{1}$ in the last equality. By rearranging we have

$$\tilde{P}_\pi^\infty (I - (1 - \eta)P_\pi) = \eta\mathbf{1}e_{s_0}^\top$$

and since $\|(1 - \eta)P_\pi\|_{\infty \rightarrow \infty} = (1 - \eta) < 1$, it is a standard fact that the matrix $(I - (1 - \eta)P_\pi)$ is invertible, and so

$$\tilde{P}_\pi^\infty = \eta\mathbf{1}e_{s_0}^\top (I - (1 - \eta)P_\pi)^{-1}$$

as desired. We can then calculate that

$$\tilde{\rho}^\pi = \tilde{P}_\pi^\infty r_\pi = \eta \mathbf{1} e_{s_0}^\top (I - (1 - \eta)P_\pi)^{-1} r_\pi = \eta \mathbf{1} e_{s_0}^\top V_{1-\eta}^\pi = \eta \mathbf{1} V_{1-\eta}^\pi(s_0).$$

Next, to compute $\tilde{h}^\pi$, we check that $\eta \mathbf{1} V_{1-\eta}^\pi(s_0)$ and $V_{1-\eta}^\pi$ satisfy the evaluation equations [Puterman, 1994, Section 8.2.3]. We have that

$$\begin{aligned} \eta \mathbf{1} V_{1-\eta}^\pi(s_0) + (I - \tilde{P}_\pi) V_{1-\eta}^\pi &= \eta \mathbf{1} V_{1-\eta}^\pi(s_0) + (I - (1 - \eta)P_\pi - \eta \mathbf{1} e_{s_0}^\top) V_{1-\eta}^\pi \\ &= (I - (1 - \eta)P_\pi) V_{1-\eta}^\pi \\ &= (I - (1 - \eta)P_\pi)(I - (1 - \eta)P_\pi)^{-1} r_\pi \\ &= r_\pi \end{aligned}$$

so by [Puterman, 1994, Corollary 8.2.7], since $\tilde{P}_\pi$ is unichain, we have that $\tilde{h}^\pi = V_{1-\eta}^\pi + c \mathbf{1}$ for some scalar c . Next, since we have already checked that $\tilde{\rho}^\pi$ is constant, we can apply Lemma 23 to obtain

$$\tilde{\rho}^\pi - \rho^\pi = P_\pi^\infty (\tilde{P}_\pi - P_\pi) \tilde{h}^\pi$$

and thus

$$\|\tilde{\rho}^\pi - \rho^\pi\|_\infty \leq \|P_\pi^\infty\|_{\infty \rightarrow \infty} \|\tilde{P}_\pi - P_\pi\|_{\infty \rightarrow \infty} \|\tilde{h}^\pi\|_\infty \leq 1 \cdot \|\eta \mathbf{1} e_{s_0}^\top - \eta P_\pi\|_{\infty \rightarrow \infty} \|\tilde{h}^\pi\|_\infty \leq 2\eta \|\tilde{h}^\pi\|_\infty.$$

Finally, assuming that ρ^π is a constant vector, we want to show that $\|\tilde{h}^\pi\|_{\text{span}} \leq 2\|h^\pi\|_{\text{span}}$. Since we have shown $\|\tilde{h}^\pi\|_{\text{span}} = \|V_{1-\eta}^\pi\|_{\text{span}}$, it suffices to bound $\|V_{1-\eta}^\pi\|_{\text{span}}$. We calculate

$$\begin{aligned} \|\tilde{h}^\pi\|_{\text{span}} &= \|V_{1-\eta}^\pi\|_{\text{span}} \\ &= \|V_{1-\eta}^\pi - \frac{1}{\eta} \rho^\pi\|_{\text{span}} && \text{because } \rho^\pi \text{ is constant} \\ &= \|(I - (1 - \eta)P_\pi)^{-1} r_\pi - \frac{1}{\eta} \rho^\pi\|_{\text{span}} \\ &= \|(I - (1 - \eta)P_\pi)^{-1} (\rho^\pi + (I - P_\pi)h^\pi) - \frac{1}{\eta} \rho^\pi\|_{\text{span}} && \rho^\pi + h^\pi = r_\pi + P_\pi h^\pi \\ &= \|(I - (1 - \eta)P_\pi)^{-1} (I - P_\pi)h^\pi\|_{\text{span}} && (I - (1 - \eta)P_\pi)^{-1} \rho^\pi = \frac{1}{\eta} \rho^\pi. \end{aligned}$$

The fact that $(I - (1 - \eta)P_\pi)^{-1} \rho^\pi = \frac{1}{\eta} \rho^\pi$ for general policies π follows from the fact that $P_\pi \rho^\pi = P_\pi P_\pi^\infty r_\pi = P_\pi^\infty r_\pi = \rho^\pi$, which implies that $P_\pi^t \rho^\pi = \rho^\pi$ which we can then combine with the Neumann series to obtain that $(I - (1 - \eta)P_\pi)^{-1} \rho^\pi = \sum_{t=0}^{\infty} (1 - \eta)^t P_\pi^t \rho^\pi = \sum_{t=0}^{\infty} (1 - \eta)^t \rho^\pi = \frac{1}{\eta} \rho^\pi$.

For convenience writing $\gamma = 1 - \eta$, using the Neumann series formula we have

$$\begin{aligned} (I - \gamma P_\pi)^{-1} (I - P_\pi) &= \sum_{t=0}^{\infty} \gamma^t P_\pi^t - \sum_{t=0}^{\infty} \gamma^t P_\pi^{t+1} \\ &= I + \sum_{t=1}^{\infty} \gamma^t P_\pi^t - \sum_{t=0}^{\infty} \gamma^t P_\pi^{t+1} \\ &= I + \sum_{t=0}^{\infty} \gamma^{t+1} P_\pi^{t+1} - \sum_{t=0}^{\infty} \gamma^t P_\pi^{t+1} \\ &= I - (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P_\pi^{t+1} \end{aligned}$$

and we note that $(1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P_\pi^{t+1}$ is a stochastic matrix (since all terms are nonnegative, and, since each row of P_π^t sums to 1 for any t , the rows all sum to $(1 - \gamma) \sum_{t=0}^{\infty} \gamma^t 1 = 1$). Therefore continuing the previous

calculation,

$$\begin{aligned}
\|\tilde{h}^\pi\|_{\text{span}} &= \|(I - (1 - \eta)P_\pi)^{-1}(I - P_\pi)h^\pi\|_{\text{span}} \\
&= \|h^\pi - (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P_\pi^{t+1} h^\pi\|_{\text{span}} \\
&\leq \|h^\pi\|_{\text{span}} + \left\| \left((1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P_\pi^{t+1} \right) h^\pi \right\|_{\text{span}} \\
&\leq 2\|h^\pi\|_{\text{span}}
\end{aligned}$$

where the last inequality is because for any stochastic matrix P' , $\|P'h^\pi\|_{\text{span}} \leq \|h^\pi\|_{\text{span}}$.

Now we verify statement 3. First we show that $\tilde{\mathcal{T}}$ is a $(1 - \eta)$ -span contraction. This follows from existing results, since the fact that all states have probability $\geq \eta$ of transitioning to s_0 means we could apply [Puterman, 1994, Theorem 8.5.2]. However, we will provide a direct proof due to its simplicity. Letting $h, h' \in \mathbb{R}^S$ be arbitrary, we can calculate

$$\begin{aligned}
\|\tilde{\mathcal{T}}(h) - \tilde{\mathcal{T}}(h')\|_{\text{span}} &= \|M(r + \tilde{P}h) - M(r + \tilde{P}h')\|_{\text{span}} \\
&= \|M(r + (1 - \eta)Ph + \eta \mathbf{1}h(s_0)) - M(r + (1 - \eta)Ph' + \eta \mathbf{1}h'(s_0))\|_{\text{span}} \\
&= \|M(r + (1 - \eta)Ph) + \eta \mathbf{1}h(s_0) - M(r + (1 - \eta)Ph') - \eta \mathbf{1}h'(s_0)\|_{\text{span}} \\
&= \|M(r + (1 - \eta)Ph) - M(r + (1 - \eta)Ph')\|_{\text{span}} \\
&\leq \|r + (1 - \eta)Ph - r - (1 - \eta)Ph'\|_{\text{span}} \\
&= (1 - \eta)\|Ph - Ph'\|_{\text{span}} \\
&\leq (1 - \eta)\|h - h'\|_{\text{span}}
\end{aligned}$$

where we used the fact that M is $\|\cdot\|_{\text{span}}$ -nonexpansive, which we verify now. Letting $x, x' \in \mathbb{R}^{SA}$ be arbitrary and letting π and π' satisfy $M(x) = M^\pi x$ and $M(x') = M^{\pi'} x'$, we have

$$M(x) - M(x') = M^\pi x - M^{\pi'} x' \leq M^\pi x - M^\pi x' = M^\pi(x - x') \leq \left(\max_{s \in S} x(s) - x'(s) \right) \mathbf{1}$$

and analogously

$$M(x) - M(x') = M^\pi x - M^{\pi'} x' \geq M^{\pi'} x - M^{\pi'} x' = M^{\pi'}(x - x') \geq \left(\min_{s \in S} x(s) - x'(s) \right) \mathbf{1}$$

so $\|M(x) - M(x')\|_{\text{span}} \leq \|x - x'\|_{\text{span}}$ as desired. Now we check that the claimed forms of $\tilde{\rho}^*$ and $\tilde{h}^*$ satisfy the (unichain) average optimality equation. We have

$$\begin{aligned}
M(r + \tilde{P}V_{1-\eta}^*) &= M(r + (1 - \eta)PV_{1-\eta}^* + \eta \mathbf{1}V_{1-\eta}^*(s_0)) \\
&= M(r + (1 - \eta)PV_{1-\eta}^*) + \eta \mathbf{1}V_{1-\eta}^*(s_0) \\
&= V_{1-\eta}^* + \eta \mathbf{1}V_{1-\eta}^*(s_0)
\end{aligned}$$

(using the discounted Bellman equation in the final equality), so indeed $\tilde{\rho}^* = \eta \mathbf{1}V_{1-\eta}^*(s_0)$ and $\tilde{h}^* = V_{1-\eta}^* + c\mathbf{1}$ for some scalar c [Puterman, 1994, Theorem 8.4.3]. (In general satisfying the average optimality equation only determines the optimal gain, but in unichain models the optimal bias is also determined up to a constant by the optimality equation [Puterman, 1994, Section 8.4.2] Schweitzer and Federgruen [1978]. In our setting it is also possible to show this directly as a consequence of the span-non-expansiveness of $\tilde{\mathcal{T}}$.)

Finally we check that $\|\tilde{h}^*\|_{\text{span}} \leq 2\|h^*\|_{\text{span}}$ in the case that ρ^* is constant. Note that $\|\tilde{h}^*\|_{\text{span}} = \|V_{1-\eta}^*\|_{\text{span}} = \|V_{1-\eta}^{\pi_{1-\eta}^*}\|_{\text{span}}$, whereas $\|h^*\|_{\text{span}} = \|h^{\pi^*}\|_{\text{span}}$, so we are comparing two different policies. Essentially the same bound has appeared in prior work, for instance [Wei et al., 2020, Lemma 2], but for

completeness we reprove it with a manner of calculation very similar to the previous case concerning the a fixed policy. First, notice that (letting $\gamma = 1 - \eta$ for notational convenience)

$$\begin{aligned}
V_{1-\eta}^* &\geq V_{1-\eta}^{\pi^*} \\
&= (I - \gamma P_{\pi^*})^{-1} r_{\pi^*} \\
&= (I - \gamma P_{\pi^*})^{-1} (\rho^* + (I - P_{\pi^*})h^*) && \rho^* + h^* = r_{\pi^*} + P_{\pi^*}h^* \\
&= (I - \gamma P_{\pi^*})^{-1} \rho^* + (I - \gamma P_{\pi^*})^{-1} (I - P_{\pi^*})h^* \\
&= \frac{1}{1-\gamma} \rho^* + (I - \gamma P_{\pi^*})^{-1} (I - P_{\pi^*})h^*
\end{aligned}$$

where the last equality holds because ρ^* is a state-independent constant. Next, using the fact that

$$\rho^* + h^* = M(r + Ph^*) \geq M^{\pi_\gamma^*}(r + Ph^*) = r_{\pi_\gamma^*} + P_{\pi_\gamma^*}h^*,$$

we have

$$\begin{aligned}
V_{1-\eta}^* &= (I - \gamma P_{\pi_\gamma^*})^{-1} r_{\pi_\gamma^*} \\
&\leq (I - \gamma P_{\pi_\gamma^*})^{-1} (\rho^* + (I - P_{\pi_\gamma^*})h^*) && \text{above inequality, monotonicity of } (I - \gamma P_{\pi_\gamma^*})^{-1} \\
&= (I - \gamma P_{\pi_\gamma^*})^{-1} \rho^* + (I - \gamma P_{\pi_\gamma^*})^{-1} (I - P_{\pi_\gamma^*})h^* \\
&= \frac{1}{1-\gamma} \rho^* + (I - \gamma P_{\pi_\gamma^*})^{-1} (I - P_{\pi_\gamma^*})h^* && \rho^* \text{ is a constant vector.}
\end{aligned}$$

Combining these two calculations we have

$$\frac{1}{1-\gamma} \rho^* + (I - \gamma P_{\pi^*})^{-1} (I - P_{\pi^*})h^* \leq V_{1-\eta}^* \leq \frac{1}{1-\gamma} \rho^* + (I - \gamma P_{\pi_\gamma^*})^{-1} (I - P_{\pi_\gamma^*})h^*.$$

We can also reuse our previous calculation that for any π , $(I - \gamma P_\pi)^{-1} (I - P_\pi) = I - Q$ for some stochastic matrix Q , to obtain

$$\frac{1}{1-\gamma} \rho^* + h^* - Q_1 h^* \leq V_{1-\eta}^* \leq \frac{1}{1-\gamma} \rho^* + h^* - Q_2 h^*$$

(for stochastic matrices Q_1, Q_2). We have the elementwise bounds $h^* - Q_2 h^* \leq h^* - (\min_s h^*(s)) \mathbf{1} \leq \|h^*\|_{\text{span}} \mathbf{1}$ and likewise $h^* - Q_1 h^* \geq -\|h^*\|_{\text{span}} \mathbf{1}$, which combined with the above display inequalities imply that $\|\hat{h}^*\|_{\text{span}} = \|V_{1-\eta}^*\|_{\text{span}} \leq 2\|h^*\|_{\text{span}}$. $\square$

Lemma 25. *Let $\hat{\gamma}^* = \inf \left\{ \gamma \in [0, 1) : \exists c \in \mathbb{R} \text{ such that } \left\| \hat{V}_\gamma^* - \hat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n} \right\}$. Then if $\hat{P}$ is weakly communicating, then the above set is nonempty and the above infimum is attained, that is there exists $c \in \mathbb{R}$ such that $\left\| \hat{V}_{\hat{\gamma}^*}^* - \hat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n}$ (and so $\hat{\gamma}^*$ may be defined as the smallest discount factor satisfying this property).*

Proof. If $\hat{P}$ is weakly communicating, then we have that $\hat{\rho}^*$ is a constant vector. Letting $\hat{\pi}^*$ be a Blackwell-optimal policy for $(\hat{P}, r)$ and letting $\hat{\gamma}_{\text{BW}} < 1$ be the Blackwell discount factor, we have that $\hat{\rho}^* = \hat{\rho}^{\hat{\pi}^*}$, that $\hat{h}^* = \hat{h}^{\hat{\pi}^*}$, and that $\hat{V}_\gamma^* = \hat{V}_\gamma^{\hat{\pi}^*}$ for all $\gamma \in [\hat{\gamma}_{\text{BW}}, 1)$ [Puterman \[1994\]](#). By the well-known Laurent series expansion (e.g. [\[Puterman, 1994, Corollary 8.2.4\]](#)), we also have $\hat{V}_\gamma^{\hat{\pi}^*} = \frac{1}{1-\gamma} \hat{\rho}^{\hat{\pi}^*} + \hat{h}^{\hat{\pi}^*} + g(\gamma)$ where $g(\gamma) \rightarrow 0$ as $\gamma \uparrow 1$. Combining these facts, we have that for all $\gamma \geq \hat{\gamma}_{\text{BW}}$,

$$\hat{V}_\gamma^* = \hat{V}_\gamma^{\hat{\pi}^*} = \frac{1}{1-\gamma} \hat{\rho}^{\hat{\pi}^*} + \hat{h}^{\hat{\pi}^*} + g(\gamma) = \frac{1}{1-\gamma} \hat{\rho}^* + \hat{h}^* + g(\gamma)$$

and also that $\hat{\rho}^*$ is a constant vector. Therefore, there exists sufficiently large γ (such that the $g(\gamma)$ term is bounded by $\frac{1}{n}$ in $\|\cdot\|_\infty$ norm) such that $\left\| \hat{V}_\gamma^* - \hat{h}^* - \frac{1}{1-\gamma} \hat{\rho}^* \mathbf{1} \right\|_\infty \leq \frac{1}{n}$, and thus the set in the definition of $\hat{\gamma}^*$ is nonempty.

Now we argue that the infimum is attained. We have already argued that the set

$$\left\{ \gamma \in [0, 1) : \exists c \in \mathbb{R} \text{ such that } \left\| \widehat{V}_\gamma^* - \widehat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n} \right\}$$

is nonempty and thus contains some $\overline{\gamma} \in [0, 1)$. Therefore we can write

$$\widehat{\gamma}^* = \inf \left\{ \gamma \in [0, \overline{\gamma}] : \exists c \in \mathbb{R} \text{ such that } \left\| \widehat{V}_\gamma^* - \widehat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n} \right\}. \quad (38)$$

Additionally, $\widehat{h}^*$ must have some entry which is ≤ 0 and some entry which is ≥ 0 (since $\widehat{P}_{\widehat{\pi}^*}^\infty \widehat{h}^* = 0$ and each row of $\widehat{P}_{\widehat{\pi}^*}^\infty$ is a probability distribution; note these may be the same entry). Furthermore, for all $\gamma \in [0, 1)$, $\widehat{V}_\gamma^* \in [0, \frac{1}{1-\gamma}]$. Thus if for some γ there exists $c \in \mathbb{R}$ such that $\left\| \widehat{V}_\gamma^* - \widehat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n}$, then letting $\widehat{h}^*(s) \leq 0$, we must have

$$\widehat{V}_\gamma^*(s) - \widehat{h}^*(s) - c \leq \left\| \widehat{V}_\gamma^* - \widehat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n}$$

which implies that

$$c \geq \widehat{V}_\gamma^*(s) - \widehat{h}^*(s) - \frac{1}{n} \geq 0 - 0 - \frac{1}{n} = -\frac{1}{n}.$$

Likewise looking at s such that $\widehat{h}^*(s) \geq 0$, we must have

$$\widehat{V}_\gamma^*(s) - \widehat{h}^*(s) - c \geq -\left\| \widehat{V}_\gamma^* - \widehat{h}^* - c\mathbf{1} \right\|_\infty \geq -\frac{1}{n}$$

which implies that

$$c \leq \widehat{V}_\gamma^*(s) - \widehat{h}^*(s) + \frac{1}{n} \leq \frac{1}{1-\gamma} - 0 + \frac{1}{n} = \frac{1}{1-\gamma} + \frac{1}{n}.$$

Therefore, the set

$$\left\{ \gamma \in [0, \overline{\gamma}], c \in \mathbb{R} : \left\| \widehat{V}_\gamma^* - \widehat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n} \right\} \quad (39)$$

is bounded since it is contained within $[0, \overline{\gamma}] \times \left[-\frac{1}{n}, \frac{1}{1-\overline{\gamma}} + \frac{1}{n} \right]$. Therefore it remains to show that the set (39) is closed, since this would imply that the set (39) is compact, and then since the continuous image of a compact set is compact and the projection of the set (39) onto its first coordinate is exactly the set in the expression (38) for $\widehat{\gamma}^*$, meaning that the infimum of this set is contained within the set, which is what we are trying to prove.

To show that the set (39) is closed, we first show that $\gamma \mapsto \widehat{V}_\gamma^*$ (with domain restricted to $[0, \overline{\gamma}]$) is a continuous function. This is a known result but we prove it for completeness. First, if π is fixed, then letting $\gamma, \gamma' \in [0, \overline{\gamma}]$, we have

$$\begin{aligned} \widehat{V}_\gamma^\pi - \widehat{V}_{\gamma'}^\pi &= (I - \gamma \widehat{P}_\pi)^{-1} r_\pi - (I - \gamma' \widehat{P}_\pi)^{-1} r_\pi \\ &= (I - \gamma \widehat{P}_\pi)^{-1} (I - \gamma' \widehat{P}_\pi) (I - \gamma' \widehat{P}_\pi)^{-1} r_\pi - (I - \gamma \widehat{P}_\pi)^{-1} (I - \gamma \widehat{P}_\pi) (I - \gamma' \widehat{P}_\pi)^{-1} r_\pi \\ &= (I - \gamma \widehat{P}_\pi)^{-1} \left[(I - \gamma' \widehat{P}_\pi) - (I - \gamma \widehat{P}_\pi) \right] (I - \gamma' \widehat{P}_\pi)^{-1} r_\pi \\ &= (I - \gamma \widehat{P}_\pi)^{-1} (\gamma - \gamma') \widehat{P}_\pi \widehat{V}_{\gamma'}^\pi \end{aligned}$$

and thus

$$\begin{aligned} \left\| \widehat{V}_\gamma^\pi - \widehat{V}_{\gamma'}^\pi \right\|_\infty &\leq |\gamma - \gamma'| \left\| (I - \gamma \widehat{P}_\pi)^{-1} \right\|_{\infty \rightarrow \infty} \left\| \widehat{P}_\pi \right\|_{\infty \rightarrow \infty} \left\| \widehat{V}_{\gamma'}^\pi \right\|_\infty \\ &\leq |\gamma - \gamma'| \frac{1}{1-\gamma} \frac{1}{1-\gamma'} \\ &\leq \frac{|\gamma - \gamma'|}{(1-\overline{\gamma})^2} \end{aligned}$$

so the function $\gamma \mapsto \widehat{V}_\gamma^\pi$ is Lipschitz and thus continuous. Now $\widehat{V}_\gamma^*$ is equal to the maximum over (the finite number of) all Markovian deterministic policies π , and thus $\gamma \mapsto \widehat{V}_\gamma^*$ is also continuous. This means that the function $f : [0, \overline{\gamma}] \times \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(\gamma, c) = \left\| \widehat{V}_\gamma^* - \widehat{h}^* - c\mathbf{1} \right\|_\infty$ is a continuous function, since we have shown that $\gamma \mapsto \widehat{V}_\gamma^*$ is continuous, and $\|\cdot\|_\infty$ and addition are continuous functions. Therefore the preimage of $[1, \frac{1}{n}]$ under f is a closed set since $[0, \frac{1}{n}]$ is closed. (Technically it we only immediately know that it is closed in the topology of the domain of f , $[0, \overline{\gamma}] \times \mathbb{R}$, but its closed sets are exactly the closed sets of $\mathbb{R}^2$ intersected with $[0, \overline{\gamma}] \times \mathbb{R}$ Pugh [2015].) Thus we have shown that the set (39) is closed, and thus as argued we can conclude that the infimum in the definition of $\widehat{\gamma}^*$ is attained. $\square$

We also remark that the asymptotic (partial) Laurent series expansion, used to bound $\widehat{\gamma}^*$ within this proof, could be replaced with a non-asymptotic version to give a more explicit bound.

Lemma 26. *If there exists $h \in \mathbb{R}^S$ such that π is greedy with respect to $r + \widehat{P}h$, then*

$$\widehat{V}_{1-\frac{1}{n}}^* - \widehat{V}_{1-\frac{1}{n}}^\pi \leq (n-1) \|\widehat{h}^* - h\|_{\text{span}}.$$

Proof. First we note a classic result for γ -discounted MDPs, that if there exists V such that π is greedy with respect to $r + \gamma \widehat{P}V$, then

$$\widehat{V}_\gamma^* - \widehat{V}_\gamma^\pi \leq \frac{2\gamma \|V - \widehat{V}_\gamma^*\|_\infty}{1 - \gamma} \quad (40)$$

[Singh and Yee, 1994]. The desired result will follow from (40) and the connection between AMDP and DMDP provided by Lemma 24. Specifically, we will try to find a vector V such that π is greedy with respect to $r + \gamma \widehat{P}V$ and such that $\|V - \widehat{V}_\gamma^*\|_\infty$ is small and bounded in terms of $\|\widehat{h}^* - h\|_{\text{span}}$.

Let $\gamma = 1 - \frac{1}{n}$. First note that if π is greedy with respect to $r + \widehat{P}h$, then it is also greedy with respect to $r + \widehat{P}h + \alpha\mathbf{1}$ for any α (since this shifts all entries by the same amount α). Now we try to choose α to meet the aforementioned conditions. First, note that

$$r + \widehat{P}h + \alpha\mathbf{1} = r + \gamma \widehat{P}h + (1 - \gamma)h(s_0)\mathbf{1} + \alpha\mathbf{1} = r + \gamma \widehat{P}(h + (1 - \gamma)h(s_0)\mathbf{1} + \alpha\mathbf{1})$$

(since $\widehat{P}\mathbf{1} = \mathbf{1}$) so we can define $V = h + (1 - \gamma)h(s_0)\mathbf{1} + \alpha\mathbf{1}$ and then try to minimize the quantity $\|V - \widehat{V}_\gamma^*\|_\infty$ by appropriately choosing α (as V is a function of α). We also know that $\widehat{h}^* = \widehat{V}_\gamma^* + \alpha^*\mathbf{1}$ for some α^* from Lemma 24. Therefore if we choose

$$\alpha = \frac{\max_s(\widehat{h}^*(s) - h(s)) + \min_s(\widehat{h}^*(s) - h(s))}{2} \mathbf{1} - \alpha^*\mathbf{1} - (1 - \gamma)h(s_0)\mathbf{1},$$

then

$$\begin{aligned} \|\widehat{V}_\gamma^* - V\|_\infty &= \|\widehat{h}^* - \alpha^*\mathbf{1} - h - (1 - \gamma)h(s_0)\mathbf{1} - \alpha\mathbf{1}\|_\infty \\ &= \left\| \widehat{h}^* - h - \frac{\max_s(\widehat{h}^*(s) - h(s)) + \min_s(\widehat{h}^*(s) - h(s))}{2} \mathbf{1} \right\|_\infty \\ &= \frac{\|\widehat{h}^* - h\|_{\text{span}}}{2}. \end{aligned}$$

Now we can conclude by applying (40) and noting that $\frac{\gamma}{1-\gamma} = \frac{1-\frac{1}{n}}{1-1+\frac{1}{n}} = n-1$. $\square$

Lemma 27. *For any $h \in \mathbb{R}^S$, let $\widehat{\mathcal{T}}(h) := M(r + \widehat{P}h)$ be the average-reward Bellman optimality operator for the anchored MDP $\widehat{\mathcal{P}} = (1 - \eta)\widehat{P} + \eta\mathbf{1}e_{s_0}^\top$. Suppose that one of the following conditions are satisfied for some policy π .*

1. π is greedy with respect to $r + \hat{P}h$ for some h such that $\|\hat{h}^\star - h\|_{\text{span}} \leq \frac{1}{n^2}$.

2. $\hat{\mathcal{T}}(\hat{h}^\pi) \leq \hat{h}^\pi + \hat{\rho}^\pi + \frac{1}{n^2}\mathbf{1}$.

3. $\hat{\rho}^\pi \geq \hat{\rho}^\star - \frac{1}{3n^2}$ and $\|\hat{h}^\pi - \hat{h}^\star\|_\infty \leq \frac{1}{3n^2}$.

Then we have that

$$\left\| \hat{V}_{1-\frac{1}{n}}^\star - \hat{V}_{1-\frac{1}{n}}^\pi \right\|_\infty \leq \frac{1}{n}. \quad (41)$$

Furthermore, if inequality (41) holds, we have $\hat{\rho}^\pi \geq \hat{\rho}^\star - \frac{1}{n^2}$ and $\|\hat{h}^\pi - \hat{h}^\star\|_{\text{span}} \leq \frac{2}{n}$.

Proof. The fact that the first condition implies (41) follows immediately from Lemma 26.

For the second condition, similar to the proof of Lemma 26, we first note an optimality condition for DMDPs, which we will later verify using the second condition. Letting $\hat{\mathcal{T}}_\gamma(V) := M(r + \gamma \hat{P}V)$ be the Bellman optimality operator for the γ -discounted MDP $\hat{P}$, for any policy π , we have

$$\left\| \hat{V}_\gamma^\star - \hat{V}_\gamma^\pi \right\|_\infty \leq \frac{\left\| \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi \right\|_\infty}{1 - \gamma}. \quad (42)$$

Inequality (42) is well-known but we give a proof for completeness: using γ -contractivity of $\hat{\mathcal{T}}_\gamma$,

$$\begin{aligned} \left\| \hat{V}_\gamma^\star - \hat{V}_\gamma^\pi \right\|_\infty &= \left\| \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\star) - \hat{V}_\gamma^\pi \right\|_\infty \\ &\leq \left\| \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\star) - \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) \right\|_\infty + \left\| \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi \right\|_\infty \\ &\leq \gamma \left\| \hat{V}_\gamma^\star - \hat{V}_\gamma^\pi \right\|_\infty + \left\| \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi \right\|_\infty \end{aligned}$$

which implies $\left\| \hat{V}_\gamma^\star - \hat{V}_\gamma^\pi \right\|_\infty \leq \frac{\left\| \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi \right\|_\infty}{1 - \gamma}$ after rearranging. Now we relate condition 2 from the lemma to the quantity $\left\| \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi \right\|_\infty$ from (42). Letting $\gamma = 1 - \frac{1}{n}$, by Lemma 24 we have that $\hat{V}_\gamma^\pi = \hat{h}^\pi + c\mathbf{1}$ for some scalar c . We also have that

$$\hat{\mathcal{L}}(x) = M(r + \hat{P}x) = M(r + (1 - \eta)\hat{P}x + \eta x(s_0)\mathbf{1}) = M(r + (1 - \eta)\hat{P}x) + \eta x(s_0)\mathbf{1} = \hat{\mathcal{T}}_\gamma(x) + \eta x(s_0)\mathbf{1}$$

for any $x \in \mathbb{R}^S$. Then we can calculate

$$\begin{aligned} \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi &= \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{h}^\pi - c\mathbf{1} \\ &= \hat{\mathcal{L}}(\hat{V}_\gamma^\pi) - \eta \hat{V}_\gamma^\pi(s_0)\mathbf{1} - \hat{h}^\pi - c\mathbf{1} \\ &= \hat{\mathcal{L}}(\hat{h}^\pi + c\mathbf{1}) - \eta \hat{V}_{1-\eta}^\pi(s_0)\mathbf{1} - \hat{h}^\pi - c\mathbf{1} \\ &= \hat{\mathcal{L}}(\hat{h}^\pi) - \eta \hat{V}_{1-\eta}^\pi(s_0)\mathbf{1} - \hat{h}^\pi \\ &= \hat{\mathcal{L}}(\hat{h}^\pi) - \hat{\rho}^\pi - \hat{h}^\pi \end{aligned} \quad (43)$$

where in the last step we used the fact that $\eta \hat{V}_{1-\eta}^\pi(s_0)\mathbf{1} = \hat{\rho}^\pi$ by Lemma 24. Also we have that $\hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi \geq 0$ (this is a standard fact, but to see this note that $\hat{V}_\gamma^\pi$ satisfies the Bellman equation $\hat{V}_\gamma^\pi = M^\pi(r + \gamma \hat{P}\hat{V}_\gamma^\pi)$, and thus $\hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi = M(r + \gamma \hat{P}\hat{V}_\gamma^\pi) - M^\pi(r + \gamma \hat{P}\hat{V}_\gamma^\pi) \geq M^\pi(r + \gamma \hat{P}\hat{V}_\gamma^\pi) - M^\pi(r + \gamma \hat{P}\hat{V}_\gamma^\pi) = 0$.) Thus combining this with (43), we have the equivalence

$$\left\| \hat{\mathcal{T}}_\gamma(\hat{V}_\gamma^\pi) - \hat{V}_\gamma^\pi \right\|_\infty \leq \varepsilon \iff \hat{\mathcal{L}}(\hat{h}^\pi) - \hat{\rho}^\pi - \hat{h}^\pi \leq \varepsilon \mathbf{1}.$$

Therefore, by this equivalence, if $\widehat{\mathcal{T}}(\widehat{h}^\pi) - \widehat{\rho}^\pi - \widehat{h}^\pi \leq \frac{1}{n^2} \mathbf{1}$, then $\|\widehat{\mathcal{T}}_\gamma(\widehat{V}_\gamma^\pi) - \widehat{V}_\gamma^\pi\|_\infty \leq \frac{1}{n^2}$, and plugging into (42), we obtain that

$$\widehat{V}_\gamma^\star - \widehat{V}_\gamma^\pi \leq \frac{\|\widehat{\mathcal{T}}(\widehat{V}_\gamma^\pi) - \widehat{V}_\gamma^\pi\|_\infty}{1 - \gamma} \leq \frac{1/n^2}{1/n} = \frac{1}{n}$$

as desired. Thus we have justified the second condition.

Next we will show that the third condition implies the second condition, and thus it also implies that $\|\widehat{V}_\gamma^\star - \widehat{V}_\gamma^\pi\|_\infty \leq \frac{1}{n}$. We can write

$$\begin{aligned} \widehat{\mathcal{T}}(\widehat{h}^\pi) - \widehat{h}^\pi - \widehat{\rho}^\pi &= \widehat{\mathcal{T}}(\widehat{h}^\pi) - \widehat{\mathcal{T}}(\widehat{h}^\star) + \widehat{\mathcal{T}}(\widehat{h}^\star) - \widehat{h}^\star + \widehat{h}^\star - \widehat{h}^\pi - \widehat{\rho}^\star + \widehat{\rho}^\star - \widehat{\rho}^\pi \\ &\leq \|\widehat{\mathcal{T}}(\widehat{h}^\pi) - \widehat{\mathcal{T}}(\widehat{h}^\star)\|_\infty \mathbf{1} + \widehat{\mathcal{T}}(\widehat{h}^\star) - \widehat{h}^\star - \widehat{\rho}^\star + \|\widehat{h}^\star - \widehat{h}^\pi\|_\infty \mathbf{1} + \|\widehat{\rho}^\star - \widehat{\rho}^\pi\|_\infty \mathbf{1} \\ &\leq \gamma \|\widehat{h}^\star - \widehat{h}^\pi\|_\infty \mathbf{1} + 0 + \|\widehat{h}^\star - \widehat{h}^\pi\|_\infty \mathbf{1} + \|\widehat{\rho}^\star - \widehat{\rho}^\pi\|_\infty \mathbf{1} \\ &\leq \frac{\gamma}{3n^2} \mathbf{1} + \frac{1}{3n^2} \mathbf{1} + \frac{1}{3n^2} \mathbf{1} \leq \frac{1}{n^2} \mathbf{1} \end{aligned}$$

where in the penultimate inequality we use the assumptions of condition 3.

Finally, we assume that condition (41) holds and try to show that $\widehat{\rho}^\pi \geq \widehat{\rho}^\star - \frac{1}{n^2}$ and $\|\widehat{h}^\pi - \widehat{h}^\star\|_{\text{span}} \leq \frac{2}{n}$. First, by Lemma 24 we have $\widehat{\rho}^\pi = \frac{\widehat{V}_{1-\frac{1}{n}}^\pi(s_0)}{n} \mathbf{1}$ and $\widehat{\rho}^\star = \frac{\widehat{V}_{1-\frac{1}{n}}^\star(s_0)}{n} \mathbf{1}$. Combining with condition (41) we can obtain that

$$\begin{aligned} \widehat{\rho}^\pi &= \frac{\widehat{V}_{1-\frac{1}{n}}^\pi(s_0)}{n} \mathbf{1} \geq \frac{\widehat{V}_{1-\frac{1}{n}}^\star(s_0) - |\widehat{V}_{1-\frac{1}{n}}^\star(s_0) - \widehat{V}_{1-\frac{1}{n}}^\pi(s_0)|}{n} \mathbf{1} \\ &\geq \frac{\widehat{V}_{1-\frac{1}{n}}^\star(s_0) - \|\widehat{V}_{1-\frac{1}{n}}^\star - \widehat{V}_{1-\frac{1}{n}}^\pi\|_\infty}{n} \mathbf{1} \geq \widehat{\rho}^\star - \frac{1}{n^2} \mathbf{1}. \end{aligned}$$

Second, note that by Lemma 24 we have $\|\widehat{h}^\pi - \widehat{h}^\star\|_{\text{span}} = \|\widehat{V}_{1-\frac{1}{n}}^\pi - \widehat{V}_{1-\frac{1}{n}}^\star\|_{\text{span}}$, so we have that

$$\|\widehat{h}^\pi - \widehat{h}^\star\|_{\text{span}} = \|\widehat{V}_{1-\frac{1}{n}}^\pi - \widehat{V}_{1-\frac{1}{n}}^\star\|_{\text{span}} \leq 2 \|\widehat{V}_{1-\frac{1}{n}}^\pi - \widehat{V}_{1-\frac{1}{n}}^\star\|_\infty \leq \frac{2}{n}.$$

□

Lemma 28. *Under the same event that the conclusions of Theorem 6 hold, there is a unique Markovian bias-optimal policy $\widehat{\pi}_p^\star$ for the AMDP $(\widehat{P}, \widehat{r})$, and thus this policy is also the unique Markovian Blackwell optimal policy. The policy $\widehat{\pi}_p^\star$ is deterministic. Furthermore, $\widehat{\pi}_p^\star$ can be computed with $K = n \lceil \log \left(\frac{12n^3 S A^2}{\xi \delta} \right) \rceil$ iterations of average-reward value iteration: Letting ${}_p \widehat{\mathcal{T}}(h) := M(\widehat{r} + \widehat{P}h)$ be the average-reward Bellman operator for the AMDP $(\widehat{P}, \widehat{r})$, $\widehat{\pi}_p^\star$ is the greedy policy with respect to ${}_p \widehat{\mathcal{T}}^{(K)}(0)$.*

Proof. First we note that under the event that the conclusions of Theorem 6 hold, we have the separation property described in Theorem 21 with $\gamma = 1 - \frac{1}{n}$, namely that the optimal Markovian policy $\widehat{\pi}_p^\star$ in the DMDP $(\widehat{P}, \widehat{r}, 1 - \frac{1}{n})$ is unique and deterministic, and for all $s \in \mathcal{S}, a \in \mathcal{A}$ such that $a \neq \widehat{\pi}_p^\star(s)$,

$${}_p \widehat{Q}_{1-1/n}^\star(s, \widehat{\pi}_p^\star(s)) - {}_p \widehat{Q}_{1-1/n}^\star(s, a) \geq \frac{\xi \delta}{3n S A^2}. \quad (44)$$

(This property is used within the proof of Theorem 6, namely within Lemma 22).

First we check that $\widehat{\pi}_p^\star$ is the unique bias-optimal policy. Note that since Blackwell-optimal implies bias-optimal, this would imply that $\widehat{\pi}_p^\star$ is also the unique Blackwell-optimal policy. Since any bias-optimal policy π satisfies ${}_p \widehat{h}^\pi = {}_p \widehat{h}^\star$, and ${}_p \widehat{\mathcal{T}}({}_p \widehat{h}^\star) = \rho^\star + {}_p \widehat{h}^\star$, and also $\rho^\star + {}_p \widehat{h}^\pi = M^\pi(r + \widehat{P} {}_p \widehat{h}^\pi)$, we must have that π

is greedy with respect to $\tilde{r} + \hat{\underline{P}}_p \hat{\underline{h}}^*$, that is $M(\tilde{r} + \hat{\underline{P}}_p \hat{\underline{h}}^*) = M^\pi(\tilde{r} + \hat{\underline{P}}_p \hat{\underline{h}}^*)$. By the definition of ${}_p\hat{Q}_{1-1/n}^*$, the definition of $\hat{\underline{P}}$, and Lemma 24 (which ensures ${}_p\hat{V}^* = {}_p\hat{\underline{h}}^* + c\mathbf{1}$), we have

$$\begin{aligned} {}_p\hat{Q}_{1-1/n}^*(s, a) &= \tilde{r} + (1 - \frac{1}{n})\hat{\underline{P}}_p \hat{\underline{V}}_{1-1/n}^* = \tilde{r} + \hat{\underline{P}}_p \hat{\underline{V}}_{1-1/n}^* - \eta \mathbf{1} {}_p\hat{V}_{1-1/n}^*(s_0) \\ &= \tilde{r} + \hat{\underline{P}}_p \hat{\underline{h}}^* + \hat{\underline{P}}_p c\mathbf{1} - \eta \mathbf{1} {}_p\hat{V}^*(s_0) = \tilde{r} + \hat{\underline{P}}_p \hat{\underline{h}}^* + c'\mathbf{1} \end{aligned}$$

for some scalars c, c' . Therefore π is also greedy with respect to ${}_p\hat{Q}_{1-1/n}^*$, and by the separation property (44), this implies that we must have $\pi = \hat{\pi}_p^*$. Therefore the unique bias- and Blackwell-optimal policy is $\hat{\pi}_p^*$.

To show finite convergence we can combine the separation condition (44) with Lemma 26, which also holds with the perturbed $\tilde{r}$ reward function and thus guarantees that if there exists $h \in \mathbb{R}^S$ such that π is greedy with respect to $\tilde{r} + \hat{\underline{P}}h$, then

$${}_p\hat{V}_{1-\frac{1}{n}}^* - {}_p\hat{V}_{1-\frac{1}{n}}^\pi \leq (n-1)\|{}_p\hat{\underline{h}}^* - h\|_{\text{span}}. \quad (45)$$

where ${}_p\hat{\underline{h}}^*$ is the optimal bias function of the perturbed anchored empirical AMDP $(\hat{\underline{P}}, \tilde{r})$. From Lemma 24 we have that ${}_p\hat{\underline{T}}$ is a $1 - \eta = 1 - \frac{1}{n}$ span-contraction, and it has fixed point ${}_p\hat{\underline{h}}^*$. Thus

$$\begin{aligned} \|{}_p\hat{\underline{T}}^{(K)}(0) - {}_p\hat{\underline{h}}^*\|_{\text{span}} &= \|{}_p\hat{\underline{T}}^{(K)}(0) - {}_p\hat{\underline{T}}^{(K)}({}_p\hat{\underline{h}}^*)\|_{\text{span}} \\ &\leq \left(1 - \frac{1}{n}\right)^K \|{}_p\hat{\underline{h}}^* - 0\|_{\text{span}} \\ &\leq \left(\left(1 - \frac{1}{n}\right)^n\right)^{\lceil \log\left(\frac{12n^3SA^2}{\xi\delta}\right) \rceil} 2n \\ &\leq \left(\frac{1}{e}\right)^{\log\left(\frac{12n^3SA^2}{\xi\delta}\right)} 2n \\ &\leq \frac{\xi\delta}{6n^2SA^2} \end{aligned} \quad (46)$$

where we used that

$$\|{}_p\hat{\underline{h}}^* - 0\|_{\text{span}} = \|{}_p\hat{\underline{h}}^*\|_{\text{span}} = \|{}_p\hat{V}_{1-1/n}\|_{\text{span}} \leq \|{}_p\hat{V}_{1-1/n}\|_\infty \leq \frac{1}{1 - (1 - \frac{1}{n})} \|\tilde{r}\|_\infty \leq n(1 + \frac{1}{n}) \leq 2n$$

and also that $(1 - \frac{1}{n})^n \leq \frac{1}{e}$. Now combining the bound (46) with (45), if π is chosen to be deterministic and greedy with respect to $\tilde{r} + \hat{\underline{P}}_p \hat{\underline{T}}^{(K)}(0)$, we have that ${}_p\hat{V}_{1-\frac{1}{n}}^* - {}_p\hat{V}_{1-\frac{1}{n}}^\pi \leq \frac{\xi\delta}{6nSA^2}$. By (44), if π is a deterministic policy, then this implies $\pi = \hat{\pi}_p^*$. However, if π is not deterministic, then it must be possible to write π as the convex combination of some distinct policies which are deterministic and such that each of these policies is also greedy with respect to $\tilde{r} + \hat{\underline{P}}_p \hat{\underline{T}}^{(K)}(0)$, but then each of these deterministic policies must also be equal to $\hat{\pi}_p^*$, giving a contradiction, so π must have been deterministic. $\square$

C.2 Proof of Theorem 6

Here we complete the proof of Theorem 6. Since Lemma 24 relates the gains and bias functions of anchored AMDPs to DMDPs, we can prove this theorem by simply combining Lemma 24 with Theorem 10.

Proof of Theorem 6. Following the conditions listed in Theorem 6, let $\hat{\pi}$ be an exact Blackwell-optimal policy of the AMDP $(\hat{\underline{P}}, \tilde{r})$.

First, by using triangle inequality twice, we have the elementwise inequality

$$\rho^{\pi^*} - \rho^{\hat{\pi}} \leq \left\| \rho^{\hat{\pi}} - \underline{\rho}^{\hat{\pi}} \right\|_\infty \mathbf{1} + \left(\underline{\rho}^{\pi^*} - \underline{\rho}^{\hat{\pi}} \right) + \left\| \underline{\rho}^{\pi^*} - \rho^{\pi^*} \right\|_\infty \mathbf{1}. \quad (47)$$

Note that the first and third terms of (47) are controlled by Lemma 24 as

$$\left\| \rho^{\pi^*} - \underline{\rho}^{\pi^*} \right\|_{\infty} \leq \frac{2 \|\underline{h}^{\pi^*}\|_{\text{span}}}{n} \leq \frac{4 \|h^{\pi^*}\|_{\text{span}}}{n} = \frac{4 \|h^*\|_{\text{span}}}{n} \quad (48)$$

(because $\rho^{\pi^*} = \rho^*$ is constant, for the second inequality) and

$$\left\| \rho^{\hat{\pi}} - \underline{\rho}^{\hat{\pi}} \right\|_{\infty} \leq \frac{2 \|\underline{h}^{\hat{\pi}}\|_{\text{span}}}{n}. \quad (49)$$

For the middle term on the RHS of (47), using the identity for the gain in the anchored AMDP from Lemma 24, we have

$$\begin{aligned} \underline{\rho}^{\pi^*} - \underline{\rho}^{\hat{\pi}} &= \frac{V_{1-\frac{1}{n}}^{\pi^*}(s_0)}{n} - \frac{V_{1-\frac{1}{n}}^{\hat{\pi}}(s_0)}{n} \\ &\leq \frac{V_{1-\frac{1}{n}}^*(s_0)}{n} - \frac{V_{1-\frac{1}{n}}^{\hat{\pi}}(s_0)}{n} \\ &\leq \frac{1}{n} \left\| V_{1-\frac{1}{n}}^* - V_{1-\frac{1}{n}}^{\hat{\pi}} \right\|_{\infty}. \end{aligned} \quad (50)$$

On the high-probability event in the conclusion of Theorem 10, by Lemma 28, the policy $\hat{\pi}$ (defined as a Blackwell-optimal policy for $(\hat{P}, \hat{r})$) is identical to the optimal policy for the DMDP $(\hat{P}, \hat{r}, 1-1/n)$. Therefore Theorem 10 (which is stated for the optimal policy of the DMDP $(\hat{P}, \hat{r}, 1-1/n)$) also applies to $\hat{\pi}$ as defined in this proof. Thus by Theorem 10, with probability at least $1 - \delta$,

$$\begin{aligned} \left\| V_{1-\frac{1}{n}}^* - V_{1-\frac{1}{n}}^{\hat{\pi}} \right\|_{\infty} &\leq \frac{1}{1 - (1 - \frac{1}{n})} \sqrt{\frac{C_2 \log^3 \left(\frac{SAn}{(1-(1-\frac{1}{n}))\delta\xi} \right)}{n} \left(\|V_{1-\frac{1}{n}}^*\|_{\text{span}} + \|V_{1-\frac{1}{n}}^{\hat{\pi}}\|_{\text{span}} + 1 \right)} \\ &= n \sqrt{\frac{C_2 \log^3 \left(\frac{SAn^2}{\delta\xi} \right)}{n} \left(\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right)} \\ &\leq n \sqrt{\frac{C_2 \log^3 \left(\frac{SAn^2}{\delta\xi} \right)}{n} \left(2\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right)} \end{aligned}$$

(using Lemma 24 in the second two steps, specifically the facts that $\|V_{1-\frac{1}{n}}^*\|_{\text{span}} = \|\underline{h}^*\|_{\text{span}}$, $\|V_{1-\frac{1}{n}}^{\hat{\pi}}\|_{\text{span}} = \|\underline{h}^{\hat{\pi}}\|_{\text{span}}$, and then that $\|\underline{h}^*\|_{\text{span}} \leq 2\|h^*\|_{\text{span}}$ since ρ^* is constant). Combining this with (50),

$$\underline{\rho}^{\pi^*} - \underline{\rho}^{\hat{\pi}} \leq \sqrt{\frac{C_2 \log^3 \left(\frac{SAn^2}{\delta\xi} \right)}{n} \left(2\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right)}. \quad (51)$$

Combining (48), (49), and (51) with (47), and then simplifying, we obtain

$$\begin{aligned} \rho^{\pi^*} - \rho^{\hat{\pi}} &\leq \frac{4\|h^*\|_{\text{span}}}{n} + \sqrt{\frac{C_2 \log^3 \left(\frac{SAn^2}{\delta\xi} \right)}{n} \left(2\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right)} + \frac{2\|\underline{h}^{\hat{\pi}}\|_{\text{span}}}{n} \\ &\leq \sqrt{\frac{4\|h^*\|_{\text{span}}}{n}} + \sqrt{\frac{C_2 \log^3 \left(\frac{SAn^2}{\delta\xi} \right)}{n} \left(2\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right)} + \sqrt{\frac{2\|\underline{h}^{\hat{\pi}}\|_{\text{span}}}{n}} \\ &\leq 3\sqrt{\frac{4\|h^*\|_{\text{span}}}{n} + \frac{C_2 \log^3 \left(\frac{SAn^2}{\delta\xi} \right)}{n} \left(2\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right) + \frac{2\|\underline{h}^{\hat{\pi}}\|_{\text{span}}}{n}} \\ &\leq \sqrt{\frac{C_3 \log^3 \left(\frac{SAn}{\delta\xi} \right)}{n} \left(2\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right)} \end{aligned}$$

where the second inequality holds since the terms $\frac{4\|h^*\|_{\text{span}}}{n}, \frac{2\|\underline{h}\|_{\text{span}}}{n}$ must be ≤ 1 or else the bound holds trivially, the third inequality uses $\sqrt{a} + \sqrt{b} + \sqrt{c} \leq 3\sqrt{a+b+c}$, and the final inequality sets C_3 sufficiently large. $\square$

C.3 Higher-order variance bounds

Lemma 29.

$$\left(P_\pi(h^\pi)^{\circ 2^k}\right)^{\circ 2} \geq (P_\pi h)^{\circ 2^{k+1}}.$$

Proof. For arbitrary $\ell \geq 0$, we have (elementwise)

$$0 \leq \mathbb{V}_{P_\pi} \left[(h^\pi)^{\circ 2^\ell} \right] = P_\pi(h^\pi)^{\circ 2^{\ell+1}} - \left(P_\pi(h^\pi)^{\circ 2^\ell}\right)^{\circ 2}$$

so $\left(P_\pi(h^\pi)^{\circ 2^\ell}\right)^{\circ 2} \leq P_\pi(h^\pi)^{\circ 2^{\ell+1}}$. Now we apply this fact k times to obtain the desired conclusion:

$$\begin{aligned} \left(P_\pi(h^\pi)^{\circ 2^k}\right)^{\circ 2} &\geq \left(\left(P_\pi(h^\pi)^{\circ 2^{k-1}}\right)^{\circ 2}\right)^{\circ 2} = \left(P_\pi(h^\pi)^{\circ 2^{k-1}}\right)^{\circ 2^2} \\ &\vdots \\ &\geq \left(P_\pi(h^\pi)^{\circ 2^{k-j}}\right)^{\circ 2^{j+1}} \\ &\vdots \\ &\geq (P_\pi h)^{\circ 2^{k+1}}. \end{aligned}$$

$\square$

Lemma 30. Fix an integer $k \geq 0$ and let $\bar{h} = h^\pi - (\min_s h^\pi(s)) \mathbf{1}$. If

$$\left|(\hat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^k}\right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_\pi}[(\bar{h})^{\circ 2^k}]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \quad (52)$$

holds, then

$$\begin{aligned} &\left(\frac{\alpha}{n}\right)^{1-2^{-k}} \left|\hat{P}_\pi^\infty(\hat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^k}\right|^{2^{-k}} \\ &\leq \left(\frac{\alpha}{n}\right)^{1-2^{-(k+1)}} \left|\hat{P}_\pi^\infty(\hat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^{k+1}}\right|^{2^{-(k+1)}} + 2 \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n}\right)^{1-2^{-(k+1)}} \mathbf{1} + \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1) \mathbf{1}. \end{aligned}$$

Proof.

$$\left|\hat{P}_\pi^\infty(\hat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^k}\right| \leq \hat{P}_\pi^\infty \left|(\hat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^k}\right| \quad (53)$$

$$\leq \hat{P}_\pi^\infty \sqrt{\frac{\alpha \mathbb{V}_{P_\pi}[(\bar{h})^{\circ 2^k}]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \quad (54)$$

using the fact that all entries of $\hat{P}_\pi^\infty$ are non-negative, and then the condition (52). Focusing on the first term in (54),

$$\hat{P}_\pi^\infty \sqrt{\mathbb{V}_{P_\pi}[(\bar{h})^{\circ 2^k}]} \leq \sqrt{\hat{P}_\pi^\infty \mathbb{V}_{P_\pi}[(\bar{h})^{\circ 2^k}]} \quad (55)$$

$$= \sqrt{\hat{P}_\pi^\infty \left(P_\pi(\bar{h})^{\circ 2^{k+1}} - (P_\pi(\bar{h})^{\circ 2^k})^{\circ 2} \right)} \quad (56)$$

$$\leq \sqrt{\hat{P}_\pi^\infty \left(P_\pi(\bar{h})^{\circ 2^{k+1}} - (P_\pi \bar{h})^{\circ 2^{k+1}} \right)} \quad (57)$$

$$= \sqrt{\hat{P}_\pi^\infty \left(P_\pi(\bar{h})^{\circ 2^{k+1}} - (\bar{h} + P_\pi^\infty r_\pi - r_\pi)^{\circ 2^{k+1}} \right)} \quad (58)$$

where we used Jensen's inequality since each row of $\hat{P}_\pi^\infty$ is a probability distribution, the definition of $\mathbb{V}_{P_\pi}[(\bar{h})^{\circ 2^k}]$, the non-negativity of $\hat{P}_\pi^\infty$ along with the inequality from Lemma 29, and then finally the Bellman equation $h^\pi + P_\pi^\infty r_\pi = r_\pi + P_\pi h^\pi$, which after subtracting $(\min_s h^\pi(s)) \mathbf{1}$ from both sides yields

$$\begin{aligned} \bar{h} + P_\pi^\infty r_\pi &= h^\pi - \left(\min_s h^\pi(s) \right) \mathbf{1} + P_\pi^\infty r_\pi \\ &= r_\pi + P_\pi h^\pi - \left(\min_s h^\pi(s) \right) \mathbf{1} \\ &= r_\pi + P_\pi \left(h^\pi - \left(\min_s h^\pi(s) \right) \mathbf{1} \right) \\ &= r_\pi + P_\pi \bar{h} \end{aligned}$$

since $P_\pi \mathbf{1} = \mathbf{1}$. Now note that if we expand $(\bar{h} + P_\pi^\infty r_\pi - r_\pi)^{\circ 2^{k+1}}$ into the sum of 2^{k+1} individual terms, the leading term will be $(\bar{h})^{\circ 2^{k+1}}$, while all other terms will be (entrywise) bounded in magnitude by $\max\{\|\bar{h}\|_\infty, 1\}^{2^{k+1}-1}$ since $\|P_\pi^\infty r_\pi - r_\pi\|_\infty \leq 1$. Thus resumming from (58), we have

$$\begin{aligned} &\sqrt{\hat{P}_\pi^\infty \left(P_\pi(\bar{h})^{\circ 2^{k+1}} - (\bar{h} + P_\pi^\infty r_\pi - r_\pi)^{\circ 2^{k+1}} \right)} \\ &\leq \sqrt{\hat{P}_\pi^\infty \left(P_\pi(\bar{h})^{\circ 2^{k+1}} - (\bar{h})^{\circ 2^{k+1}} + 2^{2^{k+1}} \max\{\|\bar{h}\|_\infty, 1\}^{2^{k+1}-1} \mathbf{1} \right)} \\ &\leq \sqrt{\hat{P}_\pi^\infty \left(P_\pi(\bar{h})^{\circ 2^{k+1}} - (\bar{h})^{\circ 2^{k+1}} + 2^{2^{k+1}} (\|\bar{h}\|_\infty + 1)^{2^{k+1}-1} \mathbf{1} \right)} \\ &= \sqrt{\hat{P}_\pi^\infty (P_\pi - I) (\bar{h})^{\circ 2^{k+1}} + 2^{2^{k+1}} (\|\bar{h}\|_\infty + 1)^{2^{k+1}-1} \mathbf{1}} \\ &= \sqrt{\hat{P}_\pi^\infty (P_\pi - \hat{P}_\pi) (\bar{h})^{\circ 2^{k+1}} + 2^{2^{k+1}} (\|\bar{h}\|_\infty + 1)^{2^{k+1}-1} \mathbf{1}} \\ &\leq \left| \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{1/2} + 2^{2^k} (\|\bar{h}\|_\infty + 1)^{\frac{1}{2}(2^{k+1}-1)} \mathbf{1} \end{aligned}$$

where in the final equality step we use that $\hat{P}_\pi^\infty \hat{P}_\pi = \hat{P}_\pi^\infty$, and then in the final inequality step we use that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$. Combining these steps we have that

$$\begin{aligned} &\left| \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^k} \right| \\ &\leq \hat{P}_\pi^\infty \sqrt{\frac{\alpha \mathbb{V}_{P_\pi}[(\bar{h})^{\circ 2^k}]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \\ &\leq \sqrt{\frac{\alpha}{n}} \left(\left| \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{1/2} + 2^{2^k} (\|\bar{h}\|_\infty + 1)^{\frac{1}{2}(2^{k+1}-1)} \mathbf{1} \right) + \frac{\alpha \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \\ &= \left| \frac{\alpha}{n} \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{1/2} + 2^{2^k} \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)^{2^{k+1}-1}}{n} \right)^{1/2} \mathbf{1} + \frac{\alpha \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1}. \end{aligned}$$

Therefore

$$\begin{aligned}
& \left(\frac{\alpha}{n}\right)^{1-2^{-k}} \left| \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^k} \right|^{2^{-k}} \\
& \leq \left(\frac{\alpha}{n}\right)^{1-2^{-k}} \left(\left| \frac{\alpha}{n} \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{1/2} + 2^{2^k} \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)^{2^{k+1}-1}}{n} \right)^{1/2} \mathbf{1} \right. \\
& \quad \left. + \frac{\alpha \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \right)^{2^{-k}} \\
& \leq \left(\frac{\alpha}{n}\right)^{1-2^{-k}} \left(\left| \frac{\alpha}{n} \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{2^{-k}/2} + 2^{2^k \cdot 2^{-k}} \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)^{2^{k+1}-1}}{n} \right)^{2^{-k}/2} \mathbf{1} \right. \\
& \quad \left. + \left(\frac{\alpha \cdot 2^k}{n} \right)^{2^{-k}} (\|\bar{h}\|_\infty + 1)^{2^k \cdot 2^{-k}} \mathbf{1} \right) \\
& \leq \left(\frac{\alpha}{n}\right)^{1-2^{-k}} \left(\left| \frac{\alpha}{n} \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{2^{-k}/2} + 2^{2^k \cdot 2^{-k}} \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)^{2^{k+1}-1}}{n} \right)^{2^{-k}/2} \mathbf{1} \right. \\
& \quad \left. + 2 \left(\frac{\alpha}{n} \right)^{2^{-k}} (\|\bar{h}\|_\infty + 1)^{2^k \cdot 2^{-k}} \mathbf{1} \right) \\
& = \left(\frac{\alpha}{n}\right)^{1-2^{-k}} \left(\left| \frac{\alpha}{n} \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{2^{-(k+1)}} + 2 \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)^{2^{k+1}-1}}{n} \right)^{2^{-(k+1)}} \mathbf{1} \right. \\
& \quad \left. + 2 \left(\frac{\alpha}{n} \right)^{2^{-k}} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \right) \\
& = \left(\frac{\alpha}{n}\right)^{1-2^{-k}} \left(\left| \frac{\alpha}{n} \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{2^{-(k+1)}} + 2 \left(\frac{\alpha}{n} \right)^{2^{-(k+1)}} (\|\bar{h}\|_\infty + 1)^{1-2^{-(k+1)}} \mathbf{1} \right. \\
& \quad \left. + 2 \left(\frac{\alpha}{n} \right)^{2^{-k}} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \right) \\
& = \left(\frac{\alpha}{n}\right)^{1-2^{-(k+1)}} \left| \hat{P}_\pi^\infty (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^{k+1}} \right|^{2^{-(k+1)}} + 2 \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)^{1-2^{-(k+1)}}}{n} \right)^{1-2^{-(k+1)}} \mathbf{1} + \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1}
\end{aligned}$$

as desired, where for the inequality steps we used the previous calculations, then that $(a + b + c)^{2^{-k}} \leq a^{2^{-k}} + b^{2^{-k}} + c^{2^{-k}}$, and then that $2^k \leq 2^{2^k}$ so $(2^k)^{2^{-k}} \leq 2^{2^k \cdot 2^{-k}} = 2$. $\square$

Lemma 31. Let $\bar{h} = h^\pi - (\min_s h^\pi(s)) \mathbf{1}$ and $\ell = \lceil \log_2 \log_2 (\|\bar{h}\|_\infty + 4) \rceil$. Suppose that for some $\alpha \in \mathbb{R}$, the inequalities

$$\left| (\hat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_\pi} [(\bar{h})^{\circ 2^k}]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1}$$

hold for all $k = 0, \dots, \ell$. Also suppose that ρ^π is constant. Then

$$\|\widehat{\rho}^\pi - \rho^\pi\|_\infty \leq 2(\ell + 1) \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n} \right)^{\frac{1}{2}} + (\ell + 1) \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1).$$

We also note that this Lemma is purely algebraic, and thus we can accordingly replace the transition matrices $P, \widehat{P}$ (and all their derived quantities, that is, $h^\pi, \rho^\pi, \widehat{\rho}^\pi$) with any other pair of transition matrices.

Proof. Similarly to the proof of Lemma 16, first we give a weaker but non-recursive bound which can be used on the final term. Note that

$$\begin{aligned} \left| (\widehat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^\ell} \right| &\leq \sqrt{\frac{\alpha \mathbb{V}_{P_\pi}[(\bar{h})^{\circ 2^\ell}]}{n}} \mathbf{1} + \frac{\alpha \cdot 2^\ell}{n} (\|\bar{h}\|_\infty + 1)^{2^\ell} \mathbf{1} \\ &\leq \sqrt{\frac{\alpha}{n}} \left\| (\bar{h})^{\circ 2^\ell} \right\|_\infty \mathbf{1} + \frac{\alpha \cdot 2^\ell}{n} (\|\bar{h}\|_\infty + 1)^{2^\ell} \mathbf{1} \\ &= \sqrt{\frac{\alpha}{n}} \|\bar{h}\|_\infty^{2^\ell} \mathbf{1} + \frac{\alpha \cdot 2^\ell}{n} (\|\bar{h}\|_\infty + 1)^{2^\ell} \mathbf{1} \end{aligned} \quad (59)$$

and so (elementwise)

$$\begin{aligned} \left| \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^\ell} \right| &\leq \widehat{P}_\pi^\infty \left| (\widehat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^\ell} \right| \\ &\leq \widehat{P}_\pi^\infty \left(\sqrt{\frac{\alpha}{n}} \|\bar{h}\|_\infty^{2^\ell} \mathbf{1} + \frac{\alpha \cdot 2^\ell}{n} (\|\bar{h}\|_\infty + 1)^{2^\ell} \mathbf{1} \right) \\ &= \sqrt{\frac{\alpha}{n}} \|\bar{h}\|_\infty^{2^\ell} \mathbf{1} + \frac{\alpha \cdot 2^\ell}{n} (\|\bar{h}\|_\infty + 1)^{2^\ell} \mathbf{1} \end{aligned}$$

using that all entries of $\widehat{P}_\pi^\infty$ are non-negative in the first inequality, then the bound (59), then the fact $\widehat{P}_\pi^\infty \mathbf{1} = \mathbf{1}$. Therefore

$$\begin{aligned} \left(\frac{\alpha}{n} \right)^{1-2^{-\ell}} \left| \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^\ell} \right|^{2^{-\ell}} &\leq \left(\frac{\alpha}{n} \right)^{1-2^{-\ell}} \left(\sqrt{\frac{\alpha}{n}} \|\bar{h}\|_\infty^{2^\ell} + \frac{\alpha \cdot 2^\ell}{n} (\|\bar{h}\|_\infty + 1)^{2^\ell} \right)^{2^{-\ell}} \mathbf{1} \\ &\leq \left(\frac{\alpha}{n} \right)^{1-2^{-\ell}} \left(\left(\sqrt{\frac{\alpha}{n}} \|\bar{h}\|_\infty^{2^\ell} \right)^{2^{-\ell}} + \left(\frac{\alpha \cdot 2^\ell}{n} (\|\bar{h}\|_\infty + 1)^{2^\ell} \right)^{2^{-\ell}} \right) \mathbf{1} \\ &\leq \left(\frac{\alpha}{n} \right)^{1-2^{-\ell}} \left(\left(\sqrt{\frac{\alpha}{n}} \|\bar{h}\|_\infty^{2^\ell} \right)^{2^{-\ell}} + \left(\frac{\alpha \cdot 2^{2^\ell}}{n} (\|\bar{h}\|_\infty + 1)^{2^\ell} \right)^{2^{-\ell}} \right) \mathbf{1} \\ &= \left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} \|\bar{h}\|_\infty \mathbf{1} + 2 \frac{\alpha}{n} (\|\bar{h}\|_\infty + 1) \mathbf{1}. \end{aligned} \quad (60)$$

Since we have assumed that ρ^π is constant, using Lemma 23 and then the fact that $(\widehat{P}_\pi - P_\pi) \mathbf{1} = 0$, we have

$$\begin{aligned} \widehat{\rho}^\pi - \rho^\pi &= \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi) h^\pi \\ &= \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi) \left(h^\pi - \left(\min_s h^\pi(s) \right) \mathbf{1} \right) \\ &= \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi) \bar{h}. \end{aligned}$$

Now, using this equation, then Lemma 30 ℓ times (for $k = 0, \dots, \ell - 1$), then assuming $n \geq \alpha (\|\bar{h}\|_\infty + 1)$,

then using the bound (60) for the final term, we obtain

$$\begin{aligned}
\|\widehat{\rho}^\pi - \rho^\pi\|_\infty &= \left\| \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi) \bar{h} \right\|_\infty \\
&\leq \sum_{k=0}^{\ell-1} \left(2 \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n} \right)^{1-2^{-(k+1)}} + \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1) \right) \\
&\quad + \left(\frac{\alpha}{n} \right)^{1-2^{-\ell}} \left| \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^\ell} \right|^{2^{-\ell}} \\
&\leq \sum_{k=0}^{\ell-1} \left(2 \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n} \right)^{1-\frac{1}{2}} + \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1) \right) \\
&\quad + \left(\frac{\alpha}{n} \right)^{1-2^{-\ell}} \left| \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^\ell} \right|^{2^{-\ell}} \\
&= 2\ell \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n} \right)^{\frac{1}{2}} + \ell \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1) + \left(\frac{\alpha}{n} \right)^{1-2^{-\ell}} \left| \widehat{P}_\pi^\infty (\widehat{P}_\pi - P_\pi) (\bar{h})^{\circ 2^\ell} \right|^{2^{-\ell}} \\
&\leq 2\ell \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n} \right)^{\frac{1}{2}} + \ell \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1) + \left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} \|\bar{h}\|_\infty + 2\frac{\alpha}{n} (\|\bar{h}\|_\infty + 1) \\
&= 2\ell \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n} \right)^{\frac{1}{2}} + (\ell+1) \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1) + \left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} \|\bar{h}\|_\infty. \tag{61}
\end{aligned}$$

Note that the assumption $n \geq \alpha(\|\bar{h}\|_\infty + 1)$ was used to guarantee that the largest term in the initial summation was the $k = 0$ term.

Finally, we need to ensure that $\ell = \lceil \log_2 \log_2 (\|\bar{h}\|_\infty + 4) \rceil$ is sufficiently large so that the rightmost term in (61) is bounded by $2 \left(\frac{\alpha \|\bar{h}\|_\infty}{n} \right)^{\frac{1}{2}}$. This rightmost term can be bounded as

$$\left(\frac{\alpha}{n} \right)^{1-2^{-(\ell+1)}} \|\bar{h}\|_\infty = \left(\frac{\alpha \|\bar{h}\|_\infty}{n} \right)^{1-2^{-(\ell+1)}} \|\bar{h}\|_\infty^{2^{-(\ell+1)}} \leq \left(\frac{\alpha \|\bar{h}\|_\infty}{n} \right)^{1/2} \|\bar{h}\|_\infty^{2^{-(\ell+1)}}$$

(again using the assumption that $n \geq \alpha(\|\bar{h}\|_\infty + 1)$), and then we have the equivalences

$$\begin{aligned}
&\left(\frac{\alpha \|\bar{h}\|_\infty}{n} \right)^{1/2} \|\bar{h}\|_\infty^{2^{-(\ell+1)}} \leq 2 \left(\frac{\alpha \|\bar{h}\|_\infty}{n} \right)^{\frac{1}{2}} \\
&\iff \|\bar{h}\|_\infty^{2^{-(\ell+1)}} \leq 2 \\
&\iff 2^{-(\ell+1)} \log_2 (\|\bar{h}\|_\infty) \leq 1 \\
&\iff \log_2 (\|\bar{h}\|_\infty) \leq 2^{\ell+1} \\
&\iff \ell + 1 \geq \log_2 \log_2 (\|\bar{h}\|_\infty).
\end{aligned}$$

The final inequality is true for our definition of ℓ , so we have that the rightmost term in (61) is bounded by

$2 \left(\frac{\alpha \|\bar{h}\|_\infty}{n} \right)^{\frac{1}{2}}$ as desired. Thus combining this fact with (61), we have

$$\begin{aligned}
\|\widehat{\rho}^\pi - \rho^\pi\|_\infty &\leq 2\ell \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n} \right)^{\frac{1}{2}} + (\ell+1) \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1) + 2 \left(\frac{\alpha \|\bar{h}\|_\infty}{n} \right)^{\frac{1}{2}} \\
&\leq 2(\ell+1) \left(\frac{\alpha (\|\bar{h}\|_\infty + 1)}{n} \right)^{\frac{1}{2}} + (\ell+1) \frac{2\alpha}{n} (\|\bar{h}\|_\infty + 1)
\end{aligned}$$

as desired. $\square$

C.4 Bernstein-like inequalities

First we check the Bernstein-like inequality required for the proof of Theorem 2.

Lemma 32. *Fix a policy π , and let $\bar{h} = h^\pi - (\min_s h^\pi(s)) \mathbf{1}$. With probability at least $1 - \delta$, for all $k = 0, \dots, \lceil \log_2 \log_2 (\|\bar{h}\|_\infty + 4) \rceil$ we have*

$$\left| (\hat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{P_\pi} [(\bar{h})^{\circ 2^k}]}{n}} + \frac{\alpha}{n} \|\bar{h}\|_\infty^{2^k} \mathbf{1}$$

where $\alpha = 2 \log \left(\frac{3SA \log_2 \log_2 (\|h^\pi\|_{\text{span}} + 4)}{\delta} \right)$.

Proof. First, note that we only need to check this inequality for a fixed k , and then the desired result follows by taking a union bound and adjusting the failure probability. Fix k . Also fix $s \in \mathcal{S}$ and $a \in \mathcal{A}$. Using Bernstein's inequality (e.g. [Maurer and Pontil, 2009, Theorem 3]), we have that with probability at least $1 - 2\delta'$,

$$\begin{aligned} \left| (\hat{P}_{sa} - P_{sa})(\bar{h})^{\circ 2^k} \right| &\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{sa}} [(\bar{h})^{\circ 2^k}]}{n}} + \frac{\log \left(\frac{1}{\delta'} \right)}{3n} \left\| (\bar{h})^{\circ 2^k} \right\|_\infty \\ &\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{sa}} [(\bar{h})^{\circ 2^k}]}{n}} + \frac{\log \left(\frac{1}{\delta'} \right)}{3n} \|\bar{h}\|_\infty^{2^k} \\ &\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_{sa}} [(\bar{h})^{\circ 2^k}]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right)}{n} \|\bar{h}\|_\infty^{2^k}. \end{aligned}$$

Now taking a union bound over all possible $s \in \mathcal{S}$ and $a \in \mathcal{A}$, we have the elementwise inequality

$$\left| (\hat{P} - P)(\bar{h})^{\circ 2^k} \right| \leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_P [(\bar{h})^{\circ 2^k}]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right)}{n} \|\bar{h}\|_\infty^{2^k} \mathbf{1} \quad (62)$$

with probability at least $1 - 2SA\delta'$. We can use this to obtain that

$$\begin{aligned} \left| (\hat{P}_\pi - P_\pi)(\bar{h})^{\circ 2^k} \right| &= \left| M^\pi (\hat{P} - P)(\bar{h})^{\circ 2^k} \right| \\ &\leq M^\pi \left| (\hat{P} - P)(\bar{h})^{\circ 2^k} \right| \\ &\leq M^\pi \left(\sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_P [(\bar{h})^{\circ 2^k}]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right)}{n} \|\bar{h}\|_\infty^{2^k} \mathbf{1} \right) \\ &\leq \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) M^\pi \mathbb{V}_P [(\bar{h})^{\circ 2^k}]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right)}{n} \|\bar{h}\|_\infty^{2^k} \mathbf{1} \\ &= \sqrt{\frac{2 \log \left(\frac{1}{\delta'} \right) \mathbb{V}_{P_\pi} [(\bar{h})^{\circ 2^k}]}{n}} + \frac{2 \log \left(\frac{1}{\delta'} \right)}{n} \|\bar{h}\|_\infty^{2^k} \mathbf{1} \end{aligned}$$

where we used Jensen's inequality for the first inequality step (since each row of M^π is a probability distribution), then (62), then Jensen's inequality again since $\sqrt{\cdot}$ is concave and the fact that $M^\pi \mathbf{1} = \mathbf{1}$. Now taking a union bound over all values of k , of which there are at most

$$1 + \lceil \log_2 \log_2 (\|\bar{h}\|_\infty + 4) \rceil \leq 2 + \log_2 \log_2 (\|\bar{h}\|_\infty + 4) \leq 3 \log_2 \log_2 (\|\bar{h}\|_\infty + 4),$$

we can set $\delta' = \frac{\delta}{3SA \log_2 \log_2 (\|\bar{h}\|_\infty + 4)}$ and $\alpha = 2 \log \frac{1}{\delta'}$ to complete the proof. (Note $\|\bar{h}\|_\infty = \|h^\pi\|_{\text{span}}$.) $\square$

Now we set out to check the Bernstein-like inequalities required for the proof of Theorem 3. While this could be done by essentially copying the arguments of Lemma 20 but replacing $\hat{P}$ with $\tilde{P}$, we can instead reuse Lemma 20 more directly. The following Lemma 33 will help us do so.

Lemma 33. *Let π be an arbitrary policy, P be an arbitrary MDP transition matrix, let $\tilde{P} = (1 - \eta)P + \eta 1e_{s_0}^\top$ be an anchored version of P , and let $x \in \mathbb{R}^S$. Then*

$$\mathbb{V}_{P_\pi}[x] \leq \frac{1}{1 - \eta} \mathbb{V}_{\tilde{P}_\pi}[x].$$

Proof. Since the desired inequality is an elementwise inequality, it suffices to show for an arbitrary entry s . Thus let $p = (P_\pi)_s$ and $\tilde{p} = (\tilde{P}_\pi)_s$ be row vectors denoting the s th row of P_π and $\tilde{P}$, respectively. Note that $\tilde{p} = (1 - \eta)p + \eta e_{s_0}^\top$. Then we can calculate

$$\begin{aligned} \left(\mathbb{V}_{\tilde{P}_\pi}[x] \right)_s &= \tilde{p}(x^{\circ 2}) - (\tilde{p}x)^2 \\ &= (1 - \eta)p(x^{\circ 2}) + \eta(x(s_0))^2 - ((1 - \eta)px + \eta x(s_0))^2 \\ &= (1 - \eta)p(x^{\circ 2}) + \eta(x(s_0))^2 - (1 - \eta)^2(px)^2 - \eta^2(x(s_0))^2 - 2\eta(1 - \eta)(px)x(s_0) \\ &= (1 - \eta)(p(x^{\circ 2}) - (px)^2) + \eta(x(s_0))^2 + \eta(1 - \eta)(px)^2 - \eta^2(x(s_0))^2 - 2\eta(1 - \eta)(px)x(s_0) \\ &= (1 - \eta)(p(x^{\circ 2}) - (px)^2) + \eta(1 - \eta)((x(s_0))^2 + (px)^2 - 2(px)x(s_0)) \\ &\geq (1 - \eta)(p(x^{\circ 2}) - (px)^2) \\ &= (1 - \eta)(\mathbb{V}_{P_\pi}[x])_s \end{aligned}$$

where the inequality step is by the AM-GM inequality. $\square$

Using the above lemma, as well as the connection between discounted value functions and the bias functions in anchored MDPs, we are able to repurpose Lemma 34 to verify the Bernstein-like inequality conditions for the bias function of a near-optimal policy in an AMDP using anchoring.

Lemma 34. *If $n \geq 4$, then with probability at least $1 - \delta$, for all $\hat{\pi}$ which satisfy $\|\hat{V}_{1-\frac{1}{n}}^{\hat{\pi}} - \hat{V}_{1-\frac{1}{n}}^*\|_\infty \leq \frac{1}{n}$, letting $\bar{h} = \hat{\underline{h}}^{\hat{\pi}} - \left(\min_s \hat{\underline{h}}^{\hat{\pi}}(s) \right) \mathbf{1}$, for all $k = 0, \dots, \left\lceil \log_2 \log_2 \left(\|\hat{\underline{h}}^{\hat{\pi}}\|_{\text{span}} + 4 \right) \right\rceil$, we have*

$$\left| \left(\hat{\underline{P}}_{\hat{\pi}} - P_{\hat{\pi}} \right) (\bar{h})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha \mathbb{V}_{\hat{\underline{P}}_{\hat{\pi}}}[(\bar{h})^{\circ 2^k}]}{n}} + \frac{\alpha \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1}$$

where $\alpha = 18 \log \left(12 \frac{SAn^3}{\delta} \right)$.

Proof. First, we can directly use Lemma 20 to obtain that with probability at least $1 - \delta$, for all $\hat{\pi}$ which satisfy $\|\hat{V}_{1-\frac{1}{n}}^{\hat{\pi}} - \hat{V}_{1-\frac{1}{n}}^*\|_\infty \leq \frac{1}{n}$, we have

$$\left| \left(\hat{\underline{P}}_{\hat{\pi}} - P_{\hat{\pi}} \right) (\bar{V})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha' \mathbb{V}_{\hat{\underline{P}}_{\hat{\pi}}}[(\bar{V})^{\circ 2^k}]}{n}} + \frac{\alpha' \cdot 2^k}{n} (\|\bar{V}\|_\infty + 1)^{2^k} \mathbf{1} \quad (63)$$

for all $k = 0, \dots, \left\lceil \log_2 \log_2 \left(\|\hat{V}_{1-\frac{1}{n}}^{\hat{\pi}}\|_{\text{span}} + 4 \right) \right\rceil$, where $\bar{V} = \hat{V}_{1-\frac{1}{n}}^{\hat{\pi}} - \left(\min_s \hat{V}_{1-\frac{1}{n}}^{\hat{\pi}}(s) \right) \mathbf{1}$ and

$$\alpha' = 16 \log \left(12 \frac{SAn}{(1 - (1 - \frac{1}{n}))^2 \delta} \right) = 16 \log \left(12 \frac{SAn^3}{\delta} \right).$$

Now using Lemma 24, we have that $\hat{\underline{h}}^{\hat{\pi}} = \hat{V}_{1-\frac{1}{n}}^{\hat{\pi}} - c\mathbf{1}$ for some c , which immediately implies that $\bar{V} = \bar{h}$ and also that $\|\hat{V}_{1-\frac{1}{n}}^{\hat{\pi}}\|_{\text{span}} = \|\hat{\underline{h}}^{\hat{\pi}}\|_{\text{span}}$. Applying these facts to (63) we obtain that (under the same event)

$$\left| \left(\hat{\underline{P}}_{\hat{\pi}} - P_{\hat{\pi}} \right) (\bar{h})^{\circ 2^k} \right| \leq \sqrt{\frac{\alpha' \mathbb{V}_{\hat{\underline{P}}_{\hat{\pi}}}[(\bar{h})^{\circ 2^k}]}{n}} + \frac{\alpha' \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \quad (64)$$

for all $k = 0, \dots, \left\lceil \log_2 \log_2 \left(\|\hat{\underline{h}}^{\hat{\pi}}\|_{\text{span}} + 4 \right) \right\rceil$. Therefore it remains to replace $\hat{P}$ with $\hat{\underline{P}}$ within (64). Fix k . First notice that

$$\hat{\underline{P}}_{\pi} - P_{\pi} = \left(1 - \frac{1}{n}\right) \hat{P}_{\pi} + \frac{1}{n} \mathbf{1} e_{s_0}^{\top} - P_{\pi} = \left(1 - \frac{1}{n}\right) (\hat{P}_{\pi} - P_{\pi}) + \frac{1}{n} (\mathbf{1} e_{s_0}^{\top} - P_{\pi})$$

so

$$\begin{aligned} \left| (\hat{\underline{P}}_{\hat{\pi}} - P_{\hat{\pi}}) (\bar{h})^{\circ 2^k} \right| &\leq \left(1 - \frac{1}{n}\right) \left| (\hat{P}_{\hat{\pi}} - P_{\hat{\pi}}) (\bar{h})^{\circ 2^k} \right| + \frac{1}{n} \|\mathbf{1} e_{s_0}^{\top} - P_{\pi}\|_{\infty \rightarrow \infty} \left\| (\bar{h})^{\circ 2^k} \right\|_{\infty} \mathbf{1} \\ &\leq \left(1 - \frac{1}{n}\right) \left| (\hat{P}_{\hat{\pi}} - P_{\hat{\pi}}) (\bar{h})^{\circ 2^k} \right| + \frac{2}{n} \|\bar{h}\|_{\infty}^{2^k} \mathbf{1}. \end{aligned} \quad (65)$$

Also, we can use Lemma 33 to obtain that

$$\mathbb{V}_{\hat{P}_{\hat{\pi}}} \left[(\bar{V})^{\circ 2^k} \right] \leq \frac{1}{1 - \frac{1}{n}} \mathbb{V}_{\hat{\underline{P}}_{\hat{\pi}}} \left[(\bar{V})^{\circ 2^k} \right]. \quad (66)$$

Now combining inequalities (64), (65), and (66), we obtain

$$\begin{aligned} \left| (\hat{\underline{P}}_{\hat{\pi}} - P_{\hat{\pi}}) (\bar{h})^{\circ 2^k} \right| &\leq \left(1 - \frac{1}{n}\right) \left| (\hat{P}_{\hat{\pi}} - P_{\hat{\pi}}) (\bar{h})^{\circ 2^k} \right| + \frac{2}{n} \|\bar{h}\|_{\infty}^{2^k} \mathbf{1} \\ &\leq \left(1 - \frac{1}{n}\right) \sqrt{\frac{\alpha' \mathbb{V}_{\hat{P}_{\hat{\pi}}} \left[(\bar{h})^{\circ 2^k} \right]}{n}} + \frac{\alpha' \cdot 2^k}{n} (\|\bar{h}\|_{\infty} + 1)^{2^k} \mathbf{1} + \frac{2}{n} \|\bar{h}\|_{\infty}^{2^k} \mathbf{1} \\ &\leq \left(1 - \frac{1}{n}\right) \sqrt{\frac{\alpha' \frac{1}{1 - \frac{1}{n}} \mathbb{V}_{\hat{\underline{P}}_{\hat{\pi}}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha' \cdot 2^k}{n} (\|\bar{h}\|_{\infty} + 1)^{2^k} \mathbf{1} + \frac{2}{n} \|\bar{h}\|_{\infty}^{2^k} \mathbf{1} \\ &\leq \sqrt{1 - \frac{1}{n}} \sqrt{\frac{\alpha' \mathbb{V}_{\hat{\underline{P}}_{\hat{\pi}}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{\alpha' \cdot 2^k}{n} (\|\bar{h}\|_{\infty} + 1)^{2^k} \mathbf{1} + \frac{2}{n} (\|\bar{h}\|_{\infty} + 1)^{2^k} \mathbf{1} \\ &\leq \sqrt{\frac{\alpha' \mathbb{V}_{\hat{\underline{P}}_{\hat{\pi}}} \left[(\bar{V})^{\circ 2^k} \right]}{n}} + \frac{(\alpha' + 2) \cdot 2^k}{n} (\|\bar{h}\|_{\infty} + 1)^{2^k} \mathbf{1}. \end{aligned}$$

Therefore we can set $\alpha = 18 \log \left(12 \frac{SAn}{(1-\gamma)^2} \right) \geq \alpha' + 2$ and obtain the desired conclusion. $\square$

Now we check the Bernstein-like inequalities used within the proof of Theorem 1.

Lemma 35. *Suppose $\hat{\pi}^*$ is a bias-optimal policy in the AMDP $(\hat{P}, r)$. If $n \geq 4$, then with probability at least $1 - \delta$, letting $\bar{h} = \hat{h}^* - \left(\min_s \hat{h}^*(s) \right) \mathbf{1}$, for all $k = 0, \dots, \left\lceil \log_2 \log_2 \left(\|\hat{h}^*\|_{\text{span}} + 4 \right) \right\rceil$, we have*

$$\left| (\hat{P}_{\hat{\pi}^*} - P_{\hat{\pi}^*}) (\bar{h})^{\circ 2^k} \right| \leq \sqrt{\frac{16 \log \left(\frac{18 \cdot 8 SAn}{\delta (1-\hat{\gamma}^*)^3} \right) \mathbb{V}_{\hat{P}_{\hat{\pi}^*}} \left[(\bar{h})^{\circ 2^k} \right]}{n}} + \frac{16 \log \left(\frac{18 \cdot 8 SAn}{\delta (1-\hat{\gamma}^*)^3} \right) \cdot 2^k}{n} (\|\bar{h}\|_{\infty} + 1)^{2^k} \mathbf{1}.$$

Proof. To handle the fact that $\hat{\gamma}^*$ is random, we will prove a version of the inequality for each γ such that $\frac{1}{1-\gamma} = 2^m$ for some integer $m \geq 0$, and then we will adjust the failure probability for each m so that the overall failure probability is bounded by δ . First, we fix γ , and we seek to show that with probability at least $1 - \delta'$, for all vectors $x \in \mathbb{R}^S$ such that $\left\| x - \hat{V}_{\gamma}^* \right\|_{\infty} \leq \frac{1}{n}$, letting $\bar{x} = x - (\min_s x(s)) \mathbf{1}$, we have that

$$\left| (\hat{P}_{\hat{\pi}^*} - P_{\hat{\pi}^*}) (\bar{x})^{\circ 2^k} \right| \leq \sqrt{\frac{16 \log \left(\frac{18 SAn}{(1-\gamma)^{2\delta'}} \right) \mathbb{V}_{\hat{P}_{\hat{\pi}^*}} \left[(\bar{x})^{\circ 2^k} \right]}{n}} + \frac{16 \log \left(\frac{18 SAn}{(1-\gamma)^{2\delta'}} \right) \cdot 2^k}{n} (\|\bar{x}\|_{\infty} + 1)^{2^k} \mathbf{1} \quad (67)$$

for all $k = 0, \dots, \lceil \log_2 \log_2 (\|x\|_{\text{span}} + 4) \rceil$. We will argue that (67) follows from an identical argument to Lemma 20. Specifically, we will argue that we can replace $\bar{V} = \hat{V}^{\hat{\pi}} - (\min_s \hat{V}^{\hat{\pi}}(s)) \mathbf{1}$ within the proof of Lemma 20 by $\bar{x}$. We observe that the proof of Lemma 20 only uses the following properties of the vector $\hat{V}^{\hat{\pi}}$: that $\|\hat{V}^{\hat{\pi}}\|_{\text{span}} \leq \frac{1}{1-\gamma}$, and that $\|\hat{V}^{\hat{\pi}} - \hat{V}^{\star}\|_{\infty} \leq \frac{1}{n}$. Furthermore, the bound $\|\hat{V}^{\hat{\pi}}\|_{\text{span}} \leq \frac{1}{1-\gamma}$ is only used to coarsely upper-bound the number of values of k for which the desired inequality must be checked. We can instead use the fact that $\|x - \hat{V}_{\gamma}^{\star}\|_{\infty} \leq \frac{1}{n}$ to obtain that

$$\|x\|_{\text{span}} \leq \|\hat{V}_{\gamma}^{\star}\|_{\text{span}} + \|\hat{V}_{\gamma}^{\star} - x\|_{\text{span}} \leq \|\hat{V}_{\gamma}^{\star}\|_{\text{span}} + 2 \|\hat{V}_{\gamma}^{\star} - x\|_{\infty} \leq \frac{1}{1-\gamma} + \frac{2}{n} \leq \frac{1}{1-\gamma} + 1 \leq \frac{2}{1-\gamma}.$$

Thus repeating an argument similar to (18), using that $\log_2 \log_2 (x+4) \leq 2x$ for $x \geq 1$, we can bound

$$\begin{aligned} \lceil \log_2 \log_2 (\|x\|_{\text{span}} + 4) \rceil &\leq \left\lceil \log_2 \log_2 \left(\frac{2}{1-\gamma} + 4 \right) \right\rceil \\ &\leq 1 + \log_2 \log_2 \left(\frac{2}{1-\gamma} + 4 \right) \\ &\leq 1 + 4 \frac{1}{1-\gamma} \\ &\leq 5 \frac{1}{1-\gamma} \end{aligned}$$

and thus (counting $k = 0$) there are $\leq 1 + 5 \frac{1}{1-\gamma} \leq \frac{6}{1-\gamma}$ values of k to check the inequality for if we check it for all values up to the upper bound $5 \frac{1}{1-\gamma}$. Comparing with the bound (27), this will cause us to obtain a factor of

$$16 \log \left(\frac{3SA|U|}{\delta'} 6 \frac{1}{1-\gamma} \right) = 16 \log \left(\frac{18SAn}{(1-\gamma)^2 \delta'} \right)$$

(rather than the $16 \log \left(\frac{12SAn}{(1-\gamma)^2 \delta'} \right)$ which appears in Lemma 20). The rest of the proof of Lemma 20 only uses the fact that $\|\hat{V}^{\hat{\pi}} - \hat{V}^{\star}\|_{\infty} \leq \frac{1}{n}$, and thus goes through unchanged if we replace $\hat{V}^{\hat{\pi}}$ with x (and thus also $\bar{V}$ with $\bar{x}$), which, following the proof up to the bound (25), yields

$$\left| (\hat{P} - P) (\bar{x})^{\circ 2^k} \right| \leq \sqrt{\frac{16 \log \left(\frac{18SAn}{(1-\gamma)^2 \delta'} \right) \mathbb{V}_{\hat{P}} \left[(\bar{x})^{\circ 2^k} \right]}{n}} + \frac{16 \log \left(\frac{18SAn}{(1-\gamma)^2 \delta'} \right) \cdot 2^k}{n} (\|\bar{x}\|_{\infty} + 1)^{2^k} \mathbf{1}$$

for all $k = 0, \dots, \lceil \log_2 \log_2 (\|x\|_{\text{span}} + 4) \rceil$. From here, we can use identical steps as to the end of the proof of the inequality (26) within Lemma 20 (but with $\hat{\pi}^{\star}$ rather than $\hat{\pi}$) to conclude (67) as desired.

Now applying (67) with $\frac{1}{1-\gamma} = 2^m$ and $\delta' = \frac{\delta}{2^m}$ for each $m = 1, 2, \dots$, and taking a union bound over all m , we obtain that with probability at least $1 - \sum_{m=1}^{\infty} \frac{\delta}{2^m} = 1 - \delta$, we have that for all integers $m \geq 1$, for all x such that $\|x - \hat{V}_{1-2^{-m}}^{\star}\|_{\infty}$, that

$$\begin{aligned} \left| (\hat{P}_{\hat{\pi}^{\star}} - P_{\hat{\pi}^{\star}}) (\bar{x})^{\circ 2^k} \right| &\leq \sqrt{\frac{16 \log \left(\frac{18SAn}{(2^{-m})^2 \frac{\delta}{2^m}} \right) \mathbb{V}_{\hat{P}_{\hat{\pi}^{\star}}} \left[(\bar{x})^{\circ 2^k} \right]}{n}} + \frac{16 \log \left(\frac{18SAn}{(2^{-m})^2 \frac{\delta}{2^m}} \right) \cdot 2^k}{n} (\|\bar{x}\|_{\infty} + 1)^{2^k} \mathbf{1} \\ &= \sqrt{\frac{16 \log \left(\frac{18SAn 2^{3m}}{\delta} \right) \mathbb{V}_{\hat{P}_{\hat{\pi}^{\star}}} \left[(\bar{x})^{\circ 2^k} \right]}{n}} + \frac{16 \log \left(\frac{18SAn 2^{3m}}{\delta} \right) \cdot 2^k}{n} (\|\bar{x}\|_{\infty} + 1)^{2^k} \mathbf{1}. \quad (68) \end{aligned}$$

On this event, recalling we have defined $\hat{\gamma}^{\star}$ as the smallest discount factor such that for all $\gamma \geq \hat{\gamma}^{\star}$, there exists $c \in \mathbb{R}$ such that

$$\|\hat{V}_{\gamma}^{\star} - \hat{h}^{\star} - c\mathbf{1}\|_{\infty} \leq \frac{1}{n}, \quad (69)$$

we now define $\hat{\gamma}^\sharp$ as the smallest γ such that $\gamma \geq \hat{\gamma}^*$ and also there exists an integer $m \geq 1$ such that $\frac{1}{1-\hat{\gamma}^\sharp} = 2^m$. Thus we have $\frac{1}{1-\hat{\gamma}^*} \leq \frac{1}{1-\hat{\gamma}^\sharp} \leq \frac{2}{1-\hat{\gamma}^*}$. Also since $\hat{\gamma}^\sharp \geq \hat{\gamma}^*$, by (69) we have that there exists some (random) scalar c such that

$$\left\| \hat{V}_{\hat{\gamma}^\sharp}^* - \hat{h}^* - c\mathbf{1} \right\|_\infty \leq \frac{1}{n}. \quad (70)$$

By (70), we may apply (68) to $x = \hat{h}^* + c\mathbf{1}$. Also note that for this choice of x , $\bar{x} = x - (\min_s x(s))\mathbf{1} = \hat{h}^* + c\mathbf{1} - (\min_s \hat{h}^*(s) + c)\mathbf{1} = \hat{h}^* - (\min_s \hat{h}^*(s))\mathbf{1} = \bar{h}$. Also $\|x\|_{\text{span}} = \|\hat{h}^*\|_{\text{span}}$. Thus, plugging these observations into (68), we obtain (still on the aforementioned event) that

$$\begin{aligned} \left| \left(\hat{P}_{\hat{\pi}^*} - P_{\hat{\pi}^*} \right) (\bar{h})^{\circ 2^k} \right| &= \sqrt{\frac{16 \log \left(\frac{18SA n}{\delta(1-\hat{\gamma}^\sharp)^3} \right) \mathbb{V}_{\hat{P}_{\hat{\pi}^*}} \left[(\bar{h})^{\circ 2^k} \right]}{n}} + \frac{16 \log \left(\frac{18SA n}{\delta(1-\hat{\gamma}^\sharp)^3} \right) \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1} \\ &\leq \sqrt{\frac{16 \log \left(\frac{18 \cdot 8SA n}{\delta(1-\hat{\gamma}^*)^3} \right) \mathbb{V}_{\hat{P}_{\hat{\pi}^*}} \left[(\bar{h})^{\circ 2^k} \right]}{n}} + \frac{16 \log \left(\frac{18 \cdot 8SA n}{\delta(1-\hat{\gamma}^*)^3} \right) \cdot 2^k}{n} (\|\bar{h}\|_\infty + 1)^{2^k} \mathbf{1}. \end{aligned}$$

□

C.5 Proof of Theorem 2

Proof of Theorem 2. Combining Lemma 32 with Lemma 31, we obtain that with probability at least $1 - \delta$,

$$\|\hat{\rho}^\pi - \rho^\pi\|_\infty \leq 2(\ell + 1) \left(\frac{\alpha (\|h^\pi\|_{\text{span}} + 1)}{n} \right)^{\frac{1}{2}} + (\ell + 1) \frac{2\alpha}{n} (\|h^\pi\|_{\text{span}} + 1) \quad (71)$$

where $\alpha = 2 \log \left(\frac{3SA \log_2 \log_2 (\|h^\pi\|_{\text{span}} + 4)}{\delta} \right)$ and $\ell = \lceil \log_2 \log_2 (\|h^\pi\|_{\text{span}} + 4) \rceil$. We can additionally assume without loss of generality that $n \geq \alpha (\|h^\pi\|_{\text{span}} + 1)$, since otherwise the desired theorem conclusion still holds trivially since we always have $\|\hat{\rho}^\pi - \rho^\pi\|_\infty \leq 1$. Then since $\frac{\alpha (\|h^\pi\|_{\text{span}} + 1)}{n} \leq 1$, we have $\left(\frac{\alpha (\|h^\pi\|_{\text{span}} + 1)}{n} \right)^{\frac{1}{2}} \geq \frac{\alpha (\|h^\pi\|_{\text{span}} + 1)}{n}$, and then we can use this to simplify the bound (71) to obtain

$$\begin{aligned} \|\hat{\rho}^\pi - \rho^\pi\|_\infty &\leq 2(\ell + 1) \left(\frac{\alpha (\|h^\pi\|_{\text{span}} + 1)}{n} \right)^{\frac{1}{2}} + (\ell + 1) \frac{2\alpha}{n} (\|h^\pi\|_{\text{span}} + 1) \\ &\leq 2(\ell + 1) \left(\frac{\alpha (\|h^\pi\|_{\text{span}} + 1)}{n} \right)^{\frac{1}{2}} + 2(\ell + 1) \left(\frac{\alpha (\|h^\pi\|_{\text{span}} + 1)}{n} \right)^{\frac{1}{2}} \\ &= 4(\lceil \log_2 \log_2 (\|h^\pi\|_{\text{span}} + 4) \rceil + 1) \sqrt{\frac{2 \log \left(\frac{3SA \log_2 \log_2 (\|h^\pi\|_{\text{span}} + 4)}{\delta} \right)}{n}} (\|h^\pi\|_{\text{span}} + 1) \\ &\leq 12 (\log_2 \log_2 (\|h^\pi\|_{\text{span}} + 4)) \sqrt{\frac{2 \log \left(\frac{3SA \log_2 \log_2 (\|h^\pi\|_{\text{span}} + 4)}{\delta} \right)}{n}} (\|h^\pi\|_{\text{span}} + 1) \\ &\leq \sqrt{\frac{C_4 \log^3 \left(\frac{SA n}{\delta} \right)}{n}} (\|h^\pi\|_{\text{span}} + 1) \end{aligned}$$

where in the final inequality we use the upper-bound $\|h^\pi\|_{\text{span}} \leq n$ (which follows from $\frac{\alpha (\|h^\pi\|_{\text{span}} + 1)}{n} \leq 1$) and choose a sufficiently large constant C_4 . □

C.6 Proof of Theorem 3

First, we show that the result follows from bounding certain “policy evaluation error” terms.

Lemma 36. *Under the conditions of Theorem 3,*

$$\rho^{\hat{\pi}} \geq \rho^{\pi^*} - \left(\left\| \hat{\rho}^{\pi^*_{1-\frac{1}{n}}} - \rho^{\pi^*_{1-\frac{1}{n}}} \right\|_{\infty} + \left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty} + \frac{\|h^*\|_{\text{span}} + 1/n}{n} \right) \mathbf{1}. \quad (72)$$

Proof. Note that by Lemma 27 and the conditions on the SolveAMDP procedure used in the statement of Theorem 3, we have that $\hat{\rho}^{\hat{\pi}} \geq \hat{\rho}^* - \frac{1}{n^2} \mathbf{1}$. Also we recall that $\pi^*_{1-\frac{1}{n}}$ is defined as the optimal policy for the DMDP $(P, r, 1 - \frac{1}{n})$, and by Lemma 24 this policy has optimal gain in the anchored AMDP with transition matrix $\underline{P} = (1 - \frac{1}{n})P + \frac{1}{n} \mathbf{1} e_{s_0}^\top$. For notational convenience we let $\gamma = 1 - \frac{1}{n}$ so that we can abbreviate $\pi^*_\gamma = \pi^*_{1-\frac{1}{n}}$. Then we can calculate that

$$\begin{aligned} \rho^{\hat{\pi}} &\geq \hat{\rho}^{\hat{\pi}} - \left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty} \mathbf{1} && \text{triangle inequality} \\ &\geq \hat{\rho}^* - \frac{1}{n^2} \mathbf{1} - \left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty} \mathbf{1} && \hat{\rho}^{\hat{\pi}} \geq \hat{\rho}^* - \frac{1}{n^2} \mathbf{1} \\ &\geq \hat{\rho}^{\pi^*_\gamma} - \frac{1}{n^2} \mathbf{1} - \left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty} \mathbf{1} && \hat{\rho}^* \geq \hat{\rho}^{\pi^*_\gamma} \\ &\geq \underline{\rho}^{\pi^*_\gamma} - \left\| \hat{\rho}^{\pi^*_\gamma} - \underline{\rho}^{\pi^*_\gamma} \right\|_{\infty} \mathbf{1} - \frac{1}{n^2} \mathbf{1} - \left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty} \mathbf{1} && \text{triangle inequality} \\ &\geq \underline{\rho}^{\pi^*} - \left\| \hat{\rho}^{\pi^*_\gamma} - \underline{\rho}^{\pi^*_\gamma} \right\|_{\infty} \mathbf{1} - \frac{1}{n^2} \mathbf{1} - \left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty} \mathbf{1} && \underline{\rho}^{\pi^*_{1-\frac{1}{n}}} = \underline{\rho}^* \geq \underline{\rho}^{\pi^*} \\ &\geq \rho^{\pi^*} - \left\| \underline{\rho}^{\pi^*} - \rho^{\pi^*} \right\|_{\infty} \mathbf{1} - \left\| \hat{\rho}^{\pi^*_\gamma} - \underline{\rho}^{\pi^*_\gamma} \right\|_{\infty} \mathbf{1} - \frac{1}{n^2} \mathbf{1} - \left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty} \mathbf{1} && \text{triangle inequality} \\ &\geq \rho^{\pi^*} - \frac{\|h^*\|_{\text{span}}}{n} \mathbf{1} - \left\| \hat{\rho}^{\pi^*_\gamma} - \underline{\rho}^{\pi^*_\gamma} \right\|_{\infty} \mathbf{1} - \frac{1}{n^2} \mathbf{1} - \left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty} \mathbf{1}. && \text{Lemma 24, } \eta = \frac{1}{n} \end{aligned}$$

□

We remark that with very similar arguments we could replace the term $\left\| \hat{\rho}^{\pi^*_{1-\frac{1}{n}}} - \rho^{\pi^*_{1-\frac{1}{n}}} \right\|_{\infty}$ with the term $\left\| \hat{\rho}^{\pi^*} - \rho^{\pi^*} \right\|_{\infty}$ or the term $\left\| \hat{\rho}^{\pi^*} - \rho^{\pi^*} \right\|_{\infty}$ and it would still be possible to carry out the arguments, however as will be seen shortly, the term $\left\| \hat{\rho}^{\pi^*_{1-\frac{1}{n}}} - \rho^{\pi^*_{1-\frac{1}{n}}} \right\|_{\infty}$ enables us to reuse bounds from our DMDP results.

Now we complete the proof of the theorem.

Proof of Theorem 3. By Lemma 36, it suffices to bound the terms $\left\| \hat{\rho}^{\pi^*_{1-\frac{1}{n}}} - \rho^{\pi^*_{1-\frac{1}{n}}} \right\|_{\infty}$ and $\left\| \hat{\rho}^{\hat{\pi}} - \rho^{\hat{\pi}} \right\|_{\infty}$ with high probability.

First we handle the easier term $\left\| \hat{\rho}^{\pi^*_{1-\frac{1}{n}}} - \rho^{\pi^*_{1-\frac{1}{n}}} \right\|_{\infty}$. By similar observations as those used in the proof of Theorem 6, this term can be directly related to a difference of discounted value functions using Lemma

24. Lemma 24 shows that $\hat{\rho}^{\pi^*_{1-\frac{1}{n}}} = \frac{\hat{V}_{1-\frac{1}{n}}^{\pi^*_{1-\frac{1}{n}}}(s_0)}{n}$ and that $\rho^{\pi^*_{1-\frac{1}{n}}} = \frac{V_{1-\frac{1}{n}}^{\pi^*_{1-\frac{1}{n}}}(s_0)}{n}$, which implies that

$$\begin{aligned} \left\| \hat{\rho}^{\pi^*_{1-\frac{1}{n}}} - \rho^{\pi^*_{1-\frac{1}{n}}} \right\|_{\infty} &= \left\| \frac{\hat{V}_{1-\frac{1}{n}}^{\pi^*_{1-\frac{1}{n}}}(s_0)}{n} - \frac{V_{1-\frac{1}{n}}^{\pi^*_{1-\frac{1}{n}}}(s_0)}{n} \right\|_{\infty} \\ &\leq \frac{1}{n} \left\| \hat{V}_{1-\frac{1}{n}}^{\pi^*_{1-\frac{1}{n}}} - V_{1-\frac{1}{n}}^{\pi^*_{1-\frac{1}{n}}} \right\|_{\infty}. \end{aligned} \quad (73)$$

Now we can reuse part of the proof of Theorem 10 which bounds $\left\| \hat{V}_{1-\frac{1}{n}}^{\pi^*_{1-\frac{1}{n}}} - V_{1-\frac{1}{n}}^{\pi^*_{1-\frac{1}{n}}} \right\|_{\infty}$. Specifically, setting $\gamma = 1 - \frac{1}{n}$, then it is shown in inequality (34) from the proof of Theorem 10 that with probability at least

$1 - \delta$, we have

$$\begin{aligned} \left\| \widehat{V}_\gamma^{\pi_\gamma^*} - V_\gamma^{\pi_\gamma^*} \right\|_\infty &\leq \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha_1 \left(\|V_\gamma^{\pi_\gamma^*}\|_{\text{span}} + 1 \right)}{n}} \\ &\leq n 24 (\log_2 \log_2 (n+4)) \sqrt{\frac{2\alpha_1 (\|h^*\|_{\text{span}} + 1)}{n}} \end{aligned} \quad (74)$$

where $\alpha_1 = 2 \log \left(\frac{6S \log_2 \log_2 (\|V_\gamma^{\pi_\gamma^*}\|_{\text{span}} + 4)}{\delta} \right)$, and in the second inequality we used that $\|V_\gamma^{\pi_\gamma^*}\|_{\text{span}} = \|h^*\|_{\text{span}} \leq 2\|h^*\|_{\text{span}}$, both steps of which follow from Lemma 24 (the inequality step because ρ^* is constant). Combining (74) with (73), we obtain

$$\left\| \widehat{\rho}_{1-\frac{1}{n}}^{\pi_\gamma^*} - \rho_{1-\frac{1}{n}}^{\pi_\gamma^*} \right\|_\infty \leq 24 (\log_2 \log_2 (n+4)) \sqrt{\frac{2\alpha_1 (\|h^*\|_{\text{span}} + 1)}{n}}. \quad (75)$$

Now we bound the term $\left\| \widehat{\rho}^{\widehat{\pi}} - \rho^{\widehat{\pi}} \right\|_\infty$. First note that by Lemma 27, the requirement (2) implies that $\left\| \widehat{V}_{1-\frac{1}{n}}^{\widehat{\pi}} - \widehat{V}_{1-\frac{1}{n}}^{\pi^*} \right\|_\infty \leq \frac{1}{n}$. Thus, if we assume for now that $n \geq 4$, the conditions of Lemma 34 are satisfied, and thus by combining it with Lemma 31, we have that with (additional) failure probability at most δ ,

$$\left\| \widehat{\rho}^{\widehat{\pi}} - \rho^{\widehat{\pi}} \right\|_\infty \leq 2(\ell+1) \left(\frac{\alpha_2 \left(\|\widehat{h}^{\widehat{\pi}}\|_{\text{span}} + 1 \right)}{n} \right)^{\frac{1}{2}} + (\ell+1) \frac{2\alpha_2}{n} \left(\|\widehat{h}^{\widehat{\pi}}\|_{\text{span}} + 1 \right) \quad (76)$$

where $\alpha_2 = 18 \log \left(12 \frac{SA n^3}{\delta} \right)$ and $\ell = \left\lceil \log_2 \log_2 \left(\|\widehat{h}^{\widehat{\pi}}\|_{\text{span}} + 4 \right) \right\rceil$. Also note that the use of Lemma 31 requires that $\widehat{\rho}^{\widehat{\pi}}$ is constant, which follows from Lemma 24. By following arguments which are analogous to the bounds in the proof of Theorem 2, we can simplify (76) and obtain

$$\left\| \widehat{\rho}^{\widehat{\pi}} - \rho^{\widehat{\pi}} \right\|_\infty \leq 12 \left(\log_2 \log_2 \left(\|\widehat{h}^{\widehat{\pi}}\|_{\text{span}} + 4 \right) \right) \sqrt{\frac{\alpha_2}{n} \left(\|\widehat{h}^{\widehat{\pi}}\|_{\text{span}} + 1 \right)}. \quad (77)$$

We can also assume without loss of generality that $\|\widehat{h}^{\widehat{\pi}}\|_{\text{span}} \leq n$, since otherwise the RHS of (77) is greater than 1 and so the inequality (77) still holds since trivially always $\left\| \widehat{\rho}^{\widehat{\pi}} - \rho^{\widehat{\pi}} \right\|_\infty \leq 1$. Thus we can bound

$$\log_2 \log_2 \left(\|\widehat{h}^{\widehat{\pi}}\|_{\text{span}} + 4 \right) \leq \log_2 \log_2 (n+4) \leq \log 4n \leq \alpha_2$$

since $\log_2 \log_2 (x+4) \leq \log 4x$ for $x \geq 1$. Also since $\|V_\gamma^{\pi_\gamma^*}\|_{\text{span}} \leq \frac{1}{1-\gamma} = n$ and also $\log_2 \log_2 (x+4) \leq 2x$ for $x \geq 1$,

$$2\alpha_1 = 4 \log \left(\frac{6S \log_2 \log_2 \left(\|V_\gamma^{\pi_\gamma^*}\|_{\text{span}} + 4 \right)}{\delta} \right) \leq 4 \log \left(\frac{12Sn}{\delta} \right) \leq \alpha_2.$$

Using these bounds and combining inequalities (75) and (77) with Lemma 36, we obtain that with probability at least $1 - 2\delta$,

$$\rho^{\widehat{\pi}} \geq \rho^{\pi^*} - \left(24\alpha_2 \sqrt{\frac{\alpha_2}{n} (\|h^*\|_{\text{span}} + 1)} + 12\alpha_2 \sqrt{\frac{\alpha_2}{n} \left(\|\widehat{h}^{\widehat{\pi}}\|_{\text{span}} + 1 \right)} + \frac{\|h^*\|_{\text{span}} + 1/n}{n} \right) \mathbf{1}. \quad (78)$$

We can also bound

$$\frac{\|h^*\|_{\text{span}} + 1/n}{n} \leq \frac{\|h^*\|_{\text{span}} + 1}{n} \leq \sqrt{\frac{\|h^*\|_{\text{span}} + 1}{n}}$$

(where similarly to before we are assuming without loss of generality that $\|h^*\|_{\text{span}} + 1 \leq n$ for the last inequality). Combining this fact with (78), using (2) to bound $\|\hat{h}^{\hat{\pi}}\|_{\text{span}} \leq \|\hat{h}^*\|_{\text{span}} + \|\hat{h}^{\hat{\pi}} - \hat{h}^*\|_{\text{span}} \leq \|\hat{h}^*\|_{\text{span}} + 2\|\hat{h}^{\hat{\pi}} - \hat{h}^*\|_{\infty} \leq \|\hat{h}^*\|_{\text{span}} + \frac{2}{3n^2} \leq \|\hat{h}^*\|_{\text{span}} + 1$, and using the fact that $\sqrt{a} + \sqrt{b} \leq 2\sqrt{a+b}$, we conclude

$$\begin{aligned}
\rho^{\hat{\pi}} &\geq \rho^{\pi^*} - \left(24\alpha_2 \sqrt{\frac{\alpha_2}{n} (\|h^*\|_{\text{span}} + 1)} + 12\alpha_2 \sqrt{\frac{\alpha_2}{n} (\|\hat{h}^{\hat{\pi}}\|_{\text{span}} + 1)} + \sqrt{\frac{\|h^*\|_{\text{span}} + 1}{n}} \right) \mathbf{1} \\
&\geq \rho^{\pi^*} - \left(25\alpha_2 \sqrt{\frac{\alpha_2}{n} (\|h^*\|_{\text{span}} + 1)} + 12\alpha_2 \sqrt{\frac{\alpha_2}{n} (\|\hat{h}^{\hat{\pi}}\|_{\text{span}} + 1)} \right) \mathbf{1} \\
&\geq \rho^{\pi^*} - \left(25\alpha_2 \sqrt{\frac{\alpha_2}{n} (\|h^*\|_{\text{span}} + 1)} + 12\alpha_2 \sqrt{\frac{\alpha_2}{n} (\|\hat{h}^*\|_{\text{span}} + 2)} \right) \mathbf{1} \\
&\geq \rho^{\pi^*} - 50\alpha_2 \sqrt{\frac{\alpha_2}{n} (\|h^*\|_{\text{span}} + \|\hat{h}^*\|_{\text{span}} + 3)} \mathbf{1} \\
&\geq \rho^{\pi^*} - C_5 \sqrt{\frac{\log(\frac{SA_n}{\delta})}{n}} (\|h^*\|_{\text{span}} + \|\hat{h}^*\|_{\text{span}} + 1) \mathbf{1},
\end{aligned}$$

also adjusting δ make the total failure probability δ rather than 2δ . Also note that we assumed $n \geq 4$ to derive this bound, but the inequality is also trivially true if $n \leq 3$ (since always $\rho^{\hat{\pi}} \geq \rho^{\pi^*} - 1$), so this assumption can be removed without changing the result. $\square$

C.7 Proof of Theorem 8

First we provide a helper lemma to show that optimal policies are still near-optimal in perturbed DMDPs.

Lemma 37. *Let P be any transition matrix and fix a discount factor γ . Let $r_1, r_2 \in \mathbb{R}^{SA}$ be two reward vectors. Let V_{γ, r_1}^{π} denote the value function of policy π in the DMDP (P, r_1, γ) , and likewise let V_{γ, r_2}^{π} denote the value function of policy π in the DMDP (P, r_2, γ) . Also let π_1^* denote the optimal policy in (P, r_1, γ) and let π_2^* denote the optimal policy in (P, r_2, γ) . Let $V_{\gamma, r_1}^* = V_{\gamma, r_1}^{\pi_1^*}$ and $V_{\gamma, r_2}^* = V_{\gamma, r_2}^{\pi_2^*}$ denote the respective optimal value functions. Then*

$$V_{\gamma, r_2}^{\pi_1^*} \geq V_{\gamma, r_2}^* - 2 \frac{\|r_1 - r_2\|_{\infty}}{1 - \gamma} \mathbf{1}.$$

Proof. Using the definitions for value functions as well as the facts that $\|(I - \gamma P_{\pi})^{-1}\|_{\infty \rightarrow \infty} = \frac{1}{1 - \gamma}$ and

$\|M^\pi\|_{\infty \rightarrow \infty} = 1$ for any policy π , we can calculate

$$\begin{aligned}
V_{\gamma, r_2}^{\pi_1^*} &= (I - \gamma P_{\pi_1^*})^{-1} M^{\pi_1^*} r_2 \\
&= (I - \gamma P_{\pi_1^*})^{-1} M^{\pi_1^*} r_1 + (I - \gamma P_{\pi_1^*})^{-1} M^{\pi_1^*} (r_2 - r_1) \\
&\geq (I - \gamma P_{\pi_1^*})^{-1} M^{\pi_1^*} r_1 - \|(I - \gamma P_{\pi_1^*})^{-1}\|_{\infty \rightarrow \infty} \|M^{\pi_1^*}\|_{\infty \rightarrow \infty} \|r_2 - r_1\|_\infty \mathbf{1} \\
&= V_{\gamma, r_1}^* - \frac{\|r_2 - r_1\|_\infty}{1 - \gamma} \mathbf{1} \\
&\geq V_{\gamma, r_1}^{\pi_2^*} - \frac{\|r_2 - r_1\|_\infty}{1 - \gamma} \mathbf{1} \\
&= (I - \gamma P_{\pi_2^*})^{-1} M^{\pi_2^*} r_1 - \frac{\|r_2 - r_1\|_\infty}{1 - \gamma} \mathbf{1} \\
&= (I - \gamma P_{\pi_2^*})^{-1} M^{\pi_2^*} r_2 + (I - \gamma P_{\pi_2^*})^{-1} M^{\pi_2^*} (r_1 - r_2) - \frac{\|r_2 - r_1\|_\infty}{1 - \gamma} \mathbf{1} \\
&\geq (I - \gamma P_{\pi_2^*})^{-1} M^{\pi_2^*} r_2 + \|(I - \gamma P_{\pi_2^*})^{-1}\|_{\infty \rightarrow \infty} \|M^{\pi_2^*}\|_{\infty \rightarrow \infty} \|r_1 - r_2\|_\infty \mathbf{1} - \frac{\|r_2 - r_1\|_\infty}{1 - \gamma} \mathbf{1} \\
&= V_{\gamma, r_2}^* - 2 \frac{\|r_1 - r_2\|_\infty}{1 - \gamma} \mathbf{1}.
\end{aligned}$$

□

Proof of Theorem 8. The desired conclusion follows immediately once we verify that the guarantees within Theorems 3 and 6 both hold. The guarantees of Theorem 6 (regarding the performance of the perturbed empirical optimal policy) obviously hold, so our main task is to verify that Theorem 3 can be applied under the desired assumptions, which will be done by showing that the perturbation level $\xi = \frac{1}{2n^2}$ is sufficiently small so that, with high probability, $\hat{\pi}$ (the exact Blackwell-optimal policy of $(\hat{P}, \tilde{r})$) is also near-optimal for the unperturbed AMDP $(\hat{P}, r)$.

Instead of checking condition (2) on the optimality of $\hat{\pi}$ for the unperturbed AMDP (which could also be done, with smaller ξ and more effort), we instead notice that the proof of Theorem 3 only uses condition (2) to apply Lemma 27, which in turn verifies that (41) holds. We can thus instead directly check condition (41), which we recall is $\|\hat{V}_{1-\frac{1}{n}}^* - \hat{V}_{1-\frac{1}{n}}^{\hat{\pi}}\|_\infty \leq \frac{1}{n}$ or equivalently $\hat{V}_{1-\frac{1}{n}}^{\hat{\pi}} \geq \hat{V}_{1-\frac{1}{n}}^* - \frac{1}{n} \mathbf{1}$. Applying Lemma 37 (with $P = \hat{P}$, $r_2 = r$, $r_1 = \tilde{r}$, $\gamma = 1 - \frac{1}{n}$, and thus π_1^* is equal to $\hat{\pi}$ since by Lemma 24 $\hat{\pi}$ is also optimal for the DMDP $(\hat{P}, \tilde{r}, 1 - \frac{1}{n})$), we immediately obtain that

$$\hat{V}_{1-\frac{1}{n}}^{\hat{\pi}} \geq \hat{V}_{1-\frac{1}{n}}^* - \frac{2\|\tilde{r} - r\|_\infty}{1 - (1 - \frac{1}{n})} \mathbf{1} \geq \hat{V}_{1-\frac{1}{n}}^* - 2n\xi \mathbf{1} \geq \hat{V}_{1-\frac{1}{n}}^* - \frac{1}{n} \mathbf{1}$$

since by construction $\|\tilde{r} - r\|_\infty \leq \xi$ and $\xi = \frac{1}{2n^2}$. Now the desired result follows from applying both Theorems 3 and 6, and simplifying the constant and log factors (in particular, applying the union bound to bound the failure probability by 2δ and then adjusting the failure probability and absorbing this factor of 2, as well as using the fact that we have chosen $\xi = \frac{1}{2n^2}$). □

C.8 Proof of Theorem 1

Proof of Theorem 1. For consistency with Lemma 35 we will use $\hat{\pi}^*$ rather than $\hat{\pi}$ to denote the bias-optimal policy of $(\hat{P}, r)$ which is returned by `SolveAMDP`. We have that

$$\begin{aligned}
\rho^{\hat{\pi}^*} &\geq \hat{\rho}^{\hat{\pi}^*} + \|\hat{\rho}^{\hat{\pi}^*} - \rho^{\hat{\pi}^*}\|_\infty \mathbf{1} \\
&\geq \hat{\rho}^{\pi^*} + \|\hat{\rho}^{\hat{\pi}^*} - \rho^{\hat{\pi}^*}\|_\infty \mathbf{1} \\
&\geq \rho^{\pi^*} + \|\hat{\rho}^{\pi^*} - \rho^{\pi^*}\|_\infty \mathbf{1} + \|\hat{\rho}^{\hat{\pi}^*} - \rho^{\hat{\pi}^*}\|_\infty \mathbf{1}
\end{aligned} \tag{79}$$

where we used the fact that since $\hat{\pi}^*$ is bias-optimal, it is also gain-optimal, and thus $\hat{\rho}^{\hat{\pi}^*} = \hat{\rho}^* \geq \hat{\rho}^{\pi^*}$. Thus it remains to bound the terms $\|\hat{\rho}^{\pi^*} - \rho^{\pi^*}\|_\infty$ and $\|\hat{\rho}^{\hat{\pi}^*} - \rho^{\hat{\pi}^*}\|_\infty$ with high probability. First, since $\rho^{\pi^*} = \rho^*$ is a constant vector, we can apply Theorem 2 to bound

$$\|\hat{\rho}^{\pi^*} - \rho^{\pi^*}\|_\infty \leq \sqrt{\frac{C_4 \log^3\left(\frac{SAn}{\delta}\right)}{n}} (\|h^{\pi^*}\|_{\text{span}} + 1) \quad (80)$$

with probability at least $1 - \delta$. Next, to bound $\|\hat{\rho}^{\hat{\pi}^*} - \rho^{\hat{\pi}^*}\|_\infty$, we can combine Lemma 35 and Lemma 31 to obtain that with probability at least $1 - \delta$, if $\hat{P}$ is weakly communicating (which ensures $\hat{\rho}^*$ is a constant vector, as required by Lemma 31) then

$$\|\hat{\rho}^{\hat{\pi}^*} - \rho^{\hat{\pi}^*}\|_\infty \leq 2(\ell + 1) \left(\frac{\alpha (\|\hat{h}^*\|_{\text{span}} + 1)}{n} \right)^{\frac{1}{2}} + (\ell + 1) \frac{2\alpha}{n} (\|\hat{h}^*\|_{\text{span}} + 1)$$

where $\alpha = 16 \log\left(\frac{18.8SAn}{\delta(1-\gamma^*)^3}\right)$ and $\ell = \lceil \log_2 \log_2 (\|\hat{h}^*\|_{\text{span}} + 4) \rceil$. As in previous proofs, we can simplify by assuming $\alpha \frac{\|\hat{h}^*\|_{\text{span}} + 1}{n} \leq 1$, in which case $\left(\frac{\alpha (\|\hat{h}^*\|_{\text{span}} + 1)}{n} \right)^{\frac{1}{2}} \geq \frac{\alpha}{n} (\|\hat{h}^*\|_{\text{span}} + 1)$ and also $\|\hat{h}^*\|_{\text{span}} \leq n$. This is because if actually $\alpha \frac{\|\hat{h}^*\|_{\text{span}} + 1}{n} > 1$, then since always $\|\hat{\rho}^{\hat{\pi}^*} - \rho^{\hat{\pi}^*}\|_\infty \leq 1$, the desired theorem conclusion follows trivially. Thus continuing with the case that $\alpha \frac{\|\hat{h}^*\|_{\text{span}} + 1}{n} \leq 1$, we can bound

$$\ell \leq 1 + \log_2 \log_2 (\|\hat{h}^*\|_{\text{span}} + 4) \leq \log 4 \|\hat{h}^*\|_{\text{span}} \leq \log 4n$$

where we used that $\log_2 \log_2 (x + 4) \leq \log 4x$ for $x \geq 1$ and also that $\|\hat{h}^*\|_{\text{span}} \leq n$. Using this in combination with $\left(\frac{\alpha (\|\hat{h}^*\|_{\text{span}} + 1)}{n} \right)^{\frac{1}{2}} \geq \frac{\alpha}{n} (\|\hat{h}^*\|_{\text{span}} + 1)$ to simplify, we have that

$$\begin{aligned} \|\hat{\rho}^{\hat{\pi}^*} - \rho^{\hat{\pi}^*}\|_\infty &\leq 4(\ell + 1) \sqrt{\frac{\alpha (\|\hat{h}^*\|_{\text{span}} + 1)}{n}} \\ &\leq 4(\log(4n) + 1) \sqrt{\frac{\alpha (\|\hat{h}^*\|_{\text{span}} + 1)}{n}} \\ &\leq 4\alpha \sqrt{\frac{\alpha (\|\hat{h}^*\|_{\text{span}} + 1)}{n}} \\ &\leq \sqrt{\frac{16\alpha^3 (\|\hat{h}^*\|_{\text{span}} + 1)}{n}}. \end{aligned} \quad (81)$$

Plugging (80) and (81) into (79), and also halving the failure probability parameter of each to get an overall failure probability of $\leq \delta$ by the union bound, we obtain

$$\begin{aligned} \rho^{\hat{\pi}^*} &\geq \rho^{\pi^*} - \sqrt{\frac{C_4 \log^3\left(\frac{SAn}{\delta}\right)}{n}} (\|h^{\pi^*}\|_{\text{span}} + 1) - \sqrt{\frac{16\alpha^3 (\|\hat{h}^*\|_{\text{span}} + 1)}{n}} \mathbf{1} \\ &\geq \rho^{\pi^*} - \sqrt{\frac{C_4 \log^3\left(\frac{SAn}{\delta}\right)}{n}} (\|h^{\pi^*}\|_{\text{span}} + 1) + \frac{16\alpha^3 (\|\hat{h}^*\|_{\text{span}} + 1)}{n} \mathbf{1} \\ &\geq \rho^{\pi^*} - \sqrt{\frac{C_7 \log^3\left(\frac{SAn}{\delta(1-\gamma^*)}\right)}{n}} (\|h^{\pi^*}\|_{\text{span}} + \|\hat{h}^*\|_{\text{span}} + 1) \mathbf{1} \end{aligned}$$

where we used that $\sqrt{a} + \sqrt{b} \leq 2\sqrt{a+b}$ and then chose C_7 sufficiently large. $\square$

D Corollaries for Bounded Diameter or Uniformly Mixing MDPs

D.1 Proof of Lemma 4

Setup Let $s \in \mathcal{S}$, and we refer to it as the target state. Define the MDP $P^{\rightarrow s}$ and the reward vector $r^{\rightarrow s}$ by

$$P_{s'a'}^{\rightarrow s} = \begin{cases} e_s^\top & s' = s \\ P_{s'a'} & s' \neq s \end{cases}$$

$$r^{\rightarrow s}(s', a') = \begin{cases} 1 & s' = s \\ 0 & s' \neq s \end{cases}.$$

Note $e_s^\top$ is a vector which is all 0 except for a 1 in state s , meaning that the target state s is absorbing in $P^{\rightarrow s}$. All other states have identical transitions as in P . Also the reward is 1 in the target state s and 0 for all other states. Thus intuitively the optimal policy in $P^{\rightarrow s}$ should try to reach state s as quickly as possible.

In the rest of this subsection, we will exclusively use the discount factor $\gamma = 1 - \frac{1}{6D}$. We define $V_{\rightarrow s}^\pi$ to be the discounted value function in the MDP $(P^{\rightarrow s}, r^{\rightarrow s})$ with respect to a policy π , and we define $\hat{V}_{\rightarrow s}^\pi$ to be the value function for the “empirical” MDP $(\hat{P}^{\rightarrow s}, r^{\rightarrow s})$ where

$$\hat{P}_{s'a'}^{\rightarrow s} = \begin{cases} e_s^\top & s' = s \\ \hat{P}_{s'a'} & s' \neq s \end{cases}.$$

We note that there are two equivalent sampling processes for generating $\hat{P}^{\rightarrow s}$. The first process, suggested by the above definition, is sampling $\hat{P}$ as usual, and then making state s an absorbing state. The second process is to treat $P^{\rightarrow s}$ as if it were the true MDP transition matrix P , and then following the usual sampling process to generate $\hat{P}$ from P . These are equivalent because in the second process, with probability one all transitions from state s will return to state s , and thus it will be absorbing with probability 1. Using this correspondence, we will later be able to apply our Theorem 9 to $\hat{V}_{\rightarrow s}^\pi$.

We make a few final definitions. Let $\hat{D}$ be the diameter of $\hat{P}$. We also define the optimal value functions $V_{\rightarrow s}^*$ and $\hat{V}_{\rightarrow s}^*$, and let $\pi_{\rightarrow s}^*$ and $\hat{\pi}_{\rightarrow s}^*$ be the policies which attain these optimal value functions, respectively.

Correspondence between value functions and diameter Next we establish some basic facts about the above-defined value functions and their relationships to the diameters of $\hat{P}$ and P .

Letting $\mathbb{E}$ be the expectation with respect to the original MDP P and $T_s = \inf\{t \geq 0 : S_t = s\}$ be the first hitting time of state s (where we allow the “hit” to occur at time 0), we have

$$\begin{aligned} V_{\rightarrow s}^\pi(s') &= \mathbb{E}_{s'}^\pi \sum_{t=0}^{\infty} \gamma^t \mathbb{I}\{T_s \leq t\} \\ &= \mathbb{E}_{s'}^\pi \sum_{t=0}^{\infty} \gamma^t (1 - \mathbb{I}\{T_s > t\}) \\ &\geq \frac{1}{1-\gamma} - \mathbb{E}_{s'}^\pi \sum_{t=0}^{\infty} \mathbb{I}\{T_s > t\} \\ &= \frac{1}{1-\gamma} - \sum_{t=0}^{\infty} \mathbb{P}_{s'}^\pi(T_s > t) \\ &= \frac{1}{1-\gamma} - \mathbb{E}_{s'}^\pi T_s \end{aligned}$$

where we used the fact that $\sum_{t=0}^{\infty} \gamma^t = \frac{1}{1-\gamma}$, the monotone convergence theorem to interchange the expectation and the infinite sum, and the fact that since T_s takes values in the non-negative integers, $\sum_{t=0}^{\infty} \mathbb{P}_{s'}^\pi(T_s > t) = \mathbb{E}_{s'}^\pi T_s$.

Then

$$V_{\rightarrow s}^*(s') = \sup_{\pi} V_{\rightarrow s}^{\pi}(s') \geq \frac{1}{1-\gamma} - \inf_{\pi} \mathbb{E}_{s'}^{\pi} T_s = \frac{1}{1-\gamma} - D.$$

Therefore $V_{\rightarrow s}^* \geq \frac{1}{1-\gamma} \mathbf{1} - D \mathbf{1}$ elementwise. Since also $V_{\rightarrow s}^* \leq \frac{1}{1-\gamma} \mathbf{1}$, we have that $\|V_{\rightarrow s}^*\|_{\text{span}} \leq D$.

Now we show another relationship, that if $\hat{V}_{\rightarrow s}^*$ is sufficiently large (for all s), then $\hat{D}$ cannot be too large. We formalize this in the following lemma.

Lemma 38. *Suppose that $\gamma = 1 - \frac{1}{6D}$ and for all $s \in \mathcal{S}$, $\hat{V}_{\rightarrow s}^* \geq \frac{2}{3} \frac{1}{1-\gamma} \mathbf{1}$. Then*

$$\hat{D} \leq 12D \log(3).$$

Proof. As an intermediate step, we will show for all $s, s' \in \mathcal{S}$ that

$$\hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq 6D \log(3)) \geq \frac{1}{2}. \quad (82)$$

Fix a constant k and states $s, s' \in \mathcal{S}$. Then (even if k is not an integer) we have

$$\hat{V}_{\rightarrow s}^*(s') \leq \frac{1}{1-\gamma} \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq k) + \frac{\gamma^k}{1-\gamma} \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s > k). \quad (83)$$

Now we want to choose k so that only a small amount of value can be contributed from the $\frac{\gamma^k}{1-\gamma} \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s > k)$ term, specifically we will choose k so that $\frac{\gamma^k}{1-\gamma} \leq \frac{1}{3} \frac{1}{1-\gamma}$. We calculate that

$$\begin{aligned} \gamma^k \leq \frac{1}{3} &\iff k \log(\gamma) \leq \log \frac{1}{3} \\ &\iff k \geq \frac{\log \frac{1}{3}}{\log(\gamma)} = \frac{\log \frac{1}{3}}{\log(1 - \frac{1}{6D})} = \frac{\log 3}{-\log(1 - \frac{1}{6D})} \\ &\iff k \geq \frac{\log 3}{\frac{1}{6D}} = 6D \log(3) \end{aligned}$$

where in the final inequality we use the fact that $\log(1-x) \leq -x$ so $-\log(1-x) \geq x$. Then if we set $k = 6D \log(3)$ in (83), we have that

$$\begin{aligned} \hat{V}_{\rightarrow s}^*(s') &\leq \frac{1}{1-\gamma} \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq 6D \log(3)) + \frac{1}{3} \frac{1}{1-\gamma} \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s > 6D \log(3)) \\ &= \frac{1}{1-\gamma} \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq 6D \log(3)) + \frac{1}{3} \frac{1}{1-\gamma} \left(1 - \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq 6D \log(3))\right) \\ &= \frac{1}{3} \frac{1}{1-\gamma} + \frac{2}{3} \frac{1}{1-\gamma} \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq 6D \log(3)). \end{aligned}$$

Now since $\hat{V}_{\rightarrow s}^*(s') \geq \frac{2}{3} \frac{1}{1-\gamma}$ by assumption, we must have that $\hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq 6D \log(3)) \geq \frac{1}{2}$, establishing (82).

Next, for convenience let $b = \lfloor 6D \log(3) \rfloor$, and note that $\hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq 6D \log(3)) = \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq b)$. Then for fixed $s, s' \in \mathcal{S}$, using the Markov property,

$$\begin{aligned} \hat{\mathbb{E}}_{s'}^{\hat{\pi}^* \rightarrow s} T_s &\leq b \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s \leq b) + \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s > b) \left(b + \sup_{s'' \in \mathcal{S}} \hat{\mathbb{E}}_{s''}^{\hat{\pi}^* \rightarrow s} T_s \right) \\ &= b + \hat{\mathbb{P}}_{s'}^{\hat{\pi}^* \rightarrow s}(T_s > b) \sup_{s'' \in \mathcal{S}} \hat{\mathbb{E}}_{s''}^{\hat{\pi}^* \rightarrow s} T_s \\ &\leq b + \frac{1}{2} \sup_{s'' \in \mathcal{S}} \hat{\mathbb{E}}_{s''}^{\hat{\pi}^* \rightarrow s} T_s \end{aligned} \quad (84)$$

using the fact that $\widehat{\mathbb{P}}_{s' \rightarrow s}^{\pi^*}(T_s > b) \leq \frac{1}{2}$ for any s' from (82). Now taking the supremum over $s' \in \mathcal{S}$ in inequality (84) and rearranging, we obtain that

$$\sup_{s' \in \mathcal{S}} \widehat{\mathbb{E}}_{s' \rightarrow s}^{\pi^*} T_s \leq 2b \leq 12D \log(3).$$

Since this holds for all $s \in \mathcal{S}$, we conclude that

$$\begin{aligned} \widehat{D} &= \sup_{s, s' \in \mathcal{S}} \inf_{\pi} \widehat{\mathbb{E}}_{s'}^{\pi} T_s \\ &\leq \sup_{s, s' \in \mathcal{S}} \widehat{\mathbb{E}}_{s' \rightarrow s}^{\pi^*} T_s \\ &\leq 12D \log(3) \end{aligned}$$

as desired. $\square$

Lower-bounding $\widehat{V}_{\rightarrow s}^*$ Now it remains to complete the proof by setting $\gamma = 1 - \frac{1}{6D}$ and showing that for sufficiently large n , with high probability, we have $\widehat{V}_{\rightarrow s}^* \geq \frac{2}{3} \frac{1}{1-\gamma}$ for all $s \in \mathcal{S}$, and thus checking the conditions of Lemma 38.

Proof of Lemma 4. As observed earlier, we may apply our theorems on discounted MDPs to the S MDPs $(\widehat{P}^{\rightarrow s})_{s \in \mathcal{S}}$ (with discount factor $\gamma = 1 - \frac{1}{6D}$). Instead of using Theorem 9, it is more direct to use the bound (31) which appears in the final stage of the proof of Theorem 9. Taking a union bound over all S MDPs, we obtain that with probability at least $1 - S\delta$, for all $s \in \mathcal{S}$ we have

$$\left\| \widehat{V}_{\rightarrow s}^{\pi^*} - V_{\rightarrow s}^{\pi^*} \right\|_{\infty} \leq \frac{24 \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right)}{1-\gamma} \sqrt{\frac{\alpha \left(\|V_{\rightarrow s}^{\pi^*}\|_{\text{span}} + 1 \right)}{n}} \quad (85)$$

where $\alpha = 16 \log \left(\frac{12SAn}{(1-\gamma)^{2\delta}} \right) = 16 \log \left(\frac{12 \cdot 6^2 SAnD^2}{\delta} \right)$. As observed above, $\|V_{\rightarrow s}^{\pi^*}\|_{\text{span}} \leq D$, so we can simplify (85) to obtain (for all $s \in \mathcal{S}$)

$$\begin{aligned} \left\| \widehat{V}_{\rightarrow s}^{\pi^*} - V_{\rightarrow s}^{\pi^*} \right\|_{\infty} &\leq 144D \log_2 \log_2 \left(\frac{1}{1-\gamma} + 4 \right) \sqrt{\frac{\alpha(D+1)}{n}} \\ &\leq 144D \log(24D) \sqrt{\frac{\alpha(D+1)}{n}} \\ &\leq 144D \log(24D) \sqrt{\frac{2\alpha D}{n}} \\ &\leq D \end{aligned}$$

where we used that $\log_2 \log_2(x+4) \leq \log 4x$ for $x > 1$, that $D \geq 1$, and in the final inequality we assume that $n \geq 2 \cdot 144^2 \alpha D \log^2(24D)$.

In this event, we thus have that

$$\widehat{V}_{\rightarrow s}^* \geq \widehat{V}_{\rightarrow s}^{\pi^*} \geq V_{\rightarrow s}^{\pi^*} - \left\| \widehat{V}_{\rightarrow s}^{\pi^*} - V_{\rightarrow s}^{\pi^*} \right\|_{\infty} \mathbf{1} \geq \frac{1}{1-\gamma} \mathbf{1} - D \mathbf{1} - D \mathbf{1} = \frac{2}{3} \frac{1}{1-\gamma} \mathbf{1}$$

since $\frac{1}{1-\gamma} = 6D$. We can thus combine this fact with Lemma 38 to conclude that $\widehat{D} \leq 12D \log(3) \leq 14D$. We conclude by choosing the constant C_8 such that $n \geq C_8 D \log^3 \left(\frac{SADn}{\delta} \right)$ implies that $n \geq 2 \cdot 144^2 \alpha D \log^2(24D)$. $\square$

D.2 Proof of optimal diameter-based complexity Corollary 5

Proof of Corollary 5. This result follows from combining Theorem 3 (on the performance of Algorithm 1 with anchoring and no perturbation) with Lemma 4 which bounds the empirical diameter. First, by Theorem 3, with probability at least $1 - \delta$ we have that

$$\rho^{\hat{\pi}} - \rho^* \leq \sqrt{\frac{C_5 \log^3 \left(\frac{SAn}{\delta} \right)}{n}} \left(\|h^*\|_{\text{span}} + \|\hat{h}^{\hat{\pi}}\|_{\text{span}} + 1 \right) \mathbf{1}. \quad (86)$$

It remains to bound the terms $\|h^*\|_{\text{span}}$ and $\|\hat{h}^{\hat{\pi}}\|_{\text{span}}$ in terms of D . First, it is well-known that $\|h^*\|_{\text{span}} \leq D$ [Bartlett and Tewari 2012]. (As pointed out in [Lattimore and Szepesvári, 2020, Exercise 38.13], the proof of this bound provided in Bartlett and Tewari [2012] is incomplete, but Lattimore and Szepesvári [2020] provide a complete proof.) Next, by applying Lemma 4, if $n \geq C_7 D \log^3 \left(\frac{SADn}{\delta} \right)$ then with additional failure probability at most δ , the diameter of $\hat{P}$, $\hat{D}$, is bounded by $14D$. By condition (2), we have $\|\hat{h}^{\hat{\pi}} - \hat{h}^*\|_{\infty} \leq \frac{1}{3n}$, which implies

$$\|\hat{h}^{\hat{\pi}}\|_{\text{span}} \leq \|\hat{h}^*\|_{\text{span}} + \|\hat{h}^{\hat{\pi}} - \hat{h}^*\|_{\text{span}} \leq \|\hat{h}^*\|_{\text{span}} + 2 \|\hat{h}^{\hat{\pi}} - \hat{h}^*\|_{\infty} \leq \|\hat{h}^*\|_{\text{span}} + \frac{2}{3n} \leq \|\hat{h}^*\|_{\text{span}} + 1.$$

Additionally, using Lemma 24 and then the bound from Bartlett and Tewari [2012] again, and then the bound on $\hat{D}$, we have that

$$\|\hat{h}^*\|_{\text{span}} \leq 2\|\hat{h}^*\|_{\text{span}} \leq 2\hat{D} \leq 28D.$$

(Note that we can apply the bound $\|\hat{h}^*\|_{\text{span}} \leq 2\|\hat{h}^*\|_{\text{span}}$ from Lemma 24 because we are operating under the event that $\hat{D} \leq 14D$, which in particular implies that $\hat{P}$ is communicating so it an optimal gain $\hat{\rho}^*$ which is a constant vector.) Plugging these bounds into (86) and simplifying, we obtain that

$$\begin{aligned} \rho^{\hat{\pi}} - \rho^* &\leq \sqrt{\frac{C_5 \log^3 \left(\frac{SAn}{\delta} \right)}{n}} (D + 28D + 1) \mathbf{1} \\ &\leq \sqrt{\frac{C_5 \log^3 \left(\frac{SAn}{\delta} \right)}{n}} 30D \mathbf{1} \\ &\leq \sqrt{\frac{\max\{30C_5, C_8\} \log^3 \left(\frac{SADn}{\delta} \right)}{n}} D \mathbf{1}. \end{aligned}$$

The last inequality ensures that whenever the above bound is non-trivial (the RHS is < 1), then the condition $n \geq C_8 D \log^3 \left(\frac{SADn}{\delta} \right)$ will be satisfied. Therefore we can conclude by choosing C_9 so that $C_9 \log^3 \left(\frac{SADn}{\delta} \right) \geq \max\{30C_5, C_8\} \log^3 \left(\frac{SADn}{\delta/2} \right)$ (where we have added the of 2 so that the total failure probability is bounded by δ).

□

D.3 Proof of optimal mixing-based complexity Corollary 7

First we collect some simple facts regarding the relationships between bias and discounted value functions and τ_{unif} for uniformly mixing MDPs. These results are completely standard but we provide their proofs for completeness.

Lemma 39. *Let P be any MDP which has a bounded uniform mixing parameter τ_{unif} . Then for any Markovian deterministic policy π ,*

1. $\|h^{\pi}\|_{\text{span}} \leq 3\tau_{\text{unif}}$.
2. For any discount factor $\gamma \in [0, 1)$, $\|V_{\gamma}^{\pi}\|_{\text{span}} \leq 3\tau_{\text{unif}}$.

(We note that the results would also hold for any randomized policy as well if τ_{unif} were defined over randomized policies.)

Proof. 1. A bound of this form is essentially claimed in [Wang et al., 2022, Lemma 9, Proposition 10], although we believe that [Wang et al., 2022, Proposition 10] is not needed, since in a mixing MDP, in the Markov chain induced by a policy π , there must be a unique stationary distribution in order for the mixing time to be defined, and thus there should only be one closed recurrent class. We also believe that there may be a missing factor of 2 in the proof of [Wang et al., 2022, Lemma 9] (in the second inequality step). Thus we choose to reprove this bound for completeness (and get a better constant), but we essentially follow their arguments.

By [Jin and Sidford, 2021, Lemma 1], by uniform mixing we have

$$\|P_\pi^k - P_\pi^\infty\|_{\infty \rightarrow \infty} \leq 2^{-\lfloor k/\tau_{\text{unif}} \rfloor} \quad (87)$$

for all $k \geq \tau_{\text{unif}}$. Also since P_π must be aperiodic since the mixing time is finite, we have [Puterman, 1994] that

$$h^\pi(s) = \lim_{T \rightarrow \infty} \mathbb{E}_s^\pi \left[\sum_{t=0}^{T-1} R_t - T\rho^\pi(s) \right] = \lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} e_s^\top (P_\pi)^t r_\pi - T\rho^\pi(s).$$

Therefore we have

$$\begin{aligned} \|h^\pi\|_{\text{span}} &= \left\| \lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} (P_\pi)^t r_\pi - T\rho^\pi \right\|_{\text{span}} \\ &= \lim_{T \rightarrow \infty} \left\| \sum_{t=0}^{T-1} (P_\pi)^t r_\pi - T\rho^\pi \right\|_{\text{span}} \\ &= \lim_{T \rightarrow \infty} \left\| \sum_{t=0}^{T-1} (P_\pi^t - P_\pi^\infty) r_\pi \right\|_{\text{span}} \\ &\leq \lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \| (P_\pi^t - P_\pi^\infty) r_\pi \|_{\text{span}} \\ &= \sum_{t=0}^{\infty} \| (P_\pi^t - P_\pi^\infty) r_\pi \|_{\text{span}}. \end{aligned}$$

Now we can bound this using (87):

$$\begin{aligned}
\sum_{t=0}^{\infty} \|(P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi}\|_{\text{span}} &\leq \sum_{t=0}^{\tau_{\text{unif}}-1} \|(P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi}\|_{\text{span}} + \sum_{t=\tau_{\text{unif}}}^{\infty} \|(P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi}\|_{\text{span}} \\
&= \sum_{t=0}^{\tau_{\text{unif}}-1} \|P_{\pi}^t r_{\pi}\|_{\text{span}} + \sum_{t=\tau_{\text{unif}}}^{\infty} \|(P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi}\|_{\text{span}} \\
&\leq \tau_{\text{unif}} + \sum_{t=\tau_{\text{unif}}}^{\infty} \|(P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi}\|_{\text{span}} \\
&\leq \tau_{\text{unif}} + \sum_{t=\tau_{\text{unif}}}^{\infty} 2 \|(P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi}\|_{\infty} \\
&\leq \tau_{\text{unif}} + \sum_{t=\tau_{\text{unif}}}^{\infty} 2 \|P_{\pi}^t - P_{\pi}^{\infty}\|_{\infty \rightarrow \infty} \|r_{\pi}\|_{\infty} \\
&\leq \tau_{\text{unif}} + 2 \sum_{t=\tau_{\text{unif}}}^{\infty} 2^{-\lfloor t/\tau_{\text{unif}} \rfloor} \\
&= \tau_{\text{unif}} + 2 \sum_{k=1}^{\infty} \tau_{\text{unif}} 2^{-k} \\
&= \tau_{\text{unif}} + 2\tau_{\text{unif}} = 3\tau_{\text{unif}}
\end{aligned}$$

where in the first equality step we used that $P_{\pi}^{\infty} r_{\pi} = \rho^{\pi}$ is a constant vector, then we used that $\|P_{\pi}^t r_{\pi}\|_{\text{span}} \leq 1$ (since $1 \geq P_{\pi}^t r_{\pi} \geq 0$ elementwise).

2. We repeat a very similar argument. By the Neumann series expansion,

$$\begin{aligned}
V_{\gamma}^{\pi} &= (I - \gamma P_{\pi})^{-1} r_{\pi} \\
&= \sum_{t=0}^{\infty} \gamma^t P_{\pi}^t r_{\pi}.
\end{aligned}$$

Since ρ^{π} is a constant vector,

$$\begin{aligned}
\|V_{\gamma}^{\pi}\|_{\text{span}} &= \|V_{\gamma}^{\pi} - \frac{1}{1-\gamma} \rho^{\pi}\|_{\text{span}} \\
&= \left\| \sum_{t=0}^{\infty} \gamma^t P_{\pi}^t r_{\pi} - \frac{1}{1-\gamma} P_{\pi}^{\infty} r_{\pi} \right\|_{\text{span}} \\
&= \left\| \sum_{t=0}^{\infty} \gamma^t (P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi} \right\|_{\text{span}} \\
&\leq \sum_{t=0}^{\infty} \|\gamma^t (P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi}\|_{\text{span}} \\
&\leq \sum_{t=0}^{\infty} \|(P_{\pi}^t - P_{\pi}^{\infty}) r_{\pi}\|_{\text{span}}
\end{aligned}$$

where we used that $\sum_{t=0}^{\infty} \gamma^t = \frac{1}{1-\gamma}$ in the third equality. We can conclude by noting that in the previous part of this lemma we have already bounded this exact final term by $3\tau_{\text{unif}}$. $\square$

Proof of Corollary 7. This follows immediately from Theorem 6, since both $\|h^*\|_{\text{span}}$ and $\|\underline{h}^{\hat{\pi}}\|_{\text{span}}$ can be bounded by $3\tau_{\text{unif}}$ (using Lemma 39, in particular noting that $\|\underline{h}^{\hat{\pi}}\|_{\text{span}} = \|V_{1/n}\|_{\hat{\pi}}^{\hat{\pi}}$ by Lemma 24 so we can apply the second part of Lemma 39.) Also $\tau_{\text{unif}} \geq 1$, so $\|h^*\|_{\text{span}} + \|\underline{h}^{\hat{\pi}}\|_{\text{span}} + 1 \leq 7\tau_{\text{unif}}$. $\square$

E Other DMDP Results

Proof of Theorem 11. Analogously to the proof of Theorem 8, we will show that for sufficiently small perturbation, the exact solution of the perturbed empirical DMDP $(\hat{P}, \tilde{r}, \gamma)$ is a sufficiently small approximate solution of the unperturbed empirical DMDP $(\hat{P}, r, \gamma)$.

Let $\hat{\pi}$ be the exact solution of the perturbed empirical DMDP $(\hat{P}, \tilde{r}, \gamma)$. By applying Lemma 37 with $r_1 = \tilde{r}$, $r_2 = r$, and $P = \hat{P}$, we obtain that

$$\hat{V}^{\hat{\pi}} \geq \hat{V}^* - 2 \frac{\|\tilde{r} - r\|_\infty}{1 - \gamma} \geq \hat{V}^* - 2 \frac{\xi}{1 - \gamma} = \hat{V}^* - \frac{1}{n}.$$

Therefore the policy $\hat{\pi}$ satisfies the conditions of both Theorem 9 and Theorem 10. Applying both theorems and taking the union bound, $\hat{\pi}$ satisfies both guarantees with probability at least $1 - 2\delta$. Adjusting the constants to make the overall failure probability δ , and absorbing $\xi = \frac{1-\gamma}{2n}$ into the other terms within the log factor of the guarantee of Theorem 10, we can immediately conclude by choosing C_{10} sufficiently large. $\square$

Proof of Lemma 12. We will use Theorem 9 to prove this result. By inspecting the proof of the theorem, specifically the chain of inequalities (33), on the same event that the theorem holds (which is a probability at least $1 - \delta$ event) we have that

$$\|V^* - \hat{V}^*\|_\infty \leq \frac{1}{1 - \gamma} \sqrt{\frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n}} \left(\|V^*\|_{\text{span}} + \|\hat{V}^{\hat{\pi}}\|_{\text{span}} + 1 \right). \quad (88)$$

Also the optimality condition on $\hat{\pi}$, equation (3), implies that $\hat{V}^* \geq \hat{V}^{\hat{\pi}} \geq \hat{V}^* - \frac{1}{n}$, which in turn implies that $\|\hat{V}^{\hat{\pi}}\|_{\text{span}} \leq \|\hat{V}^*\|_{\text{span}} + \frac{1}{n}$. Also by triangle inequality we have that

$$\|\hat{V}^*\|_{\text{span}} \leq \|V^*\|_{\text{span}} + \|\hat{V}^* - V^*\|_{\text{span}} \leq \|V^*\|_{\text{span}} + 2 \|V^* - \hat{V}^*\|_\infty.$$

Plugging both of these bounds into (88), we obtain that (again, on the event that Theorem 9 holds)

$$\begin{aligned} \|V^* - \hat{V}^*\|_\infty &\leq \frac{1}{1 - \gamma} \sqrt{\frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n}} \left(\|V^*\|_{\text{span}} + \frac{1}{n} + \|V^*\|_{\text{span}} + 2 \|V^* - \hat{V}^*\|_\infty + 1 \right) \\ &\leq \frac{1}{1 - \gamma} \sqrt{\frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n}} \left(2 \|V^*\|_{\text{span}} + 2 \|V^* - \hat{V}^*\|_\infty + 2 \right) \\ &\leq \frac{1}{1 - \gamma} \sqrt{\frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n}} \left(4 \|h^*\|_{\text{span}} + 2 \|V^* - \hat{V}^*\|_\infty + 2 \right) \end{aligned} \quad (89)$$

using that $\|V^*\|_{\text{span}} \leq 2 \|h^*\|_{\text{span}}$ [Wei et al., 2020, Lemma 2] in the last inequality, which holds since we assumed that P is weakly communicating. Now squaring both sides and rearranging, we obtain

$$\|V^* - \hat{V}^*\|_\infty^2 - \frac{2C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} \|V^* - \hat{V}^*\|_\infty - \frac{2C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} (2 \|h^*\|_{\text{span}} + 1) \leq 0.$$

Using the larger root given by the quadratic formula for this polynomial in $\|V^* - \hat{V}^*\|_\infty$, we can bound

$$\begin{aligned}
& \|V^* - \hat{V}^*\|_\infty \\
& \leq \frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} + \frac{1}{2} \sqrt{\left(\frac{2C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} \right)^2 + \frac{8C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} (2\|h^*\|_{\text{span}} + 1)} \\
& \leq \frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} + \frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} + \sqrt{\frac{2C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} (2\|h^*\|_{\text{span}} + 1)} \tag{90}
\end{aligned}$$

where in the second inequality we used that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$. Now if we assume that $n \geq \frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{(1-\gamma)^2 (\|h^*\|_{\text{span}} + 1)}$, or equivalently that $\frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n(1-\gamma)^2} \leq \|h^*\|_{\text{span}} + 1$, then plugging this into (90), we obtain that

$$\|V^* - \hat{V}^*\|_\infty \leq 2(\|h^*\|_{\text{span}} + 1) + \sqrt{2(\|h^*\|_{\text{span}} + 1)(2\|h^*\|_{\text{span}} + 1)} \leq 4(\|h^*\|_{\text{span}} + 1)$$

as desired. $\square$

Proof of Theorem 13. From the choice of γ , we have that

$$\frac{1}{1-\gamma} = \sqrt{\frac{n}{C_1 \log^3 \left(\frac{SAn^2}{\delta} \right)}} \leq \sqrt{\frac{n}{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}}$$

(where for the inequality we use the coarse bound that $\frac{1}{1-\gamma} \leq n$). We can thus immediately check from this inequality that the condition of Lemma 12 is satisfied (note we also assume P is weakly communicating), so we obtain that with probability at least $1 - \delta$, $\|\hat{V}^*\|_{\text{span}} \leq 4(\|h^*\|_{\text{span}} + 1)$. Also by the condition (3), we have $\hat{V}^* \geq \hat{V}^{\hat{\pi}} \geq \hat{V}^* - \frac{1}{n}$ so $\|\hat{V}^{\hat{\pi}}\|_{\text{span}} \leq \|\hat{V}^*\|_{\text{span}} + \frac{1}{n}$. Plugging these bounds into the guarantee from Theorem 9, as well as the bound $\|V^*\|_{\text{span}} \leq 2\|h^*\|_{\text{span}}$ from [Wei et al., 2020, Lemma 2] (which holds since P is weakly communicating), we obtain that

$$\begin{aligned}
\|V^{\hat{\pi}} - V^*\|_\infty & \leq \frac{1}{1-\gamma} \sqrt{\frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n} (\|V^*\|_{\text{span}} + \|\hat{V}^{\hat{\pi}}\|_{\text{span}} + 1)} \\
& \leq \frac{1}{1-\gamma} \sqrt{\frac{C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n} \left(2\|h^*\|_{\text{span}} + \frac{1}{n} + 4(\|h^*\|_{\text{span}} + 1) + 1 \right)} \\
& \leq \frac{1}{1-\gamma} \sqrt{\frac{6C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n} (\|h^*\|_{\text{span}} + 1)}. \tag{91}
\end{aligned}$$

Now, we pause to restate the the main AMDP-to-DMDP reduction result from Wang et al. [2022] (in a form closer more immediately useful for us):

Theorem 40. [Wang et al., 2022, Theorem 1] *If P is weakly communicating, then for any policy π , we have*

$$\rho^* - \rho^\pi \leq (1-\gamma) \left(8\|h^*\|_{\text{span}} + 3\|V_\gamma^\pi - V_\gamma^*\|_\infty \right) \mathbf{1}.$$

Plugging our bound on $\|V^{\hat{\pi}} - V^*\|_\infty$ and our choice of γ into this theorem, we obtain that

$$\rho^* - \rho^{\hat{\pi}} \leq 8\|h^*\|_{\text{span}} \sqrt{\frac{C_1 \log^3 \left(\frac{SAn^2}{\delta} \right)}{n}} + 3 \frac{1-\gamma}{1-\gamma} \sqrt{\frac{6C_1 \log^3 \left(\frac{SAn}{(1-\gamma)\delta} \right)}{n} (\|h^*\|_{\text{span}} + 1)} \mathbf{1}$$

and by using the inequality $\sqrt{a} + \sqrt{b} \leq 2\sqrt{a+b}$, that $\|h^*\|_{\text{span}} \leq \frac{1}{2}\|h^*\|_{\text{span}}^2 + \frac{1}{2}$ by AM-GM, and choosing C_{11} appropriately, the RHS can be bounded by

$$\rho^* - \rho^{\hat{\pi}} \leq \sqrt{\frac{C_{11} \log^3\left(\frac{SAn}{\delta}\right)}{n}} (\|h^*\|_{\text{span}}^2 + 1) \mathbf{1}$$

as desired. (Note that the total failure probability is actually just $\leq \delta$ since by the proof of Lemma 12, the event of Lemma 12 is contained within the event that the bound from Theorem 9 holds.) $\square$

F DMDP Reduction Approach

Here we provide theorems with identical guarantees as to those of Theorems 3 and 6, but instead of requiring solutions to the anchored (resp., perturbed) empirical AMDP $(\hat{P}, r)$ (resp., $(\hat{P}, \tilde{r})$), the optimality condition is expressed in terms of solutions to the empirical (resp., perturbed empirical) DMDPs $(\hat{P}, r, \gamma)$ (resp., $(\hat{P}, \tilde{r}, \gamma)$) with $\gamma = 1 - \frac{1}{n}$. Thus, the same conclusions hold for the DMDP reduction approach using an effective horizon of $\frac{1}{1-\gamma} = n$, which does not require prior knowledge. Because of the close connection between anchoring and horizon- n -discounted reductions, the proofs are completely trivial.

Theorem 41. *Suppose P is weakly communicating. Set $\xi = 0$ and $\gamma = 1 - \frac{1}{n}$ in Algorithm 2. Also suppose that SolveDMDP is guaranteed to return a policy $\hat{\pi}_\gamma$ satisfying*

$$\left\| \hat{V}_{1-\frac{1}{n}}^* - \hat{V}_{1-\frac{1}{n}}^{\hat{\pi}_\gamma} \right\|_\infty \leq \frac{1}{n}.$$

Then with probability at least $1 - \delta$,

$$\rho^{\hat{\pi}_\gamma} - \rho^* \leq \sqrt{\frac{C_5 \log^3\left(\frac{SAn}{\delta}\right)}{n}} \left(\|h^*\|_{\text{span}} + \|\hat{V}_{1-\frac{1}{n}}^*\|_{\text{span}} + 1 \right) \mathbf{1}.$$

Proof. Note that the proof of Theorem 3 uses Lemma 27 to show that any $\hat{\pi}$ satisfying the optimality condition (2) which appears in the statement of Theorem 3 also satisfies the condition (41), and then the rest of the proof only uses the fact that $\hat{\pi}$ satisfies the condition (41). Since the above requirement that $\left\| \hat{V}_{1-\frac{1}{n}}^* - \hat{V}_{1-\frac{1}{n}}^{\hat{\pi}_\gamma} \right\|_\infty \leq \frac{1}{n}$ is exactly the condition (41), the rest of the proof immediately goes through for $\hat{\pi}_\gamma$ in place of $\hat{\pi}$. Lastly, we can use Lemma 24 to obtain that $\|\underline{h}^*\|_{\text{span}} = \|\hat{V}_{1-\frac{1}{n}}^*\|_{\text{span}}$ and thus replace the $\|\underline{h}^*\|_{\text{span}}$ term which appears in Theorem 3. $\square$

Theorem 42. *Suppose P is weakly communicating. Set $\xi \in (0, \frac{1}{n}]$ and $\gamma = 1 - \frac{1}{n}$ in Algorithm 2. Also suppose that the policy $\hat{\pi}_\gamma$ returned by SolveDMDP is guaranteed to be the exact discounted optimal policy of the DMDP $(\hat{P}, \tilde{r}, \gamma)$. Then with probability at least $1 - \delta$,*

$$\rho^* - \rho^{\hat{\pi}_\gamma} \leq \sqrt{\frac{C_3 \log^3\left(\frac{SAn}{\delta\xi}\right)}{n}} \left(\|h^*\|_{\text{span}} + \|\hat{V}_{1-\frac{1}{n}}^{\hat{\pi}_\gamma}\|_{\text{span}} + 1 \right) \mathbf{1}.$$

Proof. As noted in the proof of Theorem 6, under the event that the theorem's guarantee holds, the policy $\hat{\pi}$ which is the exact Blackwell-optimal policy of the AMDP $(\hat{P}, \tilde{r})$ is identical to the exact discounted-optimal policy of the DMDP $(\hat{P}, \tilde{r}, 1 - \frac{1}{n})$, so we immediately obtain the bound which appears in Theorem 6. Then using Lemma 24, we have $\|\underline{h}^{\hat{\pi}_\gamma}\|_{\text{span}} = \|\hat{V}_{1-\frac{1}{n}}^{\hat{\pi}_\gamma}\|_{\text{span}}$, so we can replace the $\|\underline{h}^{\hat{\pi}_\gamma}\|_{\text{span}}$ term which appears in Theorem 6. $\square$

G Proof of Theorem 14

Proof of Theorem 14. First we provide the MDP, P as well as an MDP $\hat{P}$ which has a constant probability of being sampled from P .




Figure 1: A true MDP P and an MDP $\hat{P}$ which has constant probability of being sampled from P when n samples are drawn from each state-action pair. Dashed lines are used to indicate all possible stochastic next-state transitions after taking a given action, with each dashed line being annotated with the probability of the particular next-state transition. They differ only in state-action pair $(s, a) = (1, 1)$, for which $P(2 | 1, 1) = \frac{1}{n}$ but $\hat{P}(2 | 1, 1) = 0$.

Since there are two states we encode each Markovian deterministic policy as a tuple $(\pi(1), \pi(2))$. First we compute properties of P , starting with the gains of the Markovian deterministic policies. It is easy to see that both policies which take action 1 in state 2 will stay in state 2 once reaching it. Also it is easy to see that both policies which take action 2 in state 2 will converge to a limiting distribution which is uniform over the two states. Therefore we will have that

$$\rho^{(1,1)} = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}, \rho^{(2,1)} = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}, \rho^{(1,2)} = \begin{bmatrix} \frac{1}{4} + \frac{1}{2n} \\ \frac{1}{4} + \frac{1}{2n} \end{bmatrix}, \rho^{(2,2)} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

Thus policies $(1, 1)$ and $(2, 1)$ are both gain-optimal and the optimal gain is $\rho^* = \frac{1}{2}\mathbf{1}$. Comparing between these two policies it is clear that $(1, 1)$ is the only bias-optimal policy (and thus the only Blackwell-optimal policy), since these two gain-optimal policies induce the same distribution over states (the Markov chains $P_{(1,1)}$ and $P_{(2,1)}$ are equal), and thus the only difference is the reward obtained in state 1, which is larger if action 1 is taken. Now we compute h^* . The Markov chain $P_{(1,1)}$ is eventually absorbed in state 2, so (since we must then have $P_{(1,1)}^\infty = \mathbf{1}[0, 1]$ and we have $P_{(1,1)}^\infty h^* = 0$) it is immediate that $h^*(2) = 0$. Using the first row of the equality $\rho^* + h^* = r_{(1,1)} + P_{(1,1)}h^*$, we have

$$\begin{aligned} \frac{1}{2} + h^*(1) &= \frac{1}{2} + \frac{1}{n} + \left(1 - \frac{1}{n}\right) h^*(1) + \frac{1}{n} h^*(2) = \frac{1}{2} + \frac{1}{n} + \left(1 - \frac{1}{n}\right) h^*(1) \\ \iff \frac{1}{n} h^*(1) &= \frac{1}{n} \\ \iff h^*(1) &= 1. \end{aligned}$$

Thus $\|h^*\|_{\text{span}} = 1$.

It is immediate that the diameter of this MDP is equal to $D = n$, since the expected hitting time of state 2 from state 1 (from any policy) is the expected value of a $\text{Geom}(\frac{1}{D})$ RV (with range $\{1, 2, \dots\}$), which is D , and likewise for reaching state 1 from state 2, only action 2 leads to this state, and it again has expected hitting time of $n = D$. To calculate τ_{unif} for P , since either action taken in state 1 has the same next-state distribution, we simply need to compute the mixing times of the Markov chains with transition matrices

$$\begin{bmatrix} 1 - \frac{1}{n} & \frac{1}{n} \\ \frac{1}{n} & 1 - \frac{1}{n} \end{bmatrix}, \begin{bmatrix} 1 - \frac{1}{n} & \frac{1}{n} \\ 0 & 1 \end{bmatrix}.$$

The first of these matrices is symmetric and (as mentioned before, or is trivial to check) has stationary distribution $[\frac{1}{2}, \frac{1}{2}]$ and is irreducible, and thus its mixing time is bounded by a constant times its relaxation time [Levin and Peres, 2017, Theorem 12.4]. This matrix has trace $2 - \frac{2}{n}$ and determinant $(1 - \frac{1}{n})^2 - \frac{1}{n^2} = 1 - \frac{2}{n}$, so its eigenvalues are 1 and $1 - \frac{2}{n}$. Therefore this matrix has absolute spectral gap $\frac{2}{n}$ and thus relaxation time $\frac{n}{2}$, so it has a mixing time which is $\Theta(n)$. For the second of the matrices in the above display, since the stationary distribution is simply a distribution with all its mass on state 2, the mixing time is simply the smallest integer t such that $\mathbb{P}_1(\eta_2 \leq t) \geq \frac{1}{2}$ where η_2 is the hitting time of state 2 (and $\mathbb{P}_1$ is the probability distribution induced by this Markov chain starting in state 1). Since $\eta_2 \sim \text{Geom}(\frac{1}{n})$ (where this is the Geometric random variable taking values in the range $\{1, 2, \dots\}$), setting $t = n \ln 2$ we obtain

$$\mathbb{P}_1(\eta_2 \leq n \ln 2) = 1 - \left(\left(1 - \frac{1}{n}\right)^n \right)^{\ln 2} \geq 1 - \left(\frac{1}{e}\right)^{\ln 2} = 1 - \frac{1}{2} = \frac{1}{2}$$

where we used the fact that $n \mapsto (1 - \frac{1}{n})^n$ is an increasing function which approaches $1/e$. Thus the Markov chain associated with this second matrix also has mixing time $\Theta(n)$, so we have $\tau_{\text{unif}} = \Theta(n)$.

Later we will need the fact that $\|V_\gamma^* - V_\gamma^{(1,2)}\|_\infty \geq \frac{n^2}{5}$, which we will check now by showing that $V_\gamma^{(1,2)}(2) \leq V_\gamma^*(2) - \frac{n^2}{5}$. First we compute $V_\gamma^{(1,2)}$. For convenience we will abbreviate this as V . By the Bellman evaluation equations we have

$$\begin{aligned} V(1) &= \frac{1}{2} + \frac{1}{n} + \gamma \left(1 - \frac{1}{n}\right) V(1) + \frac{\gamma}{n} V(2) \\ &= \frac{1}{2} + \frac{1}{n} + \left(1 - \frac{1}{n^2}\right) \left(1 - \frac{1}{n}\right) V(1) + \frac{1 - \frac{1}{n^2}}{n} V(2) \\ &= \frac{1}{2} + \frac{1}{n} + \left(1 - \frac{1}{n} - \frac{1}{n^2} + \frac{1}{n^3}\right) V(1) + \frac{n^2 - 1}{n^3} V(2) \\ &= \frac{1}{2} + \frac{1}{n} + \left(1 - \frac{n^2 + n - 1}{n^3}\right) V(1) + \frac{n^2 - 1}{n^3} V(2) \end{aligned}$$

which implies

$$V(1) = \frac{n^3}{n^2 + n - 1} \left(\frac{1}{2} + \frac{1}{n} + \frac{n^2 - 1}{n^3} V(2) \right). \quad (92)$$

Writing the other evaluation equation, we have

$$\begin{aligned} V(2) &= \gamma \left(1 - \frac{1}{n}\right) V(2) + \frac{\gamma}{n} V(1) \\ \iff \left(1 - \left(1 - \frac{1}{n^2} - \frac{1}{n} + \frac{1}{n^3}\right)\right) V(2) &= \left(1 - \gamma \left(1 - \frac{1}{n}\right)\right) V(2) = \frac{\gamma}{n} V(1) = \frac{1 - \frac{1}{n^2}}{n} V(1) \\ \iff \frac{n^2 + n - 1}{n^3} V(2) &= \frac{n^2 - 1}{n^3} V(1). \end{aligned}$$

Substituting (92) into the above display,

$$\begin{aligned}
& \frac{n^2 + n - 1}{n^3} V(2) = \frac{n^2 - 1}{n^3} \frac{n^3}{n^2 + n - 1} \left(\frac{1}{2} + \frac{1}{n} + \frac{n^2 - 1}{n^3} V(2) \right) \\
& \iff \left(\frac{n^2 + n - 1}{n^3} - \frac{n^2 - 1}{n^2 + n - 1} \frac{n^2 - 1}{n^3} \right) V(2) = \frac{n^2 - 1}{n^2 + n - 1} \left(\frac{1}{2} + \frac{1}{n} \right) \\
& \iff \left(\frac{(n^2 + n - 1)^2 - (n^2 - 1)^2}{(n^2 + n - 1)n^3} \right) V(2) = \frac{n^2 - 1}{n^2 + n - 1} \left(\frac{1}{2} + \frac{1}{n} \right) \\
& \iff \left(\frac{(n^2 + n - 1)^2 - (n^2 - 1)^2}{n^3} \right) V(2) = (n^2 - 1) \left(\frac{1}{2} + \frac{1}{n} \right) \\
& \iff \left(\frac{n(2n^2 + n - 2)}{n^3} \right) V(2) = (n^2 - 1) \left(\frac{1}{2} + \frac{1}{n} \right) \\
& \iff \frac{2n^2 + n - 2}{n^2 - 1} V(2) = n^2 \left(\frac{1}{2} + \frac{1}{n} \right).
\end{aligned}$$

Therefore

$$2V(2) < \frac{2n^2 + n - 2}{n^2 - 1} V(2) = n^2 \left(\frac{1}{2} + \frac{1}{n} \right) < n^2 \left(\frac{1}{2} + \frac{1}{10} \right)$$

(using that $n \geq 10$) so $V(2) < n^2 \frac{1}{2} \left(\frac{1}{2} + \frac{1}{10} \right) = n^2 \frac{3}{10}$.

Now we can easily observe $V_\gamma^*(2) \geq \frac{1}{1-\gamma} \frac{1}{2} = n^2 \frac{1}{2}$ since there exists a policy which takes action 1 in that state and is absorbed there, collecting reward $\frac{1}{2}$ at all times (in fact this action is optimal but we don't need to check this). Therefore we have that

$$V_\gamma^*(2) - V_\gamma^{(1,2)}(2) \geq n^2 \frac{1}{2} - n^2 \frac{3}{10} = \frac{n^2}{5},$$

which implies that $\|V_\gamma^* - V_\gamma^{(1,2)}\|_\infty \geq \frac{n^2}{5}$.

Now we check that the probability of $\hat{P}$ being equal to the instance displayed above is at least $\frac{1}{25}$. There are 4 state-action pairs which are sampled independently so we can compute the probability for each state-action pair separately. There are only two states, so we can encode $\hat{P}$ with the values of a random variable $N(s, a)$ which for each s, a counts how many transitions to state 1 are observed. We have $N(1, 1) \sim \text{Binom}(n, 1 - \frac{1}{n})$, $N(1, 2) \sim \text{Binom}(n, 1 - \frac{1}{n})$, $N(2, 1) \sim \text{Binom}(n, 0)$, and $N(2, 2) \sim \text{Binom}(n, \frac{1}{n})$. With this definition of N , we get the $\hat{P}$ displayed in Figure 1 if we have $N(1, 1) = n$, $N(1, 2) = n - 1$, $N(2, 1) = 0$, $N(2, 2) = 1$. By independence and the Binomial pmf we have

$$\begin{aligned}
& \mathbb{P}(N(1, 1) = n, N(1, 2) = n - 1, N(2, 1) = 0, N(2, 2) = 1) \\
& = \mathbb{P}(N(1, 1) = n) \mathbb{P}(N(1, 2) = n - 1) \mathbb{P}(N(2, 1) = 0) \mathbb{P}(N(2, 2) = 1) \\
& = \left(1 - \frac{1}{n}\right)^n \cdot n \frac{1}{n} \left(1 - \frac{1}{n}\right)^{n-1} \cdot 1 \cdot n \frac{1}{n} \left(1 - \frac{1}{n}\right)^{n-1} \\
& = \left(1 - \frac{1}{n}\right)^{3n-2} \\
& \geq \left(\left(1 - \frac{1}{n}\right)^n\right)^3
\end{aligned}$$

and it is a standard fact that this final expression is increasing in n , so we can lower bound it by plugging in the lowest value $n = 10$ for which we obtain $\left(\frac{9}{10}\right)^{30} > 0.04 = \frac{1}{25}$. (As $n \rightarrow \infty$ this approaches $\frac{1}{e^3}$.)

From here we operate on this event that $\hat{P}$ is equal to the instance shown in Figure 1. It is easy to see that $\hat{P}$ is communicating, since in both states action 2 has positive probability of leading to either state.

First we compute the Blackwell optimal policy $\hat{\pi}^*$ of $\hat{P}$. It is easy to see that

$$\hat{\rho}^{(1,1)} = \begin{bmatrix} \frac{1}{2} + \frac{1}{n} \\ \frac{1}{2} \end{bmatrix}, \hat{\rho}^{(1,2)} = \begin{bmatrix} \frac{1}{2} + \frac{1}{n} \\ \frac{1}{2} + \frac{1}{n} \end{bmatrix}, \hat{\rho}^{(2,1)} = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}, \hat{\rho}^{(2,2)} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

Therefore the only Blackwell optimal policy is the only gain-optimal policy, $(1, 2)$. As we have already checked, this policy has suboptimality (in the true P)

$$\|\rho^{(1,2)} - \rho^*\|_\infty = \frac{1}{2} - \left(\frac{1}{4} + \frac{1}{2n}\right) \geq \frac{1}{4} - \frac{1}{20} = \frac{1}{5}.$$

Next we compute the discounted optimal policy for effective horizon $\frac{1}{1-\gamma} = n^2$. It is obvious that the optimal action in state 1 will be action 1, so we will compute and compare the value functions $\hat{V}^{(1,1)}$ and $\hat{V}^{(1,2)}$. It is easy to see that

$$\hat{V}^{(1,1)} = \frac{1}{1-\gamma} \left[\frac{1}{2} + \frac{1}{n} \right]$$

since both states are absorbing under this policy. Now we compute $\hat{V}^{(1,2)}$. First, since state 1 is absorbing, it is immediate that $\hat{V}^{(1,2)}(1) = \frac{1}{1-\gamma} \left(\frac{1}{2} + \frac{1}{n} \right) = n^2 \left(\frac{1}{2} + \frac{1}{n} \right)$. From the Bellman evaluation equation for state 2 we have that

$$\begin{aligned} \hat{V}^{(1,2)}(2) &= 0 + \gamma \frac{1}{n} \hat{V}^{(1,2)}(1) + \gamma \left(1 - \frac{1}{n}\right) \hat{V}^{(1,2)}(2) = \gamma n \left(\frac{1}{2} + \frac{1}{n}\right) + \gamma \left(1 - \frac{1}{n}\right) \hat{V}^{(1,2)}(2) \\ \iff \left(1 - \gamma \left(1 - \frac{1}{n}\right)\right) \hat{V}^{(1,2)}(2) &= \gamma n \left(\frac{1}{2} + \frac{1}{n}\right) \\ \iff \left(1 - \left(1 - \frac{1}{n^2}\right) \left(1 - \frac{1}{n}\right)\right) \hat{V}^{(1,2)}(2) &= \left(1 - \frac{1}{n^2}\right) n \left(\frac{1}{2} + \frac{1}{n}\right) \\ \iff \frac{n^2 + n - 1}{n^3} \hat{V}^{(1,2)}(2) &= \left(1 - \frac{1}{n^2}\right) n \left(\frac{1}{2} + \frac{1}{n}\right) = \frac{n}{2} + 1 - \frac{1}{2n} - \frac{1}{n^2} = \frac{\frac{n^3}{2} + n^2 - \frac{n}{2} - 1}{n^2} \\ \iff \hat{V}^{(1,2)}(2) &= \frac{n^3}{n^2 + n - 1} \frac{\frac{n^3}{2} + n^2 - \frac{n}{2} - 1}{n^2} = \frac{\frac{n^4}{2} + n^3 - \frac{n^2}{2} - n}{n^2 + n - 1} > \frac{\frac{n^4}{2} + \frac{n^3}{2} - \frac{n^2}{2}}{n^2 + n - 1} = \frac{n^2}{2} \end{aligned}$$

where the final strict inequality requires $\frac{n^3}{2} - n > 0$, which holds for all $n \geq 2$. Thus we have shown that $\hat{V}^{(1,2)}(2) > \frac{n^2}{2} = \frac{1}{1-\gamma} \frac{1}{2} = \hat{V}^{(1,1)}(2)$, so the optimal policy for the DMDP with horizon $\frac{1}{1-\gamma}$ is $(1, 2)$. As we have previously checked, this policy has suboptimality in P at least $\frac{n^2}{5}$.

Now it remains to check the final two statements of the theorem. Fix a constant $C > 0$. Then since the term $\sqrt{\frac{\log(n)}{n}}$ goes to 0 as $n \rightarrow \infty$, we can choose n sufficiently large so that $C\sqrt{\frac{\log(n)}{n}} < \frac{1}{5}$. Then considering the instance P constructed with the parameter n , since it has $\|h^*\|_{\text{span}} = 1$, we have

$$C\sqrt{\frac{\log(n)}{n}} = C\sqrt{\frac{\|h^*\|_{\text{span}} \log(\|h^*\|_{\text{span}} n)}{n}} < \frac{1}{5}.$$

As we have argued, there is probability at least $\frac{1}{25}$ that $\hat{P}$ is sampled from P , and under this event, we have both $\|\rho^* - \rho^{\hat{\pi}^*}\|_\infty \geq \frac{1}{5}$ and $\|V_\gamma^* - V_\gamma^{\hat{\pi}^*}\|_\infty \geq \frac{n^2}{5} = \frac{1}{1-\gamma} \frac{1}{5}$ (where we choose $\gamma = 1 - \frac{1}{n^2}$). Therefore the statements

$$\begin{aligned} \mathbb{P}\left(\rho^{\hat{\pi}^*} \geq \rho^* - C\sqrt{\frac{\|h^*\|_{\text{span}} \log(\|h^*\|_{\text{span}} n)}{n}}\right) &> 1 - \frac{1}{25} \\ \mathbb{P}\left(V_\gamma^{\hat{\pi}^*} \geq V_\gamma^* - C\frac{1}{1-\gamma} \sqrt{\frac{\|h^*\|_{\text{span}} \log(\|h^*\|_{\text{span}} n)}{n}}\right) &> 1 - \frac{1}{25} \end{aligned}$$

are both false. □