

Sample Complexity Bounds for Linear Constrained MDPs with a Generative Model

Xingtu Liu

Simon Fraser University

RLTHEORY@OUTLOOK.COM

Lin F. Yang

University of California, Los Angeles

LINYANG@EE.UCLA.EDU

Sharan Vaswani

Simon Fraser University

VASWANI.SHARAN@GMAIL.COM

Editors: Matus Telgarsky and Jonathan Ullman

Abstract

We consider infinite-horizon γ -discounted (linear) constrained Markov decision processes (CMDPs) where the objective is to find a policy that maximizes the expected cumulative reward subject to expected cumulative constraints. Given access to a generative model, we propose to solve CMDPs with a primal-dual framework that can leverage any black-box unconstrained MDP solver. For linear CMDPs with feature dimension d , we instantiate the framework by using mirror descent value iteration (MDVI) (Kitamura et al., 2023) an example MDP solver. We provide sample complexity bounds for the resulting CMDP algorithm in two cases: (i) relaxed feasibility, where small constraint violations are allowed, and (ii) strict feasibility, where the output policy is required to exactly satisfy the constraint. For (i), we prove that the algorithm can return an ε -optimal policy with high probability by using $\tilde{O}\left(\frac{d^2}{(1-\gamma)^4\varepsilon^2}\right)$ samples. For (ii), we show that the algorithm requires $\tilde{O}\left(\frac{d^2}{(1-\gamma)^6\varepsilon^2\zeta^2}\right)$ samples, where ζ is the problem-dependent Slater constant that characterizes the size of the feasible region. Furthermore, we prove a lower-bound of $\Omega\left(\frac{d^2}{(1-\gamma)^5\varepsilon^2\zeta^2}\right)$ for the strict feasibility setting. We note that our upper bounds under both settings exhibit a near-optimal dependence on d , ε , and ζ . Finally, we instantiate our framework for tabular CMDPs and show that it can be used to recover near-optimal sample complexities in this setting.

Keywords: Reinforcement Learning, Sample Complexity, Constrained MDPs, Linear Function Approximation

1. Introduction

Reinforcement learning (RL) (Sutton et al., 1998) is a machine learning paradigm aimed at building learning agents capable of making sequential decisions in an (unknown) environment. RL algorithms have found applications in games such as Atari (Mnih et al., 2015) or Go (Silver et al., 2016), robot manipulation tasks (Tan et al., 2018; Zeng et al., 2020), clinical trials (Schaefer et al., 2005) and more recently, aligning large language models to human preferences (Shao et al., 2024; Ouyang et al., 2022). Typical RL algorithms only focus on optimizing an unconstrained objective, although in many real-world applications, agents are often required to not only maximize cumulative rewards but also to satisfy constraints imposed by safety, fairness, or resource usage. RL with such side-constraints is typically formulated within the framework of constrained Markov decision processes (CMDPs) (Altman, 1999), where the goal is to optimize an expected reward function while ensuring that the expected cumulative cost (or utility) satisfies a given threshold. For example, in wireless

sensor networks (Buratti et al., 2009; Julian et al., 2002), the agent aims to deploy a policy that maximizes the bitrate with a constraint on its average power consumption.

Given the practical importance of constrained RL, there is a vast literature (Efroni et al., 2020; Zheng and Ratliff, 2020; Qiu et al., 2020; Brantley et al., 2020; Kalagarla et al., 2021; Yu et al., 2021; Ding et al., 2021; Gattami et al., 2021; Miryoosefi and Jin, 2022; Mondal and Aggarwal, 2024) that aims to obtain a near-optimal policy in unknown tabular CMDPs with finite states and actions. These works simultaneously tackle the exploration, estimation and planning problems and aim to minimize the regret and constraint violation in the online setting. On the other hand, recent works (HasanzadeZonuzuy et al., 2021; Wei et al., 2021; Bai et al., 2021; Vaswani et al., 2022) consider an easier, but even more fundamental problem of obtaining a near-optimal policy with access to a simulator or *generative model* (Kearns and Singh, 1999; Kakade, 2003; Agarwal et al., 2020; Sidford et al., 2018; Yang and Wang, 2019). In particular, these works assume that the agent has access to a sampling oracle (the generative model) that returns a sample of the next state when given any state-action pair as input. Depending on the application of interest, such a generative model is often available either directly for the task at hand (for example, in Atari games where the aim is to win the game) or as an proxy to the task (for example, the CARLA simulator (Dosovitskiy et al., 2017) for training autonomous vehicles). Moreover, from a theoretical perspective, since the generative model setting removes the need for exploration it has been used to characterize the statistical complexity of obtaining near-optimal policies for (C)MDPs (Azar et al., 2013; Agarwal et al., 2020; Li et al., 2020; Vaswani et al., 2022). In particular, for CMDPs, Vaswani et al. (2022) established near-optimal upper and lower-bounds on the sample complexity in two settings: (i) relaxed feasibility, where small constraint violations are allowed, and (ii) strict feasibility, where the output policy is required to exactly satisfy the constraint. For tabular CMDPs, the proposed algorithms and resulting bounds depend on the cardinality of the state-action space, and hence do not apply to modern applications involving large or infinite state spaces. Consequently, it is essential to develop provably efficient algorithms that can incorporate function approximation and go beyond the tabular case.

For unconstrained MDPs, the linear MDP assumption (e.g., (Yang and Wang, 2019; Jin et al., 2020)) is a common formalization to analyze algorithms that have access to state-action features and can incorporate linear function approximation. The assumption implies that both the rewards and transition probabilities (approximately) lie in the span of the given d -dimensional feature representation, and can be used to obtain sample complexity bounds independent of the size of the state-action space. Unconstrained linear MDPs have been extensively studied in the context of both finite-horizon regret minimization (Jin et al., 2020; Hu et al., 2022; Weisz et al., 2022; Sherman et al., 2023; Liu et al., 2023) and with access to a generative model (Kitamura et al., 2023; Taupin et al., 2023). Following the linear MDP literature, recent works consider CMDPs with linear function approximation (Jain et al., 2022; Ding et al., 2021; Miryoosefi and Jin, 2022; Ghosh et al., 2022, 2024; Liu et al., 2022; Tian et al., 2024) and assume that (in addition to the rewards and transition probabilities), the costs or utilities can also be expressed using the given features. However, all previous work on linear CMDPs considers the online regret minimization setting and the statistical complexity of the problem remains unclear. Motivated by Vaswani et al. (2022), we aim to *study the sample complexity of solving linear CMDPs with access to a generative model*. In particular, we make the following contributions.

(1) Generic primal-dual algorithm framework: In Sec. 3, we provide a generic primal-dual algorithmic framework (Alg. 1) that can be used to achieve both the *relaxed* and *strict* feasibility objectives, for both *tabular* and *linear* CMDPs. As model-based approaches (Vaswani et al., 2022)

are not applicable in the linear CMDP setting, Alg. 1 is designed to be model-free and relies on three black-box subroutines: a `DataCollection` procedure, a black-box `MDP-Solver` and a `PolicyEvaluation` oracle. We prove a meta-theorem (Thm. 4) to quantify the sample complexity of Alg. 1 in terms of that of the `MDP-Solver` and `PolicyEvaluation` oracle.

(2) Instantiating the framework for linear CMDPs: In Sec. 4.2, we instantiate the linear `MDP-Solver` with a variant of the mirror-descent value iteration (MDVI) algorithm (Kozuno et al., 2022; Kitamura et al., 2023). In contrast to the existing MDVI variants, the proposed Alg. 2 does not use entropy regularization and outputs a stationary policy, thus simplifying the algorithm design. We develop a new theoretical analysis for Alg. 2 and characterize its sample complexity for solving unconstrained linear MDPs. In Sec. 4.3, we instantiate the `PolicyEvaluation` oracle with least-squares policy evaluation (Alg. 3) and analyze the sample complexity required to evaluate the performance of a (data-dependent) policy.

(3) Upper-bound on sample complexity for linear CMDPs: In Sec. 4.5, we leverage our meta-theorem and analyze the sample complexity for the resulting CMDP algorithm that uses Algs. 2 and 3. In particular, if d is the dimension of the feature mapping, we prove that the proposed algorithm requires no more than $\tilde{O}\left(\frac{d^2}{(1-\gamma)^4\varepsilon^2}\right)$ samples to obtain an ε -optimal policy in the relaxed feasibility setting. Since the lower-bound on the sample complexity for solving unconstrained linear MDP is $\Omega\left(\frac{d^2}{(1-\gamma)^3\varepsilon^2}\right)$ (Weisz et al., 2022), our sample complexity achieves the near-optimal dependence on d and ε , and is away from the lower bound by atmost a multiplicative factor of $\tilde{O}(1/1-\gamma)$. Under strict feasibility, our algorithm requires no more than $\tilde{O}\left(\frac{d^2}{(1-\gamma)^6\varepsilon^2\zeta^2}\right)$ samples, where ζ is the problem-dependent Slater constant that characterizes the size of the feasible region and dictates the difficulty of the problem. Given the lower-bounds for tabular CMDPs in Vaswani et al. (2022), we conjecture that the dependence on d , ε , and ζ in our bounds is tight, with suboptimality arising only in the multiplicative dependence on $O(1/1-\gamma)$. To the best of our knowledge, *these are the first such sample complexity bounds with the near-optimal dependence on both d and ε* . In App. D.5, we alternatively instantiate the linear `MDP-Solver` to be the `G-Sampling-and-Stop` (GSS) algorithm (Taupin et al., 2023) and analyze the sample complexity of the resulting CMDP algorithm, thus demonstrating the flexibility of our framework.

(4) Lower-bound on sample complexity for linear CMDPs: In Sec. 4.5, we prove a problem-dependent $\Omega\left(\frac{d^2}{(1-\gamma)^5\varepsilon^2\zeta^2}\right)$ lower-bound on the sample-complexity in the strict feasibility setting. Our results thus demonstrate that the proposed algorithm is near-optimal in terms of d , ε , and ζ , with a suboptimality only in the multiplicative dependence on H . The lower bound also indicates that under strict feasibility solving linear CMDPs is inherently more difficult than solving unconstrained linear MDPs, and the problem difficulty increases as the size of the feasible region (measured in terms of ζ) decreases. To the best of our knowledge, it is the first result characterizing the difficulty of solving linear CMDPs with access to a generative model.

(5) Sample complexity bounds for tabular CMDPs: Finally, in App. G, we utilize our framework for tabular CMDPs. In particular, we instantiate Alg. 1 with tabular variants of Algs. 2 and 3 (obtained by setting $d = SA$ and considering one-hot features) and analyze the resulting CMDP algorithm. Under the relaxed and strict feasibility settings, the resulting algorithm attains sample complexity bounds of $\tilde{O}\left(\frac{|S||A|}{(1-\gamma)^3\varepsilon^2}\right)$ and $\tilde{O}\left(\frac{|S||A|}{(1-\gamma)^5\varepsilon^2\zeta^2}\right)$, respectively. These results match the near-optimal bounds attained by the model-based algorithm in Vaswani et al. (2022), and improve upon the sample-complexity of the model-free approach proposed in Bai et al. (2021).

2. Problem Formulation

An infinite-horizon discounted constrained tabular Markov decision process (CMDP) (Altman, 1999) is denoted by \mathcal{M} , and is defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, r, c, b, \rho, \gamma \rangle$ where \mathcal{S} is the set of states, \mathcal{A} is the action set, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta_{\mathcal{S}}$ is the transition probability function, $\rho \in \Delta_{\mathcal{S}}$ is the initial distribution of states and $\gamma \in [0, 1)$ is the discount factor. The primary reward to be maximized is denoted by $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, whereas the constraint reward is denoted by $c : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ ¹. If $\Delta_{\mathcal{A}}$ denotes the simplex over the action space, the expected discounted return or *reward value function* of a stationary, stochastic policy² $\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$ is defined as $V_r^\pi(\rho) = \mathbb{E}_{s_0, a_0, \dots} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right]$, where $s_0 \sim \rho, a_t \sim \pi(\cdot | s_t)$, and $s_{t+1} \sim \mathcal{P}(\cdot | s_t, a_t)$. For each state-action pair (s, a) and policy π , the reward action-value function is defined as $Q_r^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, and satisfies the relation: $V_r^\pi(s) = \langle \pi(\cdot | s), Q_r^\pi(s, \cdot) \rangle$, where $V_r^\pi(s)$ is the reward value function when the starting state is equal to s . Analogously, the *constraint value function* and constraint action-value function of policy π is denoted by $V_c^\pi(\rho)$ and Q_c^π respectively. Throughout, it will be convenient to present our results in terms of the effective horizon $H := 1/(1-\gamma)$.

In addition to the tabular CMDPs with a finite state-action space, we also consider linear (Jin et al., 2020) CMDPs where the state space can be large or possibly infinite. In this case, we assume access to a feature representation ϕ such that r, c and the transition probabilities \mathcal{P} (approximately) lie in the span of the given d -dimensional feature representation.

Assumption 1 (Linear Constrained MDP) *For the CMDP \mathcal{M} with the state-action space $\mathcal{S} \times \mathcal{A}$, we have access to a known feature map $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ that satisfies the following condition: there exist vectors $\psi_r, \psi_c \in \mathbb{R}^d$ and signed measures $\mu := (\mu_1, \dots, \mu_d)$ on \mathcal{S} such that $P(\cdot | s, a) = \langle \phi(s, a), \mu \rangle$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, $r(s, a) = \langle \phi(s, a), \psi_r \rangle$, and $c(s, a) = \langle \phi(s, a), \psi_c \rangle$. Let $\Phi := \{\phi(s, a) : (s, a) \in \mathcal{S} \times \mathcal{A}\} \subset \mathbb{R}^d$ be the set of all feature vectors. We assume that Φ is compact and spans \mathbb{R}^d .*

The objective is to return a policy that maximizes $V_r^\pi(\rho)$, while ensuring that $V_c^\pi(\rho) \geq b$. Formally,

$$\max_{\pi} V_r^\pi(\rho) \quad \text{s.t.} \quad V_c^\pi(\rho) \geq b. \quad (1)$$

The optimal stochastic policy for the above CMDP is denoted by π^* and the corresponding reward value function is denoted by $V_r^*(\rho)$. We also define $\zeta := \max_{\pi} V_c^\pi(\rho) - b > 0$ as the problem-dependent quantity referred to as the Slater constant (Ding et al., 2021; Bai et al., 2021). The Slater constant is a measure of the size of the feasible region and determines the difficulty of solving Eq. (1).

For simplicity of exposition, we assume that the rewards r and constraint rewards c are known, but the transition probabilities \mathcal{P} are unknown. We note that assuming the knowledge of the rewards does not affect the leading terms of the sample complexity since learning these is an easier problem (Azar et al., 2013; Sidford et al., 2018). Following Azar et al. (2013); Vaswani et al. (2022), we assume access to a *generative model* or simulator that allows the agent to obtain samples from the $\mathcal{P}(\cdot | s, a)$ distribution for any (s, a) .

Definition 1 (Generative Model) *A generative model Gen for an MDP is an oracle that, given any state-action pair (s, a) , returns an independent sample of the next state $s' \sim P(\cdot | s, a)$.*

1. These ranges for r and c are chosen for simplicity. Our results can be easily extended to handle other ranges.
 2. The performance of an optimal policy in a CMDP can always be achieved by a stationary, stochastic policy (Altman, 1999). On the other hand, for an MDP, it suffices to only consider stationary, deterministic policies (Puterman, 2014).

Assuming access to such a generative model, we aim to characterize the sample complexity (number of times Gen is queried) required to return a near-optimal policy $\bar{\pi}$. Specifically, given a target error $\varepsilon > 0$, we consider two different definitions of optimality.

Relaxed feasibility: We require $\bar{\pi}$ to achieve an approximately optimal reward value, while allowing it to have a small constraint violation. Formally, we aim to find a $\bar{\pi}$ such that,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - \varepsilon \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b - \varepsilon. \quad (2)$$

Strict feasibility: We require $\bar{\pi}$ to achieve an approximately optimal reward value, while simultaneously demanding zero constraint violation. Formally, we aim to find a $\bar{\pi}$ such that,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - \varepsilon \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b. \quad (3)$$

In the next section, we design a generic algorithmic framework to achieves these objectives.

3. A Generic Framework for Solving CMDPs

We first present a generic primal-dual algorithmic framework for solving CMDPs, and subsequently present a meta-theorem that quantifies its sample-complexity in the relaxed and strict feasibility settings. For this, we frame the CMDP problem in Eq. (1) as an equivalent saddle-point problem,

$$\max_{\pi} \min_{\lambda \geq 0} [V_r^{\pi}(\rho) + \lambda (V_c^{\pi}(\rho) - b)], \quad (4)$$

where, λ is the Lagrange multiplier. The solution to Eq. (4) is (π^*, λ^*) where π^* is the optimal policy to the CMDP and λ^* is the optimal Lagrange multiplier. We solve Eq. (4) iteratively, by alternatively updating the policy (primal variable) and the Lagrange multiplier (dual variable) (Ding et al., 2021; Vaswani et al., 2022).

Algorithm 1 Primal-dual CMDP framework with a generative model

Input: r (rewards), c (constraint rewards), b' (constraint RHS), U (projection upper bound), K (number of iterations), η (step-size), $\lambda_0 = 0$ (initialization), Gen (generative model), \mathcal{C} (subset of $\mathcal{S} \times \mathcal{A}$), N (sample size for each (s, a) pair in \mathcal{C}), ϕ (feature map).

Output: Mixture policy $\bar{\pi} = \frac{1}{K} \sum_{k=0}^{K-1} \pi_k$.

procedure $CMDPF(r, c, b', U, K, \eta, \text{Gen}, \mathcal{C}, N, \phi)$

Data collection procedure to populate buffer: $\mathcal{B} = \text{DataCollection}(\text{Gen}, \mathcal{C}, N)$.

for $k = 0, \dots, K - 1$ **do**

Updating the primal variable: Let $\pi_k = \text{MDP-Solver}(r + \lambda_k c, \mathcal{B}, \phi)$.

Policy Evaluation: Let $\hat{V}_c^k = \text{PolicyEvaluation}(\pi_k, c, \mathcal{B}, \phi)$.

Updating the dual variable: $\lambda_{k+1} = \mathbb{P}_{[0, U]} [\lambda_k - \eta (\hat{V}_c^k(\rho) - b')]$.

end

end

Remark 2 *The model-based framework for solving tabular CMDPs in Vaswani et al. (2022) does not directly extend to the linear MDP setting for the following reasons. In the PAC-RL framework (with access to a simulator), model-based methods build a model for the transition matrix by sampling next states for each state-action pair. In the linear (C)MDP setting, the transition kernel and reward*

functions are assumed to be approximated by a linear function. Specifically, the $SA \times S$ transition matrix can be factorized into a known $SA \times d$ matrix of features and an unknown $d \times S$ matrix to be estimated. Hence, naively estimating the transition matrix will require $O(S)$ samples, resulting in high sample complexity.

The primal and dual updates in Alg. 1 rely on three oracles, which we instantiate subsequently.

Data Collection Oracle: We first describe the mechanism of the `DataCollection` oracle in Alg. 1. This oracle takes as input a generative model `Gen`, a subset of state-action pairs $\mathcal{C} \subseteq \mathcal{S} \times \mathcal{A}$, and a sample size N . For each $(s, a) \in \mathcal{C}$, it queries the generative model `Gen` to obtain N independent next-state samples $(s'_i)_{i=1}^N$ from the distribution $\text{Gen}(\cdot | s, a)$. It then stores the resulting triplets $(s, a, s'_i)_{i=1}^N$ in a buffer \mathcal{B} . After all state-action pairs in \mathcal{C} are processed, the buffer \mathcal{B} contains N samples for each pair and is returned as the output.

MDP-Solver: The primal update at iteration k uses the `MDP-Solver`, which takes as input a buffer \mathcal{B} of samples and returns a policy π_k satisfying the following assumption.

Assumption 2 *We have access to a black-box algorithm `MDP-Solver`($\square, \mathcal{B}, \phi$) for which the inputs are the feature map ϕ and an arbitrary but bounded reward function $\square \in [0, R]$, and the output is a policy $\tilde{\pi}$ satisfying the following condition with probability $1 - \delta$,*

$$\max_{\pi} V_{\square}^{\pi}(\rho) - V_{\square}^{\tilde{\pi}}(\rho) \leq R f_{\text{mdp}}(\mathcal{B}),$$

where, $f_{\text{mdp}}(\mathcal{B})$ denotes an upper bound on the sub-optimality when given access to buffer \mathcal{B} .

Policy Evaluation Oracle: The dual update at iteration k in Alg. 1 is given as:

$$\lambda_{k+1} = \mathbb{P}_{[0, U]} \left[\lambda_k - \eta (\hat{V}_c^k(\rho) - b') \right],$$

where $\mathbb{P}_{[0, U]}$ denotes the projection onto the interval $[0, U]$, and b' is a relaxed constraint parameter that depends on b , f_{mdp} and the problem setting (relaxed or strict). The term \hat{V}_c^k is an estimate of $V_c^{\pi_k}$, computed via the `PolicyEvaluation` oracle which satisfies following assumption.

Assumption 3 *We have access to a black-box algorithm `PolicyEvaluation`($\pi, \diamond, \mathcal{B}, \phi$) for which the inputs are a possibly data-dependent (one that depends on the buffer \mathcal{B}) policy π , the feature map ϕ , a reward function $\diamond \in [0, 1]$, and the output is a value function \hat{V}_{\diamond} satisfying the following condition with probability $1 - \delta$,*

$$|\hat{V}_{\diamond}(\rho) - V_{\diamond}^{\pi}(\rho)| \leq f_{\text{eva}}(\mathcal{B}),$$

where, $f_{\text{eva}}(\mathcal{B})$ denotes an upper bound on the sub-optimality when given access to buffer \mathcal{B} .

Remark 3 *We note that while there are existing methods that update the dual variable using empirical utility value functions obtained from the MDP solver, this approach does not result in near-optimal statistical guarantees. This is because the MDP solver can only imply a concentration guarantee on the value function for the combined reward function $r + \lambda_k c$ and not for the individual value functions V_r and V_c . Proving near-optimal statistical guarantees necessitates the development of the proposed policy evaluation subroutine. Alternative ways to handle such concentration include using a uniform concentration bound (Ghosh et al., 2022) (i.e., building an ε -cover over the function*

class and using a union bound). However, such an analysis results in a sub-optimal $O(d^3)$ dependence on the dimension. Furthermore, in order to guarantee concentration, Ghosh et al. (2022) needs to guarantee Lipschitzness in the policy update. Our technique does not require this, and consequently, we can use greedy policy updates that are not Lipschitz.

After K iterations of primal and dual updates, Alg. 1 returns a mixture policy $\bar{\pi}$ which is a policy drawn uniformly at random from the set $\{\pi_0, \dots, \pi_{K-1}\}$. Given access to these oracles, we state a meta-theorem (proved in App. C) to characterize the sub-optimality of the algorithm.

Theorem 4 *Suppose Assumptions 2 and 3 hold and let $f(\mathcal{B}) := \max\{f_{\text{mdp}}(\mathcal{B}), f_{\text{eva}}(\mathcal{B})\}$. For $\delta \in (0, 1)$, Alg. 1 with $U = \frac{2}{\zeta(1-\gamma)}$, $\eta = \frac{U(1-\gamma)}{\sqrt{K}}$, $K = \frac{U^2}{[f(\mathcal{B})]^2(1-\gamma)^2}$ and $b' = b - 2f(\mathcal{B})$, returns a mixture policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,*

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - 4f(\mathcal{B}) \quad , \quad V_c^{\bar{\pi}}(\rho) \geq b - 6f(\mathcal{B}). \quad (\text{Relaxed Feasibility Setting})$$

With the same algorithm parameters, but with $b' = b + 4f(\mathcal{B})$ for $f(\mathcal{B}) \leq \frac{\zeta}{6}$, Alg. 1 returns a mixture policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - \frac{16f(\mathcal{B})}{\zeta(1-\gamma)} \quad , \quad V_c^{\bar{\pi}}(\rho) \geq b. \quad (\text{Strict Feasibility Setting})$$

The above theorem implies that, provided we can adequately control the terms $f_{\text{mdp}}(\mathcal{B})$ and $f_{\text{eva}}(\mathcal{B})$ via the three oracle procedures, both the relaxed feasibility condition (2) and the strict feasibility conditions (3) can be satisfied. Furthermore, we note that similar to (Vaswani et al., 2022), the error for the strict feasibility setting is inflated by an $O\left(\frac{1}{\zeta(1-\gamma)}\right)$ factor.

Hence, in the next section, we instantiate the subroutines `DataCollection`, `MDP-Solver` and `PolicyEvaluation` such that the quantities $f_{\text{mdp}}(\mathcal{B})$ and $f_{\text{eva}}(\mathcal{B})$ are sufficiently small.

4. Instantiating the Framework for Linear Constrained MDPs

We first describe the construction of the coreset \mathcal{C} , which serves as input to the `DataCollection` procedure. We then introduce a model-free algorithm, `LS-MDVI`, as an instantiation of the `MDP-Solver`. Finally, we present `LS-PE`, which serves as the instantiation of the `PolicyEvaluation` subroutine.

4.1. Data Collection via Core Set Construction

Recall that the `DataCollection` procedure requires as input a subset of $\mathcal{S} \times \mathcal{A}$. In the linear setting, we provide a coreset \mathcal{C} as this input. We now describe the construction of the coreset (Lattimore et al., 2020; Kitamura et al., 2023). The key properties of the coreset are that it has few elements (independent of the cardinality of \mathcal{S} and \mathcal{A}), while the features corresponding to the $(x, b) \in \mathcal{C}$ provide a good coverage of the feature space. For a distribution $\tilde{\rho}$ over $\mathcal{S} \times \mathcal{A}$, let $G \in \mathbb{R}^{d \times d}$ and $g(\tilde{\rho}) \in \mathbb{R}$ be defined as:

$$G := \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x, b) \phi(x, b) \phi(x, b)^\top \quad \text{and} \quad g(\tilde{\rho}) := \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \langle \phi(s, a), G^{-1} \phi(s, a) \rangle.$$

We refer to $\tilde{\rho}$ as the design, G as the corresponding design matrix, and define the coreset of $\tilde{\rho}$ as its support, $\mathcal{C} := \text{Supp}(\tilde{\rho})$. The task of identifying a design that minimizes g is known as the G -optimal design problem. We assume that we can construct near-optimal experimental design.

Assumption 4 (Optimal Design) We have access to an oracle called `ComputeOptimalDesign` which returns $\tilde{\rho}$, \mathcal{C} and G such that $g(\tilde{\rho}) \leq 2d$ and the coreset of $\tilde{\rho}$ has size at most $\tilde{O}(d)$.

Such a design can be obtained using the Frank-Wolfe algorithm (Todd, 2016) described in App. A. Accordingly, we first use the `ComputeOptimalDesign` procedure to construct $\tilde{\rho}$, \mathcal{C} , and the associated design matrix G , and then utilize the resulting coreset \mathcal{C} to collect data. For each state-action pair in \mathcal{C} , we collect N independent samples and store them in the buffer \mathcal{B} . Hence, the total sample complexity is $N|\mathcal{C}|$. In the subsequent section, it is convenient to consider \mathcal{B} as a union of T disjoint subsets $B_0 \cup \dots \cup B_{T-1}$, where each B_i consists of M independent samples for every state-action pair in \mathcal{C} . Consequently, we have $N = TM$.

4.2. Instantiating the MDP-Solver: Least-Squares Mirror Descent Value Iteration

We now introduce a model-free algorithm referred to as least-squares mirror descent value iteration (LS-MDVI) which serves as an instantiation of the MDP-Solver.

LS-MDVI is a generalization of MDVI (Geist et al., 2019; Vieillard et al., 2020; Kozuno et al., 2022) to the linear function approximation setting and is related to the algorithm proposed in Kitamura et al. (2023). In particular, LS-MDVI corresponds to a limiting case of policy mirror descent (Lan, 2023) when the KL regularization tends to zero (or equivalently, the step-size tends to infinity). This results in a value iteration method which we describe below.

Define $\mathcal{H}(\pi(\cdot|s))$ as the entropy of the policy π in state s and $\text{KL}(\pi(\cdot|s)||\pi'(\cdot|s))$ as the KL divergence between policies $\pi(\cdot|s)$ and $\pi'(\cdot|s)$ in state s . With a slight abuse of notation, we consider π to be an operator such that $(\pi Q)(s) := \sum_{a \in \mathcal{A}} \pi(a|s)Q(s, a)$. At iteration $t \in [T]$, LS-MDVI requires the corresponding action-value function to update the policy. Specifically, if τ is the strength of the KL regularization and κ is the entropy regularization coefficient s.t. $\alpha = \frac{\tau}{\tau+\kappa}$, $\beta = \frac{1}{\tau+\kappa}$, given Q^{t+1} for some reward function, the entropic mirror descent and LS-MDVI updates can be written as:

Entropic Mirror Descent : $\pi_{t+1}(a|s) \propto [\pi_t(a|s)]^\alpha \exp(\beta Q^{t+1}(s, a))$,

$$V^{t+1}(s) = (\pi_{t+1} Q^{t+1})(s) - \tau \text{KL}(\pi_{t+1}(\cdot|s)||\pi_t(\cdot|s)) + \kappa \mathcal{H}(\pi_{t+1}(\cdot|s)).$$

$$\text{LS-MDVI} : \pi_{t+1}(\cdot|s) = \arg \max_a \sum_{i=0}^{t+1} Q^i(s, a); V^{t+1}(s) = \left(\pi_{t+1} \sum_{i=0}^{t+1} Q^i \right)(s) - \left(\pi_t \sum_{i=0}^t Q^i \right)(s).$$

Starting from entropic mirror descent, for $\kappa = 0$ and as $\tau \rightarrow 0$, implying $\alpha = 1$, we recover the LS-MDVI update (see (Kozuno et al., 2022, App. B) for the derivation).

Remark 5 In contrast, Kitamura et al. (2023) consider both $\kappa \rightarrow 0$, $\tau \rightarrow 0$ while keeping α fixed and effectively consider an entropy-regularized update. This proposed change simplifies the algorithm design for LS-MDVI. Furthermore, while the algorithm in Kitamura et al. (2023) produces non-stationary policies, LS-MDVI outputs a stationary policy.

Next, we present Alg. 2 which implements the above LS-MDVI update, but uses the linear CMDP structure and the data collected in the buffer \mathcal{B} to estimate Q^{t+1} . Specifically, θ^{t+1} estimation in Alg. 2 corresponds to the Q^{t+1} estimation using linear regression and the last line in Alg. 2 corresponds to the above update. Similar to approximate value iteration, the \hat{Q}^{t+1} update depends

on \hat{V}^t via the Bellman equation, however, π_{t+1} depends on \tilde{Q}^{t+1} , the ‘‘soft’’ Q function formed by using the estimates up to iteration $t + 1$.

Algorithm 2 Least-Squares Mirror Descent Value Iteration (LS-MDVI)

Input: T (number of iterations), M (number of next-state samples obtained per state-action pair in each iteration), \square (rewards in MDP), $\mathcal{B} = \mathcal{B}_0 \cup \dots \cup \mathcal{B}_{T-1}$ (Buffer), $\tilde{\rho}$ (design), \mathcal{C} (coreset), ϕ (feature map).

Output: π_T where $\forall s \in \mathcal{S}, \pi_T(\cdot|s) \in \arg \max_a \tilde{Q}_\square^T(s, a)$.

Initialize: $\hat{V}_\square^0 = \mathbf{0}, \tilde{\theta}_\square^0 = \mathbf{0}$.

procedure *LS-MDVI*($T, M, \square, \mathcal{B}, \mathcal{C}, \phi$)

```

for  $t = 0, \dots, T - 1$  do
     $\forall (s, a) \in \mathcal{C}$  : Access  $(s, a, s'_m)_{m=1}^M$  from the buffer  $\mathcal{B}_t$ .
    Define regression target  $\hat{Q}_\square^{t+1}(s, a) := \square(s, a) + \gamma \frac{1}{M} \sum_{m=1}^M \hat{V}_\square^t(s'_m)$ .
     $\theta_\square^{t+1} = \arg \min_{\theta \in \mathbb{R}^d} \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x, b) (\langle \phi(x, b), \theta \rangle - \hat{Q}_\square^{t+1}(x, b))^2$ .
    Define  $\tilde{Q}_\square^{t+1} := \langle \phi, \sum_{i=0}^{t+1} \theta_\square^i \rangle$ ;  $\hat{V}_\square^{t+1}(s) := \max_a \{ \tilde{Q}_\square^{t+1}(s, a) \} - \max_a \{ \tilde{Q}_\square^t(s, a) \}$ .
end
end

```

In each iteration $t \in [T]$, Alg. 2 uses the buffer \mathcal{B}_t consisting of M samples per state-action pair in \mathcal{C} . However, since \hat{V}^{t+1} and \tilde{Q}^{t+1} depend on all the past θ^i vectors and hence, on the data collected in the previous iterations, the algorithm can effectively leverage all the data in \mathcal{B} . Furthermore, using the difference between the consecutive \tilde{Q} functions can be viewed as a form of variance reduction. This enables us to prove an $O(1/\sqrt{N})$ concentration result for \tilde{Q}^T . Moreover, since the `DataCollection` procedure constructs a coreset which ensures good coverage across the feature space, the resulting sample complexity is independent of the size of the state-action space. Formally, in App. D.2, we prove the following sub-optimality bound for $\square = r + \lambda_k c$ at iteration k of Alg. 1.

Lemma 6 For a fixed $\varepsilon \in (0, 1]$, $\delta \in (0, 1)$, and any $k \in [K]$, when using Alg. 2 at iteration k of Alg. 1 with $\square = r + \lambda_k c$, $M = \tilde{O}\left(\frac{dH^2}{\varepsilon}\right)$ and $T = O\left(\frac{H^2}{\varepsilon}\right)$, the output policy π_T satisfies the following condition with probability $1 - \delta$,

$$\max_{\pi} V_{r+\lambda_k c}^{\pi}(\rho) - V_{r+\lambda_k c}^{\pi_T}(\rho) \leq O((1 + \lambda_k)\varepsilon)$$

Hence, with a buffer \mathcal{B} of size $TM|\mathcal{C}| = \tilde{O}\left(\frac{d^2 H^4}{\varepsilon^2}\right)$, Alg. 2 guarantees an optimality gap of $f_{\text{mdp}}(\mathcal{B}) = O(\varepsilon)$, thereby satisfying Assumption 2. We note that the entropy-regularized variant of the above linear MDVI algorithm (Kitamura et al., 2023) also attains a similar guarantee but for a non-stationary policy output by the corresponding algorithm. Furthermore, in contrast to Thm. 6, the guarantee in Kitamura et al. (2023) only holds for a more restricted range of $\varepsilon \in (0, 1/H]$. In the next section, we instantiate the `PolicyEvaluation` oracle.

4.3. Instantiating the `PolicyEvaluation` oracle: Least-Squares Policy Evaluation

To understand the need for an explicit `PolicyEvaluation` oracle, note that in each iteration k , we can prove that Alg. 2 ensures a concentration guarantee for the value function corresponding

to $r + \lambda_k c$. However, this does not directly imply a concentration guarantee on the individual value functions corresponding to the reward and constraint rewards. This is in contrast to model-based approaches (Vaswani et al., 2022) for tabular CMDPs that guarantee concentration for the empirical transition probabilities, and use that to ensure concentration for both the reward and constraint reward value functions. However, since such model-based approaches cannot be used for linear MDPs, we require an additional algorithm that can compute the empirical value functions satisfying Assumption 3. To that end, we present Alg. 3 that can be used as an instantiation of the PolicyEvaluation oracle in Alg. 1. The algorithm is also based on least-squares and uses the same coreset constructed in Sec. 4.1. Furthermore, we note that Alg. 3 can be viewed as a special case of Alg. 2 for a fixed policy.

Algorithm 3 Least-Squares Policy Evaluation (LS-PE)

Input: T (number of iterations), M (number of next-state samples obtained per state-action pair in each iteration), \diamond (either r or c), $\mathcal{B} = \mathcal{B}_0 \cup \dots \cup \mathcal{B}_{T-1}$ (Buffer), π (policy to be evaluated), $\tilde{\rho}$ (design), \mathcal{C} (coreset), ϕ (feature map).

Output: $\bar{V}_\diamond^T(\rho) = \frac{1}{T} \sum_{i=1}^T \hat{V}_\diamond^i(\rho)$.

Initialize: Define $\hat{V}_\diamond^0 = \mathbf{0}$.

procedure $LS-PE(T, M, \diamond, \mathcal{B}, \pi, \tilde{\rho}, \mathcal{C}, \phi)$

for $t = 0, 1, 2, \dots, T - 1$ **do**

$\forall (s, a) \in \mathcal{C} : \text{Access}(s, a, s'_m)_{m=1}^M$ from the buffer \mathcal{B}_t .

 Define regression target $\hat{Q}_\diamond^{t+1}(s, a) := \diamond(s, a) + \gamma \frac{1}{M} \sum_{m=1}^M \hat{V}_\diamond^t(s'_m)$.

$\omega_\diamond^{t+1} = \arg \min_{\omega \in \mathbb{R}^d} \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x, b) (\langle \phi(x, b), \omega \rangle - \hat{Q}_\diamond^{t+1}(x, b))^2$.

 Define $\hat{V}_\diamond^{t+1}(s) := (\pi \langle \phi, \omega_\diamond^{t+1} \rangle)(s)$.

end

end

Remark 7 *There exists prior CMDP literature that uses PE as a subroutine (Efroni et al., 2020; Ding et al., 2021), but our PE subroutine differs in several aspects. First, Efroni et al. (2020); Ding et al. (2021) use value functions estimated via model-based policy evaluation to update both the primal and dual variables; we utilize the PE subroutine solely for updating the dual variable. This decoupling enables a more reductionist algorithm and modular proof. Moreover, the resulting algorithms in Efroni et al. (2020); Ding et al. (2021) do not achieve the near-optimal regret even for tabular CMDPs (in terms of their dependence on both S and H), and consequently sample-complexity in the generative model setting. In contrast, our method leverages averaging within PE and returns an averaged value function, which plays a crucial role in attaining the near-optimal (in d and ε for the linear setting and in all parameters for the tabular setting) sample complexity. An additional technical challenge lies in integrating the PE-induced error into the primal-dual analysis (i.e., Lemma 11) without losing statistical complexity.*

Note that LS-PE uses a fixed dataset (the buffer \mathcal{B}) to evaluate a fixed policy, and is similar to the policy evaluation algorithms in offline reinforcement learning (Duan et al., 2020). The theoretical guarantees for such offline algorithms depend on the quality of the dataset, measured in terms of metrics such as coverage or concentrability. However, in our case, we curate the dataset and choose the buffer \mathcal{B} such that it has good coverage properties that allow for fine-grained control on the algorithm’s sub-optimality. In particular, we prove the following result in App. D.3.

Lemma 8 For a fixed $\varepsilon \in (0, 1]$, $\delta \in (0, 1)$, Alg. 3 with $M = \tilde{O}\left(\frac{dH^2}{\varepsilon}\right)$ and $T = O\left(\frac{H^2}{\varepsilon}\right)$, the output \bar{V}_\diamond^T satisfies the following condition with probability $1 - \delta$,

$$|\bar{V}_\diamond^T(\rho) - V_\diamond^\pi(\rho)| \leq O(\varepsilon).$$

Hence, with a buffer \mathcal{B} of size $TM|\mathcal{C}| = \tilde{O}\left(\frac{d^2H^4}{\varepsilon^2}\right)$, Alg. 3 guarantees an optimality gap of $f_{\text{eva}}(\mathcal{B}) = O(\varepsilon)$, thereby satisfying Assumption 3.

4.4. Putting everything together

We have seen that Algs. 2 and 3 use the buffer \mathcal{B} constructed by the `DataCollection` procedure to provide control over the terms $f_{\text{mdp}}(\mathcal{B})$ and $f_{\text{eva}}(\mathcal{B})$ appearing in Thm. 4. Combining these results, we prove the following corollary in App. D.4.

Corollary 9 Using `LS-MDVI` (Alg. 2) and `LS-PE` (Alg. 3) as instantiations of the `MDP-Solver` and `PolicyEvaluation` in Alg. 1 and using the `DataCollection` oracle described in Sec. 4.1 has the following guarantee: for a fixed $\varepsilon \in (0, 1]$, $\delta \in (0, 1)$, Alg. 1 with $\tilde{O}\left(\frac{d^2H^4}{\varepsilon^2}\right)$ samples, $U = O\left(\frac{1}{\zeta(1-\gamma)}\right)$, $\eta = \frac{U(1-\gamma)}{\sqrt{K}}$, $K = O\left(\frac{1}{\varepsilon^2(1-\gamma)^2}\right)$, and $b' = b - O(\varepsilon)$, returns a mixture policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - O(\varepsilon), \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b - O(\varepsilon).$$

With the same algorithm parameters, but with $b' = b + O(\varepsilon)$ and $\tilde{O}\left(\frac{d^2H^6}{\zeta^2\varepsilon^2}\right)$ samples, Alg. 1 returns a mixture policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - O(\varepsilon), \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b.$$

Hence, the total sample complexity required to achieve the relaxed feasibility objective in Eq. (2) and the strict feasibility objective in Eq. (3) is $\tilde{O}\left(\frac{d^2H^4}{\varepsilon^2}\right)$ and $\tilde{O}\left(\frac{d^2H^6}{\varepsilon^2\zeta^2}\right)$ respectively. Since the lower bound for unconstrained linear MDPs is $\Omega\left(\frac{d^2H^3}{\varepsilon^2}\right)$ (Weisz et al., 2022), our sample complexity achieves the optimal dependence on d and ε in the relaxed setting.

Flexibility of Alg. 1. Note that instead of `LS-MDVI`, we can use other unconstrained linear MDP solvers. For example, the `G-Sampling and Stop (GSS)` algorithm from Taupin et al. (2023) uses a different `DataCollection` procedure and algorithm to return an ε -optimal policy. It requires $\tilde{O}\left(\frac{d^2H^4}{\varepsilon^2}\right)$ samples to do so, thus matching the sample complexity of `LS-MDVI`. We describe this algorithm in detail and formally instantiate Alg. 1 in App. D.5.

4.5. Lower Bound

As shown by Vaswani et al. (2022), under the relaxed feasibility setting the statistical complexity of solving CMDPs matches that of unconstrained MDPs. In contrast, solving CMDPs under the strict setting is more challenging. In the tabular case, Vaswani et al. (2022) established a lower bound using a linear-programming argument. In this section, we establish a lower bound for solving linear CMDPs under strict feasibility. The proof technique of the following result is simpler and fundamentally different from that of Vaswani et al. (2022). The proof is presented in Section E.

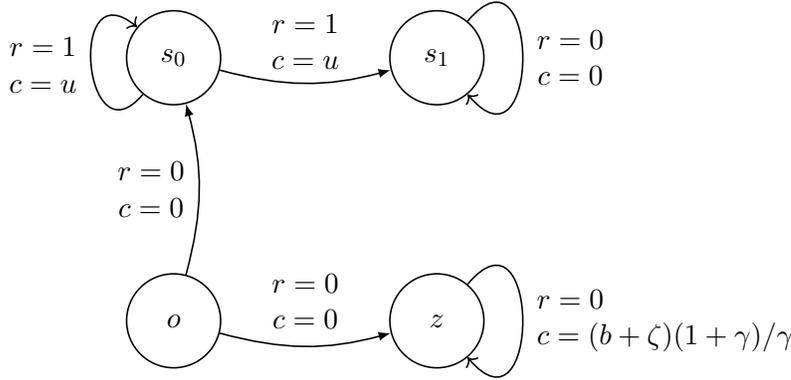


Figure 1: The lower bound instance consists of CMDPs with four states. o is the fixed starting state. At state o , taking an action will either transition to state s_0 or to the “safe” state z . At state s_0 , taking an action will either transition to state s_1 or stay in s_0 . States z and s_1 are absorbing.

Theorem 10 *Let $\delta \in (0, 0.08]$, $\gamma \in [7/12, 1)$, $H = 1/(1 - \gamma)$, $\varepsilon \in (0, 0.002)$, $\zeta \in (0, 49/2280)$, $b \in [H/2, H]$, and $d \geq 6$. There exists a class of linear constrained MDPs such that any (ε, δ) -sound algorithm requires $\Omega(d^2 H^5 / \varepsilon^2 \zeta^2)$ samples from the generative model in the worst case.*

Thus, the dependence on d, ε, ζ in our bounds for the strict feasibility setting is also tight, with a suboptimality arising only in the multiplicative dependence on H . Note that setting $\zeta = H$ reduces the above lower bound to $\Omega(d^2 H^3 / \varepsilon^2)$, which recovers the lower bound in the unconstrained linear MDPs setting.

4.6. Discussion

Why variance-weighted least squares fails in our setting. For unconstrained linear MDPs, [Kitamura et al. \(2023\)](#) provide an alternative entropy-regularized algorithm that constructs coresets that depend on the estimated empirical variance in the value function. The resulting algorithm uses variance-weighted least squares and is able to attain the near-optimal $O\left(\frac{d^2 H^3}{\varepsilon^2}\right)$ sample complexity for unconstrained linear MDPs. To the best of our knowledge, this is the only algorithm that can achieve such an optimal bound. Unfortunately, using such an idea for linear CMDPs fails. This is because in the linear CMDP setting, since the MDP reward function $r + \lambda_k c$ (and hence the MDP value function) change in every iteration k of Alg. 1, using variance-aware coresets implies that we need to construct a distinct coreset in every such iteration. This prevents the resulting algorithm from reusing data similar to Alg. 1, and actually increases the corresponding sample complexity. Resolving this issue and attaining the optimal dependence on H is an important direction for future work.

Comparing with online methods. In order to further contextualize our results, we use the state-of-the-art regret guarantees for the finite-horizon online setting ([Ghosh et al., 2022](#)) and use the reduction in [Bai et al. \(2021\)](#) to our problem setting. The reduction implies that the algorithm in ([Ghosh et al., 2022](#)) (designed and analyzed for the more difficult online regret minimization) results in an $\tilde{O}\left(\frac{d^3 H^4}{\varepsilon^2}\right)$ and $\tilde{O}\left(\frac{d^3 H^6}{\varepsilon^2 \zeta^2}\right)$ sample complexity for the relaxed and strict settings respectively. Hence, our results have a better dimension dependence. Interestingly, the analysis in ([Ghosh et al., 2022](#))

has a worse dependence on d because it uses a uniform concentration argument to get a handle on the concentration for the individual value functions corresponding to the (constraint) rewards. Recall that in Sec. 4.3, we encountered a similar issue and resolved it by using policy evaluation. We believe that our technique might be useful even for online regret minimization.

Comparing with offline methods. Since we store all the samples at the beginning (see Alg. 1), an alternative approach is to treat these stored samples as an offline dataset, and use a standard offline RL algorithm (Neu and Okolo, 2024; Hong and Tewari, 2024) for the CMDP setting. It is worth noting that the sample complexity for such a reduction from the offline constrained RL setting can be substantially worse or even vacuous. In particular, the result on linear CMDP (e.g., Theorem 9) in Hong and Tewari (2024) depends on the concentrability coefficient C^* (see Assumption 2 therein), which can become unbounded if we do not collect data that covers the optimal action in every state. We note that constructing the $O(d^2)$ size coreset as in our paper does not directly ensure such coverage for each optimal action. When using the result from Hong and Tewari (2024), this can lead to unbounded values of C^* and a vacuous bound. For the sake of argument, even if we assume that C^* is bounded by an absolute constant (independent of d and H), the bounds for the offline CMDP setting in Hong and Tewari (2024) are worse than ours. Specifically, Hong and Tewari (2024, Theorem 9) is established under the relaxed feasibility setting; however, their sample complexity bound exhibits a quadratic dependence on the Slater constant, while our bounds in this setting do not have such a dependence. It is known that the Slater constant is only relevant in the strict feasibility regime (Vaswani et al., 2022). Additionally, the dependence on the feature dimension d is d^3 in Hong and Tewari (2024), whereas our approach achieves the near-optimal dependence of d^2 .

5. Conclusion

Given access to a generative model, we proposed a generic primal-dual framework for reducing the (linear) CMDP problem to the (linear) MDP problem. Using (linear) MDVI as the MDP-Solver enabled us to obtain sample complexity bounds for both tabular and linear CMDPs with either $O(\varepsilon)$ or zero constraint violation. We obtained the first near-optimal (in d , ε , and ζ) guarantees for linear CMDPs, whereas for tabular CMDPs, we matched the existing near-optimal guarantees. We also provide a lower bound for solving linear CMDPs under strict feasibility. For linear CMDPs, improving the dependence of the sample complexity on the effective horizon H is an important direction for future work.

References

- Alekh Agarwal, Sham Kakade, and Lin F Yang. Model-based reinforcement learning with a generative model is minimax optimal. In *Conference on Learning Theory*, pages 67–83. PMLR, 2020.
- Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- 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, 2013.

- Qinbo Bai, Amrit Singh Bedi, Mridul Agarwal, Alec Koppel, and Vaneet Aggarwal. Achieving zero constraint violation for constrained reinforcement learning via primal-dual approach. *arXiv preprint arXiv:2109.06332*, 2021.
- Kianté Brantley, Miroslav Dudik, Thodoris Lykouris, Sobhan Miryoosefi, Max Simchowitz, Aleksandrs Slivkins, and Wen Sun. Constrained episodic reinforcement learning in concave-convex and knapsack settings. *arXiv preprint arXiv:2006.05051*, 2020.
- Chiara Buratti, Andrea Conti, Davide Dardari, and Roberto Verdone. An overview on wireless sensor networks technology and evolution. *Sensors*, 9(9):6869–6896, 2009.
- Dongsheng Ding, Xiaohan Wei, Zhuoran Yang, Zhaoran Wang, and Mihailo Jovanovic. Provably efficient safe exploration via primal-dual policy optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 3304–3312. PMLR, 2021.
- Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017.
- Yaqi Duan, Zeyu Jia, and Mengdi Wang. Minimax-optimal off-policy evaluation with linear function approximation. In *International Conference on Machine Learning*, pages 2701–2709. PMLR, 2020.
- Yonathan Efroni, Shie Mannor, and Matteo Pirota. Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189*, 2020.
- Ather Gattami, Qinbo Bai, and Vaneet Aggarwal. Reinforcement learning for constrained markov decision processes. In *International Conference on Artificial Intelligence and Statistics*, pages 2656–2664. PMLR, 2021.
- Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. In *International conference on machine learning*, pages 2160–2169. PMLR, 2019.
- Arnob Ghosh, Xingyu Zhou, and Ness Shroff. Provably efficient model-free constrained rl with linear function approximation. *Advances in Neural Information Processing Systems*, 35:13303–13315, 2022.
- Arnob Ghosh, Xingyu Zhou, and Ness Shroff. Towards achieving sub-linear regret and hard constraint violation in model-free rl. In *International Conference on Artificial Intelligence and Statistics*, pages 1054–1062. PMLR, 2024.
- Aria HasanzadeZonuzi, Dileep M. Kalathil, and Srinivas Shakkottai. Model-based reinforcement learning for infinite-horizon discounted constrained markov decision processes. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 2519–2525. ijcai.org, 2021.
- Kihyuk Hong and Ambuj Tewari. A primal-dual algorithm for offline constrained reinforcement learning with linear mdps. *arXiv preprint arXiv:2402.04493*, 2024.
- Pihe Hu, Yu Chen, and Longbo Huang. Nearly minimax optimal reinforcement learning with linear function approximation. In *International Conference on Machine Learning*, pages 8971–9019. PMLR, 2022.

- Arushi Jain, Sharan Vaswani, Reza Babanezhad, Csaba Szepesvari, and Doina Precup. Towards painless policy optimization for constrained mdps. *arXiv preprint arXiv:2204.05176*, 2022.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on learning theory*, pages 2137–2143. PMLR, 2020.
- Yujia Jin, Ishani Karmarkar, Aaron Sidford, and Jiayi Wang. Truncated variance reduced value iteration. *arXiv preprint arXiv:2405.12952*, 2024.
- David Julian, Mung Chiang, Daniel O’Neill, and Stephen Boyd. Qos and fairness constrained convex optimization of resource allocation for wireless cellular and ad hoc networks. In *Proceedings. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 477–486. IEEE, 2002.
- Sham Machandranath Kakade. *On the sample complexity of reinforcement learning*. University of London, University College London (United Kingdom), 2003.
- Krishna Chaitanya Kalagarla, Rahul Jain, and Pierluigi Nuzzo. A sample-efficient algorithm for episodic finite-horizon MDP with constraints. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI*, pages 8030–8037. AAAI Press, 2021.
- Michael Kearns and Satinder Singh. Finite-sample convergence rates for q-learning and indirect algorithms. *Advances in neural information processing systems*, pages 996–1002, 1999.
- Toshinori Kitamura, Tadashi Kozuno, Yunhao Tang, Nino Vieillard, Michal Valko, Wenhao Yang, Jincheng Mei, Pierre Ménard, Mohammad Gheshlaghi Azar, Rémi Munos, et al. Regularization and variance-weighted regression achieves minimax optimality in linear mdps: theory and practice. In *International Conference on Machine Learning*, pages 17135–17175. PMLR, 2023.
- Tadashi Kozuno, Eiji Uchibe, and Kenji Doya. Theoretical analysis of efficiency and robustness of softmax and gap-increasing operators in reinforcement learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2995–3003. PMLR, 2019.
- Tadashi Kozuno, Wenhao Yang, Nino Vieillard, Toshinori Kitamura, Yunhao Tang, Jincheng Mei, Pierre Ménard, Mohammad Gheshlaghi Azar, Michal Valko, Rémi Munos, et al. K1-entropy-regularized rl with a generative model is minimax optimal. *arXiv preprint arXiv:2205.14211*, 2022.
- Guanghui Lan. Policy mirror descent for reinforcement learning: Linear convergence, new sampling complexity, and generalized problem classes. *Mathematical programming*, 198(1):1059–1106, 2023.
- Tor Lattimore, Csaba Szepesvari, and Gellert Weisz. Learning with good feature representations in bandits and in rl with a generative model. In *International conference on machine learning*, pages 5662–5670. PMLR, 2020.
- 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. *Advances in neural information processing systems*, 33:12861–12872, 2020.

- Qinghua Liu, Gellért Weisz, András György, Chi Jin, and Csaba Szepesvári. Optimistic natural policy gradient: a simple efficient policy optimization framework for online rl. *Advances in Neural Information Processing Systems*, 36:3560–3577, 2023.
- Tao Liu, Ruida Zhou, Dileep Kalathil, P. R. Kumar, and Chao Tian. Policy optimization for constrained mdps with provable fast global convergence, 2022.
- Sobhan Miryoosefi and Chi Jin. A simple reward-free approach to constrained reinforcement learning. In *International Conference on Machine Learning*, pages 15666–15698. PMLR, 2022.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Washim U Mondal and Vaneet Aggarwal. Sample-efficient constrained reinforcement learning with general parameterization. *Advances in Neural Information Processing Systems*, 37:68380–68405, 2024.
- Gergely Neu and Nneka Okolo. Offline rl via feature-occupancy gradient ascent. *arXiv preprint arXiv:2405.13755*, 2024.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744, 2022.
- Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Shuang Qiu, Xiaohan Wei, Zhuoran Yang, Jieping Ye, and Zhaoran Wang. Upper confidence primal-dual reinforcement learning for cmdp with adversarial loss. *Advances in Neural Information Processing Systems*, 33:15277–15287, 2020.
- Andrew J Schaefer, Matthew D Bailey, Steven M Shechter, and Mark S Roberts. Modeling medical treatment using markov decision processes. In *Operations research and health care*, pages 593–612. Springer, 2005.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Uri Sherman, Alon Cohen, Tomer Koren, and Yishay Mansour. Rate-optimal policy optimization for linear markov decision processes. *arXiv preprint arXiv:2308.14642*, 2023.
- Aaron Sidford, Mengdi Wang, Xian Wu, Lin F Yang, and Yinyu Ye. Near-optimal time and sample complexities for solving markov decision processes with a generative model. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 5192–5202, 2018.

- Aaron Sidford, Mengdi Wang, Xian Wu, and Yinyu Ye. Variance reduced value iteration and faster algorithms for solving markov decision processes. *Naval Research Logistics (NRL)*, 70(5): 423–442, 2023.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- Jie Tan, Tingnan Zhang, Erwin Coumans, Atıl İscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. *arXiv preprint arXiv:1804.10332*, 2018.
- Jerome Taupin, Yassir Jedra, and Alexandre Proutiere. Best policy identification in linear mdps. In *2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1–8. IEEE, 2023.
- Tian Tian, Lin Yang, and Csaba Szepesvári. Confident natural policy gradient for local planning in q_π -realizable constrained mdps. *Advances in Neural Information Processing Systems*, 37: 76139–76176, 2024.
- Michael J Todd. *Minimum-volume ellipsoids: Theory and algorithms*. SIAM, 2016.
- Sharan Vaswani, Lin Yang, and Csaba Szepesvári. Near-optimal sample complexity bounds for constrained mdps. *Advances in Neural Information Processing Systems*, 35:3110–3122, 2022.
- Nino Vieillard, Tadashi Kozuno, Bruno Scherrer, Olivier Pietquin, Rémi Munos, and Matthieu Geist. Leverage the average: an analysis of kl regularization in rl. *arXiv preprint arXiv:2003.14089*, 2020.
- Mengdi Wang. Randomized linear programming solves the discounted markov decision problem in nearly-linear (sometimes sublinear) running time. *arXiv preprint arXiv:1704.01869*, 2017.
- Honghao Wei, Xin Liu, and Lei Ying. A provably-efficient model-free algorithm for constrained markov decision processes. *arXiv preprint arXiv:2106.01577*, 2021.
- Gellért Weisz, András György, Tadashi Kozuno, and Csaba Szepesvári. Confident approximate policy iteration for efficient local planning in q^π -realizable mdps. *Advances in Neural Information Processing Systems*, 35:25547–25559, 2022.
- Lin Yang and Mengdi Wang. Sample-optimal parametric q-learning using linearly additive features. In *International conference on machine learning*, pages 6995–7004. PMLR, 2019.
- Tiancheng Yu, Yi Tian, Jingzhao Zhang, and Suvrit Sra. Provably efficient algorithms for multi-objective competitive rl. In *International Conference on Machine Learning*, pages 12167–12176. PMLR, 2021.

Andy Zeng, Shuran Song, Johnny Lee, Alberto Rodriguez, and Thomas Funkhouser. Tossingbot: Learning to throw arbitrary objects with residual physics. *IEEE Transactions on Robotics*, 36(4): 1307–1319, 2020.

Liyuan Zheng and Lillian Ratliff. Constrained upper confidence reinforcement learning. In *Learning for Dynamics and Control*, pages 620–629. PMLR, 2020.

Supplementary material

A	An Instantiation of <code>ComputeOptimalDesign</code>	20
B	Table of Notation	21
C	Proof of Theorem 4	22
C.1	Proof of Lemma 11 (Primal-Dual Guarantees for Algorithm 1)	23
D	Proofs for Section 4	27
D.1	Deriving <code>LS-MDVI</code> from Entropic Mirror Descent	27
D.2	Proof of Lemma 6 (Optimality Guarantees for Algorithm 2 - Linear CMDP)	28
D.2.1	Auxiliary Lemmas	32
D.3	Proof of Lemma 8 (Optimality Guarantees for Algorithm 3 - Linear CMDP)	37
D.3.1	Auxiliary Lemmas	39
D.4	Proof of Corollary 9	39
D.5	Instantiating the <code>MDP-Solver</code> : G-Sampling-and-Stop	40
E	Lower Bound	41
F	Algorithms for Solving Tabular CMDPs	45
G	Instantiating the Framework for Tabular Constrained MDPs	45
H	Proofs for Section G	46
H.1	Proof of Lemma 29 (Optimality Guarantees for Algorithm 6 - Tabular CMDP) . . .	47
H.1.1	Proof of Lemma 33 and Lemma 34 (Proofs with Hoeffding’s Inequality) . .	47
H.1.2	Proof of Lemma 35 and Lemma 36 (Proofs with Bernstein’s Inequality) . .	51
H.1.3	Auxiliary Lemmas	53
H.2	Proof of Lemma 30 (Optimality Guarantees for Algorithm 7 - Tabular CMDP) . . .	55
H.2.1	Auxiliary Lemmas	55
H.3	Proof of Lemma 46 and Lemma 47 (Concentration Error Bounds with Bernstein’s Inequality - Tabular CMDP)	58
H.3.1	Auxiliary Lemmas for Lemma 46	60
H.3.2	Auxiliary Lemmas for Lemma 47	65
H.4	Proof of Corollary 31	67
H.5	Instantiating the <code>MDP-Solver</code> : Model-based algorithm (Li et al., 2020)	67
I	Supporting Lemmas	68
I.1	Concentration Inequalities	68
I.2	Lemmas for Variances	69
I.3	Lemmas for Constrained MDPs	70

Appendix A. An Instantiation of ComputeOptimalDesign

In this section, we present an instantiation of the ComputeOptimalDesign oracle using the Frank-Wolfe algorithm (Todd, 2016).

Algorithm 4 InitializeDesign

Choose an arbitrary nonzero $c_0 \in \mathbb{R}^d$.

Output: $\tilde{\rho}$.

procedure *InitializeDesign*

for $j = 0, 1, 2, \dots, d-1$ **do**

$(\bar{s}_j, \bar{a}_j) = \arg \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} c_j^\top \phi(s, a)$.

$(s_j, a_j) = \arg \min_{(s,a) \in \mathcal{S} \times \mathcal{A}} c_j^\top \phi(s, a)$.

$x_j = \phi(\bar{s}_j, \bar{a}_j) - \phi(s_j, a_j)$.

 Choose an arbitrary nonzero c_{j+1} orthogonal to x_0, \dots, x_j .

end

 Let $\mathcal{Z} := \{(\bar{s}_j, \bar{a}_j), (s_j, a_j) \mid j = 0, \dots, d-1\}$.

 Choose $\tilde{\rho}$ to put equal weight on each of the distinct points of \mathcal{Z} .

end

Below we present the classical Frank-Wolfe algorithm for experimental design.

Algorithm 5 Frank-Wolfe

Input: ε^{FW} . (Tolerance for algorithm)

Output: $\tilde{\rho}, \mathcal{C}, G$. (Coreset, optimal design and covariance matrix)

Initialize: Define $\hat{V}_\diamond^0 = \mathbf{0}$.

procedure *Frank-Wolfe*(ε^{FW})

$\tilde{\rho} = \text{InitializeDesign}$ by Algorithm 4.

 Define $\mathcal{U} : \tilde{\rho} \mapsto \text{diag}(\tilde{\rho}) \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times |\mathcal{S}||\mathcal{A}|}$, where $\text{diag}(\tilde{\rho})$ is a diagonal matrix with elements of $\tilde{\rho}$.

 For $(s, a) \in \mathcal{S} \times \mathcal{A}$, let $\Phi \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times d}$ be a matrix such that its $(s|\mathcal{A}| + a)$ th row is $\phi(s, a)$.

 Define $\mathcal{I} : \tilde{\rho} \mapsto (\Phi^\top \mathcal{U}(\tilde{\rho}) \Phi)^{-1}$. (defines the inverse of the covariance matrix)

 Let $\nu : (s, a, \tilde{\rho}) \mapsto \phi(s, a)^\top \mathcal{I}(\tilde{\rho}) \phi(s, a)$. (measures the variance proxy for (s, a))

 Let $\delta : \tilde{\rho} \mapsto \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} (\nu(s, a, \tilde{\rho}) - d)/d$ (computes the relative difference between the worst-case variance and d)

while $\delta(\tilde{\rho}) > \varepsilon^{FW}$ **do**

 Let $(x, b) := \arg \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \nu(s, a, \tilde{\rho})$.

 Let $\eta^* := (\nu(x, b, \tilde{\rho}) - d)/((d-1)\nu(x, b, \tilde{\rho}))$.

$\tilde{\rho}(x, b) \leftarrow \tilde{\rho}(x, b) + \eta^*$.

$\tilde{\rho} \leftarrow \tilde{\rho}/(1 + \eta^*)$.

end

 Let $\mathcal{C} := \left\{ (s, a) \mid \nu(s, a, \tilde{\rho}) \geq d \left(1 + \frac{\delta(\tilde{\rho})d}{2} - \sqrt{\delta(\tilde{\rho})(d-1) + \frac{\delta(\tilde{\rho})^2 d^2}{4}} \right) \right\}$.

(form the coreset containing state-action pairs with sufficiently high variance value)

 Let $G := \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x, b) \phi(x, b) \phi(x, b)^\top$. (calculate the corresponding covariance matrix)

end

The subroutine `InitializeDesign` returns an initial design to be used in Frank–Wolfe. `InitializeDesign` is a deterministic procedure for constructing a core set of state-action pairs that provides good coverage of the feature space in linear MDPs. The algorithm sequentially identifies informative directions in the feature space by iteratively computing difference vectors between state-action pairs with maximal and minimal feature projections along a given search direction. The algorithm iteratively updates the search direction to be orthogonal to the span of the previously discovered directions. Specifically, the vector $c_j \in \mathbb{R}^d$ is an auxiliary direction vector used to sequentially identify maximally informative state-action pairs. The next vector c_{j+1} is then chosen to be orthogonal to all previous x_0, \dots, x_j ensuring that the design explores linearly independent directions in feature space. The resulting set of state-action pairs is then used as the support for a design distribution in regression.

Appendix B. Table of Notation

Notation	Meaning
\mathcal{A}, \mathcal{S}	action space of size $ \mathcal{A} $, state space of size $ \mathcal{S} $
γ, H	discount factor in $[0, 1)$, $1/(1 - \gamma)$
P	transition matrix $P \in \mathbb{R}^{ \mathcal{S} \mathcal{A} \times \mathcal{S} }$
P_π, \hat{P}_π^t	$\pi P \in \mathbb{R}^{ \mathcal{S} \times \mathcal{S} }$, $\pi \hat{P}_\pi^t \in \mathbb{R}^{ \mathcal{S} \times \mathcal{S} }$
r, c	reward vector in $[0, 1]$ range, constraint reward vector in $[0, 1]$ range
ρ	initial distribution of states
\diamond	r or c
\square	$r + \lambda c$ where $\lambda \in \{\lambda_1, \dots, \lambda_K\}$
b, ζ	constraint value in $[0, 1/(1 - \gamma))$, Slater constant
λ, λ^*	Lagrange multiplier, the optimal Lagrange multiplier
U	projection upper bound
ϕ, d	feature map of a linear MDP and its dimension
$\tilde{\rho}, \mathcal{C}$	a design over $\mathcal{S} \times \mathcal{A}$, coreset
G	design matrix with respect to ϕ and $\tilde{\rho}$. Equal to $\sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x,b) \phi(x,b) \phi(x,b)^\top$
$W(z)$	$G^{-1} \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x,b) \phi(x,b) z(x,b)$ (solution of a least-squares estimation with features $\phi(x,b)$, weights $\tilde{\rho}$ and targets $z(x,b)$)
ε, δ	admissible suboptimality, admissible failure probability
K, T	number of outer and inner iterations
$(\hat{P}_t \hat{V}_\diamond^t)(s, a)$	$\frac{1}{M} \sum_{m=1}^M \hat{V}_\diamond^t(s'_m)$ where $s'_m \in \mathcal{B}_t$
$(P \hat{V}_\diamond^t)(s, a)$	$\mathbb{E}[\hat{V}_\diamond^t(s') \mid s_0 = s, a_0 = a]$
$\mathcal{F}_{t,m}$	σ -algebra in the filtration for Algs. 2, 3, 6 and 7
$\mathcal{T}^\pi Q$	Bellman operator $r + \gamma P(\pi Q)$
Q^π	state-action value function for policy π
\tilde{Q}_\square^t	estimated state-action value function in iteration t in Algs. 2 and 6
\hat{Q}_\diamond^t	estimated state-action value function in iteration t in Algs. 3 and 7

$\hat{V}_\square^t(s)$ (Tabular)	$\max_a \left\{ \sum_{i=0}^t \hat{Q}_\square^i(s, a) \right\} - \max_a \left\{ \sum_{i=0}^{t-1} \hat{Q}_\square^i(s, a) \right\}$ in Alg. 6
$\hat{V}_\square^t(s)$ (Linear)	$\max_a \left\{ \langle \phi, \sum_{i=0}^t \theta_\square^i \rangle (s, a) \right\} - \max_a \left\{ \langle \phi, \sum_{i=0}^{t-1} \theta_\square^i \rangle (s, a) \right\}$ in Alg. 2
\hat{Q}_\square^t (Tabular)	$\sum_{i=0}^t \hat{Q}_\square^i$ in Alg. 6
\hat{Q}_\square^t (Linear)	$\langle \phi, \sum_{i=0}^t \theta_\square^i \rangle$ in Alg. 2
$\hat{V}_\diamond^t(s)$ (Tabular)	$(\pi \hat{Q}_\diamond^t)(s)$ in Alg. 7
$\hat{V}_\diamond^t(s)$ (Linear)	$(\pi \langle \phi, \omega_\diamond^t \rangle)(s)$ in Alg. 3
$\bar{V}_\square^t(s), \bar{V}_\diamond^t(s)$	$\frac{1}{t} \sum_{i=1}^t \hat{V}_\square^i(s), \frac{1}{t} \sum_{i=1}^t \hat{V}_\diamond^i(s)$
$\hat{V}_\diamond^k, \bar{V}_\diamond^\pi$	output of the PolicyEvaluation oracle in line 5 in Algorithm 1, $\frac{1}{K} \sum_{k=0}^{K-1} \hat{V}_\diamond^k$
π_k	output policy of MDP-Solver
$\bar{\pi}$	mixture policy equal to $\frac{1}{K} \sum_{k=0}^{K-1} \pi_k$
π^*	$\operatorname{argmax}_\pi V_r^\pi(\rho)$ s.t. $V_c^\pi(\rho) \geq b$
π^{*+}	$\operatorname{argmax}_\pi V_r^\pi(\rho)$ s.t. $V_c^\pi(\rho) \geq b + 6f(\mathcal{B})$
π_k^*	$\operatorname{argmax}_\pi \{V_{r+\lambda_k c}^\pi\}$
π'_t	a non-stationary policy that follows policies π_t, π_{t-1}, \dots upto timestep t and follows π_0 thereafter
$(\pi Q)(s)$	$\sum_{a \in \mathcal{A}} \pi(a s) Q(s, a)$
$(\pi r)(s)$	$\sum_{a \in \mathcal{A}} \pi(a s) r(s, a)$
θ_\square^t	least-squares value estimate in Alg. 2
θ_\square^t	parameter that satisfies $\langle \phi, \theta_\square^t \rangle := \square + \gamma P \hat{V}_\square^{t-1}$ in the linear MDP
ω_\diamond^t	least-squares value estimate in Alg. 3
ω_\diamond^t	parameter that satisfies $\langle \phi, \omega_\diamond^t \rangle := \diamond + \gamma P \hat{V}_\diamond^{t-1}$ in the linear MDP

Appendix C. Proof of Theorem 4

Theorem 4 Suppose Assumptions 2 and 3 hold and let $f(\mathcal{B}) := \max\{f_{\text{mdp}}(\mathcal{B}), f_{\text{eva}}(\mathcal{B})\}$. For $\delta \in (0, 1)$, Alg. 1 with $U = \frac{2}{\zeta(1-\gamma)}$, $\eta = \frac{U(1-\gamma)}{\sqrt{K}}$, $K = \frac{U^2}{[f(\mathcal{B})]^2(1-\gamma)^2}$ and $b' = b - 2f(\mathcal{B})$, returns a mixture policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - 4f(\mathcal{B}) \quad , \quad V_c^{\bar{\pi}}(\rho) \geq b - 6f(\mathcal{B}). \quad (\text{Relaxed Feasibility Setting})$$

With the same algorithm parameters, but with $b' = b + 4f(\mathcal{B})$ for $f(\mathcal{B}) \leq \frac{\zeta}{6}$, Alg. 1 returns a mixture policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - \frac{16f(\mathcal{B})}{\zeta(1-\gamma)} \quad , \quad V_c^{\bar{\pi}}(\rho) \geq b. \quad (\text{Strict Feasibility Setting})$$

Proof We denote $\bar{V}_\diamond^\pi = \frac{1}{K} \sum_{k=0}^{K-1} \hat{V}_\diamond^k$ where $\diamond = r$ or c . We first prove the relaxed feasibility statement. By Lemma 11, we have $\bar{V}_c^{\bar{\pi}}(\rho) \geq b - 5f(\mathcal{B})$. Hence,

$$\begin{aligned} V_c^{\bar{\pi}}(\rho) &= V_c^{\bar{\pi}}(\rho) - \bar{V}_c^{\bar{\pi}}(\rho) + \bar{V}_c^{\bar{\pi}}(\rho) \\ &\geq b - 5f(\mathcal{B}) - |V_c^{\bar{\pi}}(\rho) - \bar{V}_c^{\bar{\pi}}(\rho)| \\ &\geq b - 5f(\mathcal{B}) - f(\mathcal{B}) && (\text{By Assumption 3 for each policy } \{\pi_k\}_{k=0}^{K-1}) \\ &= b - 6f(\mathcal{B}). \end{aligned}$$

Next, we prove $V_r^{\pi^*}(\rho) - V_r^{\bar{\pi}}(\rho) \leq 4f(\mathcal{B})$. We have

$$\begin{aligned} V_r^{\pi^*}(\rho) - V_r^{\bar{\pi}}(\rho) &= [V_r^{\pi^*}(\rho) - \bar{V}_r^{\bar{\pi}}(\rho)] + [\bar{V}_r^{\bar{\pi}}(\rho) - V_r^{\bar{\pi}}(\rho)] \\ &\leq 3f(\mathcal{B}) + |\bar{V}_r^{\bar{\pi}}(\rho) - V_r^{\bar{\pi}}(\rho)| && \text{(By Lemma 11)} \\ &\leq 3f(\mathcal{B}) + f(\mathcal{B}) && \text{(By Assumption 3 for each policy } \{\pi_k\}_{k=0}^{K-1}\text{)} \\ &= 4f(\mathcal{B}). \end{aligned}$$

Now we prove the strict feasibility statement. By Lemma 11, we have $\bar{V}_c^{\bar{\pi}}(\rho) \geq b + f(\mathcal{B})$, and thus,

$$\begin{aligned} V_c^{\bar{\pi}}(\rho) &= V_c^{\bar{\pi}}(\rho) - \bar{V}_c^{\bar{\pi}}(\rho) + \bar{V}_c^{\bar{\pi}}(\rho) \\ &\geq b + f(\mathcal{B}) - |V_c^{\bar{\pi}}(\rho) - \bar{V}_c^{\bar{\pi}}(\rho)| \\ &\geq b + f(\mathcal{B}) - f(\mathcal{B}) && \text{(By Assumption 3 for each policy } \{\pi_k\}_{k=0}^{K-1}\text{)} \\ &\geq b, \end{aligned}$$

which satisfies the constraint. Next, we prove $V_r^{\pi^*}(\rho) - V_r^{\bar{\pi}}(\rho) \leq 28f(\mathcal{B})$. We define $\pi^{*+} \in \operatorname{argmax}_{\pi} V_r^{\pi}(\rho)$ s.t. $V_c^{\pi}(\rho) \geq b + 6f(\mathcal{B})$. Note that such a policy exists by the definition of ζ and the assumption that $f(\mathcal{B}) \leq \frac{\zeta}{6}$. By Lemma 66 and Lemma 65, we know that

$$|V_r^{\pi^*}(\rho) - V_r^{\pi^{*+}}(\rho)| \leq 12f(\mathcal{B})\lambda^* \leq \frac{12f(\mathcal{B})}{\zeta(1-\gamma)}.$$

Applying Lemma 11 and Assumption 3 as before, we have

$$\begin{aligned} V_r^{\pi^*}(\rho) - V_r^{\bar{\pi}}(\rho) &= [V_r^{\pi^*}(\rho) - V_r^{\pi^{*+}}(\rho)] + [V_r^{\pi^{*+}}(\rho) - \bar{V}_r^{\bar{\pi}}(\rho)] + [\bar{V}_r^{\bar{\pi}}(\rho) - V_r^{\bar{\pi}}(\rho)] \\ &\leq \frac{12f(\mathcal{B})}{\zeta(1-\gamma)} + 3f(\mathcal{B}) + f(\mathcal{B}) \\ &\leq \frac{16f(\mathcal{B})}{\zeta(1-\gamma)}. && (\zeta(1-\gamma) \leq \frac{1-\gamma}{1-\gamma} = 1) \end{aligned}$$

This completes the proof. ■

C.1. Proof of Lemma 11 (Primal-Dual Guarantees for Algorithm 1)

Lemma 11 (Primal-Dual Guarantees for Algorithm 1) *Suppose Assumptions 2 and 3 hold and let $f(\mathcal{B}) := \max\{f_{\text{mdp}}(\mathcal{B}), f_{\text{eva}}(\mathcal{B})\}$. For $\delta \in (0, 1)$, when Alg. 1 is run with $U = \frac{2}{\zeta(1-\gamma)}$, $\eta = \frac{U(1-\gamma)}{\sqrt{K}}$, $K = \frac{U^2}{[f(\mathcal{B})]^2(1-\gamma)^2}$ and $b' = b - 2f(\mathcal{B})$, the following condition holds with probability $1 - \delta$,*

$$\frac{1}{K} \sum_{k=0}^{K-1} \hat{V}_r^k(\rho) \geq V_r^{\pi^*}(\rho) - 3f(\mathcal{B}) \quad , \quad \frac{1}{K} \sum_{k=0}^{K-1} \hat{V}_c^k(\rho) \geq b - 5f(\mathcal{B}).$$

With the same algorithm parameters, but with $b' = b + 4f(\mathcal{B})$, the following condition holds with probability $1 - \delta$,

$$\frac{1}{K} \sum_{k=0}^{K-1} \hat{V}_r^k(\rho) \geq V_r^{\pi^{*+}}(\rho) - 3f(\mathcal{B}) \quad , \quad \frac{1}{K} \sum_{k=0}^{K-1} \hat{V}_c^k(\rho) \geq b + f(\mathcal{B}).$$

Proof We begin by proving the first part of the lemma. Since both r and c are bounded by 1, we note that $r(s, a) + \lambda_k c(s, a) \leq 1 + \lambda_k$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. Define $\pi_k^* := \arg \max_{\pi} V_{r+\lambda_k c}^{\pi}$ as an optimal policy in the MDP with rewards $r + \lambda_k c$. For each iteration k in Alg. 1, by Assumption 2 with $R = 1 + \lambda_k$, we have

$$V_r^{\pi_k^*}(\rho) + \lambda_k V_c^{\pi_k^*}(\rho) - V_{r+\lambda_k c}^{\pi_k} \leq f_{\text{mdp}}(\mathcal{B})(1 + \lambda_k).$$

By Assumption 3 for policy π_k , we have

$$\begin{aligned} V_{r+\lambda_k c}^{\pi_k}(\rho) - \hat{V}_r^k(\rho) - \lambda_k \hat{V}_c^k(\rho) &= V_r^{\pi_k}(\rho) + \lambda_k V_c^{\pi_k}(\rho) - \hat{V}_r^k(\rho) - \lambda_k \hat{V}_c^k(\rho) \\ &= V_r^{\pi_k}(\rho) - \hat{V}_r^k(\rho) + \lambda_k (V_c^{\pi_k}(\rho) - \hat{V}_c^k(\rho)) \\ &\leq f_{\text{eva}}(\mathcal{B})(1 + \lambda_k). \end{aligned}$$

Combining the above inequalities and letting $f(\mathcal{B}) = \max\{f_{\text{mdp}}(\mathcal{B}), f_{\text{eva}}(\mathcal{B})\}$, we obtain

$$V_r^{\pi_k^*}(\rho) + \lambda_k V_c^{\pi_k^*}(\rho) - (\hat{V}_r^k(\rho) + \lambda_k \hat{V}_c^k(\rho)) \leq (f_{\text{mdp}}(\mathcal{B}) + f_{\text{eva}}(\mathcal{B}))(1 + \lambda_k) \leq 2f(\mathcal{B})(1 + \lambda_k).$$

By the definition of π_k^* ,

$$V_r^{\pi^*}(\rho) + \lambda_k V_c^{\pi^*}(\rho) \leq V_r^{\pi_k^*}(\rho) + \lambda_k V_c^{\pi_k^*}(\rho).$$

Therefore, by combining the above inequalities,

$$\begin{aligned} V_r^{\pi^*}(\rho) + \lambda_k V_c^{\pi^*}(\rho) &\leq \hat{V}_r^k(\rho) + \lambda_k \hat{V}_c^k(\rho) + 2f(\mathcal{B})(1 + \lambda_k) \\ \implies V_r^{\pi^*}(\rho) - \hat{V}_r^k(\rho) &\leq \lambda_k (\hat{V}_c^k(\rho) - V_c^{\pi^*}(\rho) + 2f(\mathcal{B})) + 2f(\mathcal{B}). \end{aligned} \quad (5)$$

Since $V_c^{\pi^*}(\rho) \geq b$ and $\lambda_k \geq 0$, we obtain

$$V_r^{\pi^*}(\rho) - \hat{V}_r^k(\rho) \leq \lambda_k (\hat{V}_c^k(\rho) - b + 2f(\mathcal{B})) + 2f(\mathcal{B}).$$

By taking the average, letting $b' = b - 2f(\mathcal{B})$, and adding both sides by the same term $\frac{\lambda}{K} \sum_{k=0}^{K-1} [b' - \hat{V}_c^k(\rho)]$,

$$\frac{1}{K} \sum_{k=0}^{K-1} [V_r^{\pi^*}(\rho) - \hat{V}_r^k(\rho)] + \frac{\lambda}{K} \sum_{k=0}^{K-1} [b' - \hat{V}_c^k(\rho)] \leq \frac{1}{K} \sum_{k=0}^{K-1} (\lambda_k - \lambda) (\hat{V}_c^k(\rho) - b') + 2f(\mathcal{B}).$$

Now we define $R(\lambda, K) := \sum_{k=0}^{K-1} (\lambda_k - \lambda) (\hat{V}_c^k(\rho) - b')$ as the dual regret and denote $\bar{V}_{\diamond}^{\pi} = \frac{1}{K} \sum_{k=0}^{K-1} \hat{V}_{\diamond}^k$ (where $\diamond = r$ or c). Thus, for any $\lambda \in [0, U]$,

$$V_r^{\pi^*}(\rho) - \bar{V}_r^{\pi}(\rho) + \lambda(b' - \bar{V}_c^{\pi}(\rho)) \leq \frac{R(\lambda, K)}{K} + 2f(\mathcal{B}). \quad (6)$$

Below we show that for any $\lambda \in [0, U]$, the following bound holds for the dual regret:

$$R(\lambda, K) \leq \frac{U\sqrt{K}}{1 - \gamma}.$$

Using the dual update in Alg. 1, we observe that,

$$\begin{aligned}
 |\lambda_{k+1} - \lambda|^2 &\leq \left| \lambda_k - \eta \left(\hat{V}_c^k(\rho) - b' \right) - \lambda \right|^2 && \text{(by non-expansiveness of projection)} \\
 &= |\lambda_k - \lambda|^2 - 2\eta(\lambda_k - \lambda) \left(\hat{V}_c^k(\rho) - b' \right) + \eta^2 \left(\hat{V}_c^k(\rho) - b' \right)^2 \\
 &\stackrel{(a)}{\leq} |\lambda_k - \lambda|^2 - 2\eta(\lambda_k - \lambda) \left(\hat{V}_c^k(\rho) - b' \right) + \frac{\eta^2}{(1-\gamma)^2},
 \end{aligned}$$

where (a) follows because b and the constraint value are in the $[0, 1/(1-\gamma)]$ interval. Rearranging and dividing by 2η , we get

$$(\lambda_k - \lambda) \left(\hat{V}_c^k(\rho) - b' \right) \leq \frac{|\lambda_k - \lambda|^2 - |\lambda_{k+1} - \lambda|^2}{2\eta} + \frac{\eta}{2(1-\gamma)^2}.$$

Summing from $k = 0$ to $K - 1$ and using the definition of the dual regret,

$$R(\lambda, K) \leq \frac{1}{2\eta} \sum_{k=0}^{K-1} \left[|\lambda_k - \lambda|^2 - |\lambda_{k+1} - \lambda|^2 \right] + \frac{\eta K}{2(1-\gamma)^2}.$$

Telescoping, bounding $|\lambda_0 - \lambda|$ by U and dropping a negative term gives

$$R(\lambda, K) \leq \frac{U^2}{2\eta} + \frac{\eta K}{2(1-\gamma)^2},$$

Setting $\eta = \frac{U(1-\gamma)}{\sqrt{K}}$,

$$R(\lambda, K) \leq \frac{U\sqrt{K}}{1-\gamma}. \tag{7}$$

Next, in order to bound the reward optimality gap, setting $\lambda = 0$ in Eq. (6) and using the above bound on the dual regret, we obtain

$$V_r^{\pi^*}(\rho) - \bar{V}_r^{\bar{\pi}}(\rho) \leq \frac{U}{(1-\gamma)\sqrt{K}} + 2f(\mathcal{B}). \tag{8}$$

In order to bound the constraint violation, we consider two cases. The first case is when $b' - \bar{V}_c^{\bar{\pi}}(\rho) \leq 0$. Consequently, $b - 2f(\mathcal{B}) - \bar{V}_c^{\bar{\pi}}(\rho) \leq 0$ and hence, $\bar{V}_c^{\bar{\pi}}(\rho) \geq b - 2f(\mathcal{B}) \geq b - 5f(\mathcal{B})$, which completes the proof.

The second case is when $b' - \bar{V}_c^{\bar{\pi}}(\rho) > 0$. In this case, using the notation $[x]_+ = \max\{x, 0\}$ and Eq. (6) with $\lambda = U$, we have

$$V_r^{\pi^*}(\rho) - \bar{V}_r^{\bar{\pi}}(\rho) + U [b' - \bar{V}_c^{\bar{\pi}}(\rho)]_+ \leq \frac{R(U, K)}{K} + 2f(\mathcal{B}).$$

Since U has been set such that $U > \lambda^*$, we can use Lemma 64 and obtain that,

$$[b' - \bar{V}_c^{\bar{\pi}}(\rho)]_+ \leq \frac{R(U, K)}{K(U - \lambda^*)} + \frac{2f(\mathcal{B})}{U - \lambda^*}$$

Combining the above inequality with Eq. (7) gives

$$b' - \bar{V}_c^{\bar{\pi}}(\rho) \leq [b' - \bar{V}_c^{\bar{\pi}}(\rho)]_+ \leq \frac{U}{(U - \lambda^*)(1 - \gamma)\sqrt{K}} + \frac{2f(\mathcal{B})}{U - \lambda^*}. \quad (9)$$

By Lemma 65, we know $\lambda^* \leq \frac{1}{\zeta(1-\gamma)}$. By letting $U = \frac{2}{\zeta(1-\gamma)}$, we have $U - \lambda^* \geq \frac{1}{\zeta(1-\gamma)} \geq 1$ as the Slater constant $\zeta \in (0, \frac{1}{1-\gamma}]$. Thus, $\frac{1}{U-\lambda^*} \leq 1$. Now, setting K to be

$$K = \frac{U^2}{[f(\mathcal{B})]^2(1-\gamma)^2}$$

and substituting into Eqs. (8) and (9), we obtain

$$\bar{V}_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - 3f(\mathcal{B}), \quad \text{and} \quad \bar{V}_c^{\bar{\pi}}(\rho) \geq b' - 3f(\mathcal{B}). \quad (10)$$

This establishes the first claim by substituting $b' = b - 2f(\mathcal{B})$.

Next, we prove the second claim. We define $\pi^{*+} \in \operatorname{argmax}_{\pi} V_r^{\pi}(\rho)$ s.t. $V_c^{\pi}(\rho) \geq b + 6f(\mathcal{B})$. From Eq. (11), recall that

$$V_r^{\pi_k^*}(\rho) + \lambda_k V_c^{\pi_k^*}(\rho) - (\hat{V}_r^k(\rho) + \lambda_k \hat{V}_c^k(\rho)) \leq 2f(\mathcal{B})(1 + \lambda_k)$$

As before, using the definition of π_k^* , we have

$$V_r^{\pi_k^*}(\rho) + \lambda_k V_c^{\pi_k^*}(\rho) \geq V_r^{\pi^{*+}}(\rho) + \lambda_k V_c^{\pi^{*+}}(\rho),$$

Therefore, by combining the above inequalities,

$$\begin{aligned} V_r^{\pi^{*+}}(\rho) + \lambda_k V_c^{\pi^{*+}}(\rho) &\leq \hat{V}_r^k(\rho) + \lambda_k \hat{V}_c^k(\rho) + 2f(\mathcal{B})(1 + \lambda_k) \\ \implies V_r^{\pi^{*+}}(\rho) - \hat{V}_r^k(\rho) &\leq \lambda_k (\hat{V}_c^k(\rho) - V_c^{\pi^{*+}}(\rho) + 2f(\mathcal{B})) + 2f(\mathcal{B}). \end{aligned} \quad (11)$$

Since $V_c^{\pi^{*+}}(\rho) \geq b + 6f(\mathcal{B})$, we obtain,

$$V_r^{\pi^{*+}}(\rho) - \hat{V}_r^k(\rho) \leq \lambda_k [\hat{V}_c^k(\rho) - (b + 3f(\mathcal{B}))] + 2f(\mathcal{B}).$$

As before, by taking the average, letting $b' = b + 4f(\mathcal{B})$, and adding both sides by the same term $\frac{\lambda}{K} \sum_{k=0}^{K-1} [b' - \hat{V}_c^k(\rho)]$, we obtain that for $\lambda \in [0, U]$,

$$V_r^{\pi^{*+}}(\rho) - \bar{V}_r^{\bar{\pi}}(\rho) + \lambda(b' - \bar{V}_c^{\bar{\pi}}(\rho)) \leq \frac{R(\lambda, K)}{K} + 2f(\mathcal{B}).$$

The remainder of the proof proceeds in the same manner as before. Setting K to be

$$K = \frac{U^2}{[f(\mathcal{B})]^2(1-\gamma)^2}$$

the algorithm ensures that

$$\bar{V}_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^{*+}}(\rho) - 3f(\mathcal{B}), \quad \text{and} \quad \bar{V}_c^{\bar{\pi}}(\rho) \geq b' - 3f(\mathcal{B}). \quad (12)$$

This establishes the second claim by substituting $b' = b + 4f(\mathcal{B})$. ■

Appendix D. Proofs for Section 4

The proofs in Section D.1, Section D.2 and Section D.3 are adapted from Kitamura et al. (2023); Kozuno et al. (2022) with modifications to fit our setting. Specifically, the analysis in Kitamura et al. (2023) applies to the *non-stationary policies* returned by MDVI with entropy regularization. In contrast, our analysis applies to the *stationary policy* returned by MDVI without entropy regularization. Furthermore, we also require additional analysis of the value functions returned by the LS-PE algorithm.

Throughout, we treat π as an operator that returns an $|\mathcal{S}|$ -dimensional vector s.t. for an arbitrary $|\mathcal{S}||\mathcal{A}|$ -dimensional vector u such that $(\pi u)(s) := \sum_{a \in \mathcal{A}} \pi(a|s) u(s, a)$. Furthermore, we define $P_\pi := \pi P$ where $P_\pi \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ and denotes the transition probability matrix induced by policy π .

D.1. Deriving LS-MDVI from Entropic Mirror Descent

We show that the LS-MDVI update can be derived as a limiting case of entropic mirror descent. At iteration t , given Q_t , if κ is the entropy regularization parameter and τ is the KL regularization parameter, then, the entropic mirror descent policy update Kitamura et al. (2023) is:

$$\pi_t(\cdot|s) = \arg \max_{p \in \Delta(\mathcal{A})} \sum_{a \in \mathcal{A}} p(a) \left(Q^t(s, a) - \tau \log \frac{p(a)}{\pi_{t-1}(a|s)} - \kappa \log p(a) \right), \quad \text{for all } s \in \mathcal{S},$$

The above policy update can be rewritten in a closed-form solution as follows (Kozuno et al., 2019, Equation 5)),

$$\begin{aligned} \pi_t(a|s) &= \frac{[\pi_{t-1}(a|s)]^\alpha \exp(\beta Q^t(s, a))}{\sum_{b \in \mathcal{A}} [\pi_{t-1}(b|s)]^\alpha \exp(\beta Q^t(s, b))}, \quad \text{where } \alpha := \tau/(\tau + \kappa), \beta := 1/(\tau + \kappa) \\ \implies \pi_t(a|s) &= \frac{\exp(\beta \sum_{i=0}^t \alpha^{t-i} Q^i(s, a))}{\sum_{b \in \mathcal{A}} \exp(\beta \sum_{i=0}^t \alpha^{t-i} Q^i(s, b))}. \end{aligned}$$

Since LS-MDVI does not use entropy regularization $\kappa = 0$ implying $\alpha = 1$, the resulting update is:

$$\pi_t(a|s) = \frac{\exp(\beta \sum_{i=0}^t Q^i(s, a))}{\sum_{b \in \mathcal{A}} \exp(\beta \sum_{i=0}^t Q^i(s, b))} = \frac{1}{1 + \sum_{b \neq a} \exp(\beta(\bar{Q}^t(s, b) - \bar{Q}^t(s, a)))} \quad (\text{where } \bar{Q}^t := \sum_{i=0}^t Q^i)$$

For LS-MDVI, we take the limit $\tau \rightarrow 0, \beta \rightarrow \infty$ and consider two cases.

Case 1: If $a = \arg \max_b \bar{Q}^t(s, b)$, then, $\beta(\bar{Q}^t(s, b) - \bar{Q}^t(s, a)) < 0$ for all $b \neq a$. Hence, as $\beta \rightarrow \infty, \sum_{b \neq a} \exp(\beta(\bar{Q}^t(s, b) - \bar{Q}^t(s, a))) \rightarrow 0$ and $\pi_t(a|s) \rightarrow 1$.

Case 2: If $a \neq \arg \max_b \bar{Q}^t(s, b)$, then, $\beta(\bar{Q}^t(s, b) - \bar{Q}^t(s, a)) > 0$ for the action b corresponding to the arg max action. Hence, as $\beta \rightarrow \infty, \sum_{b \neq a} \exp(\beta(\bar{Q}^t(s, b) - \bar{Q}^t(s, a))) \rightarrow \infty$ and $\pi_t(a|s) \rightarrow 0$.

Hence, as $\kappa = 0$ and $\tau \rightarrow 0, \pi_t$ is a greedy policy and for all $s \in \mathcal{S}, \pi_t(a|s) = 1$ for $a = \arg \max_b \sum_{i=0}^t Q^i(s, b)$, which recovers the policy update for LS-MDVI.

For entropic mirror descent, the value update is given as (Kitamura et al., 2023), for all $s \in \mathcal{S}$,

$$V^t(s) = \sum_a \pi_t(a|s) \left(Q^t(s, a) - \tau \log \left(\frac{\pi_t(a|s)}{\pi_{t-1}(a|s)} \right) - \kappa \ln(\pi_t(a|s)) \right)$$

$$= (\pi_t Q^t)(s) - \tau \text{KL}(\pi_t(\cdot|s) \|\pi_{t-1}(\cdot|s)) + \kappa \mathcal{H}(\pi_t(\cdot|s)).$$

Plugging the entropic mirror descent policy update and simplifying similar to (Kozuno et al., 2022, App. B), we get,

$$\begin{aligned} V^t(s) &= \frac{1}{\beta} \log \sum_{a \in \mathcal{A}} \exp(\beta Q^t(s, a) + \alpha \log \pi_{t-1}(a|s)) \\ &= \frac{1}{\beta} \log \sum_{a \in \mathcal{A}} \exp\left(\beta \sum_{i=0}^t \alpha^{t-i} Q^i(s, a)\right) - \frac{\alpha}{\beta} \log \sum_{a \in \mathcal{A}} \exp\left(\beta \sum_{i=0}^{t-1} \alpha^{t-i} Q^i(s, a)\right). \end{aligned}$$

Since LS-MDVI does not use entropy regularization i.e. $\kappa = 0$ implying $\alpha = 1$, the update is:

$$V^t(s) = \frac{1}{\beta} \log \sum_{a \in \mathcal{A}} \exp\left(\beta \sum_{i=0}^t Q^i(s, a)\right) - \frac{1}{\beta} \log \sum_{a \in \mathcal{A}} \exp\left(\beta \sum_{i=0}^{t-1} Q^i(s, a)\right).$$

For LS-MDVI, we take the limit $\tau \rightarrow 0, \beta \rightarrow \infty$. Using L'Hôpital's rule for the two terms, we get that,

$$\begin{aligned} V^t(s) &= \sum_a \pi_t(a|s) \sum_{i=0}^t Q^i(s, a) - \sum_a \pi_{t-1}(a|s) \sum_{i=0}^{t-1} Q^i(s, a) \\ &= \left(\pi_t \sum_{i=0}^t Q^i\right)(s) - \left(\pi_{t-1} \sum_{i=0}^{t-1} Q^i\right)(s), \end{aligned}$$

which recovers the value update for LS-MDVI.

D.2. Proof of Lemma 6 (Optimality Guarantees for Algorithm 2 - Linear CMDP)

Note that for each λ_k where $k \in [K]$, we run Algorithm 2 with $\square = r + \lambda_k c$. We define

$$\pi_k^* := \arg \max_{\pi} V_{r+\lambda_k c}^{\pi} \quad (13)$$

$$\bar{V}_{\square}^T := \frac{1}{T} \sum_{i=1}^T \hat{V}_{\square}^i \stackrel{\text{(by telescoping)}}{=} \frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) \stackrel{\text{(by definition)}}{=} \frac{1}{T} \left(\pi_T \left\langle \phi, \sum_{i=0}^T \theta_{\square}^i \right\rangle \right). \quad (14)$$

Throughout the proof, for any $|\mathcal{S}||\mathcal{A}|$ -dimensional vector z , we let $W(z)$ denote the solution to a weighted linear regression problem over the core set,

$$W(z) := \arg \min_{\theta} \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x, b) (z(x, b) - \langle \phi(x, b), \theta \rangle)^2. \quad (15)$$

The above problem can be solved as

$$W(z) = G^{-1} \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x, b) \phi(x, b) z(x, b) \quad (16)$$

where $G := \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x, b) \phi(x, b) \phi(x, b)^{\top}$.

Using this definition and the definition of θ_\square^i in Alg. 2, we have $\theta_\square^i = W(\hat{Q}_\square^i)$.

The linear MDP assumption ensures that there exists a vector θ_\square^t such that $\langle \phi, \theta_\square^t \rangle := \square + \gamma P \hat{V}_\square^{t-1}$. Therefore, using the definition of W , we have $\theta_\square^i = W(\langle \phi, \theta_\square^i \rangle)$.

We now present the proof of Lemma 6.

Lemma 12 *For a fixed $\varepsilon \in (0, 1]$, $\delta \in (0, 1)$, and any $k \in [K]$, when using Alg. 2 at iteration k of Alg. 1 with $\square = r + \lambda_k c$, $M = \tilde{O}\left(\frac{dH^2}{\varepsilon}\right)$ and $T = O\left(\frac{H^2}{\varepsilon}\right)$, the output policy π_T satisfies the following condition with probability $1 - \delta$,*

$$\max_{\pi} V_{r+\lambda_k c}^{\pi}(\rho) - V_{r+\lambda_k c}^{\pi_T}(\rho) \leq O((1 + \lambda_k)\varepsilon)$$

Proof Using the definition of π_k^* and that $\square = r + \lambda_k c$, we decompose the sub-optimality as:

$$V_{r+\lambda_k c}^{\pi_k^*}(\rho) - V_{r+\lambda_k c}^{\pi_T}(\rho) = [V_{r+\lambda_k c}^{\pi_k^*}(\rho) - \bar{V}_\square^T(\rho)] + [\bar{V}_\square^T(\rho) - V_{r+\lambda_k c}^{\pi_T}(\rho)]$$

Bounding the first term by Lemma 13 and the second by Lemma 14,

$$\leq \tilde{O}\left(\frac{H^2(1 + \lambda_k)}{T} + H^2(1 + \lambda_k)\sqrt{\frac{d}{TM}}\right)$$

with probability at least $1 - 2\delta$. Setting $M = \tilde{O}\left(\frac{dH^2}{\varepsilon}\right)$, $T = O\left(\frac{H^2}{\varepsilon}\right)$, and appropriately rescaling the confidence parameter δ completes the proof. \blacksquare

We now prove Thms. 13 and 14.

Lemma 13 *Let π_k^* and \bar{V}_\square^T be defined as in Eqs. (13) and (14). For any $k \in [K]$, with $\square = r + \lambda_k c$ and $M \geq \tilde{O}(dH^2)$, we have*

$$V_{r+\lambda_k c}^{\pi_k^*}(\rho) - \bar{V}_\square^T(\rho) \leq \tilde{O}\left(\frac{H^2(1 + \lambda_k)}{T} + H^2(1 + \lambda_k)\sqrt{\frac{d}{TM}}\right)$$

with probability at least $1 - \delta$.

Proof We first recall that $V_\square^{\pi_k^*} = V_{r+\lambda_k c}^{\pi_k^*}$ and $\bar{V}_\square^T = \bar{V}_{r+\lambda_k c}^T$ by the definition of \square . By the value difference lemma, we have that,

$$V_\square^{\pi_k^*} - \bar{V}_\square^T = (I - \gamma P_{\pi_k^*})^{-1}((\pi_k^* \square) + \gamma P_{\pi_k^*} \bar{V}_\square^T - \bar{V}_\square^T) \quad (17)$$

Next, from Line 6 in Algorithm 2, by the telescoping sum, and by the greediness of π_T , we have

$$\bar{V}_\square^T = \frac{1}{T} (\pi_T \tilde{Q}_\square^T) \quad (18)$$

$$\geq \frac{1}{T} (\pi_k^* \tilde{Q}_\square^T). \quad (19)$$

Now, we have

$$V_\square^{\pi_k^*} - \bar{V}_\square^T = (I - \gamma P_{\pi_k^*})^{-1}((\pi_k^* \square) + \gamma P_{\pi_k^*} \bar{V}_\square^T - \bar{V}_\square^T) \quad (\text{By Eq. (17)})$$

$$\begin{aligned}
 &\leq (I - \gamma P_{\pi_k^*})^{-1}((\pi_k^* \square) + \gamma P_{\pi_k^*} \bar{V}_\square^T - \frac{1}{T} (\pi_k^* \tilde{Q}_\square^T)) && \text{(By Eq. (19))} \\
 &= (I - \gamma P_{\pi_k^*})^{-1}((\pi_k^* \square) + \gamma P_{\pi_k^*} \frac{1}{T} (\pi_T \tilde{Q}_\square^T) - \frac{1}{T} (\pi_k^* \tilde{Q}_\square^T)) && \text{(By Eq. (18))} \\
 &= (I - \gamma P_{\pi_k^*})^{-1} \left[(\pi_k^* \square) + \gamma P_{\pi_k^*} \frac{1}{T} (\pi_T \tilde{Q}_\square^T) - (\pi_k^* \square) - \gamma P_{\pi_k^*} \frac{1}{T} (\pi_{T-1} \tilde{Q}_\square^{T-1}) \right. \\
 &\quad \left. - \left(\pi_k^* \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right) \right] \quad \text{(Using Lemma 21 for } \frac{1}{T} \tilde{Q}_\square^T) \\
 &= (I - \gamma P_{\pi_k^*})^{-1} \left[\gamma P_{\pi_k^*} \frac{1}{T} (\pi_T \tilde{Q}_\square^T) - \gamma P_{\pi_k^*} \frac{1}{T} (\pi_{T-1} \tilde{Q}_\square^{T-1}) \right. \\
 &\quad \left. - \left(\pi_k^* \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right) \right].
 \end{aligned}$$

By defining $\mathcal{H}_{\pi_k^*} := (I - \gamma P_{\pi_k^*})^{-1}$, taking the infinity norm and using the triangle inequality, we obtain

$$\begin{aligned}
 \left\| V_\square^{\pi_k^*} - \bar{V}_\square^T \right\|_\infty &\leq \underbrace{\left\| \gamma \mathcal{H}_{\pi_k^*} P_{\pi_k^*} \left(\frac{1}{T} (\pi_T \tilde{Q}_\square^T) - \frac{1}{T} (\pi_{T-1} \tilde{Q}_\square^{T-1}) \right) \right\|_\infty}_{\text{Term (i)}} \\
 &\quad + \underbrace{\left\| \mathcal{H}_{\pi_k^*} \left(\pi_k^* \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right) \right\|_\infty}_{\text{Term (ii)}}. \quad (20)
 \end{aligned}$$

In order to bound Term (i), we use Holder's inequality i.e. for a matrix A and vector x , $\|Ax\|_\infty \leq \|A\|_{1,\infty} \|x\|_\infty$, and that $\|\mathcal{H}_{\pi_k^*} P_{\pi_k^*}\|_{1,\infty} \leq H$ to obtain,

$$\begin{aligned}
 \left\| \gamma \mathcal{H}_{\pi_k^*} P_{\pi_k^*} \left(\frac{1}{T} (\pi_T \tilde{Q}_\square^T) - \frac{1}{T} (\pi_{T-1} \tilde{Q}_\square^{T-1}) \right) \right\|_\infty &\leq H \left\| \left(\frac{1}{T} (\pi_T \tilde{Q}_\square^T) - \frac{1}{T} (\pi_{T-1} \tilde{Q}_\square^{T-1}) \right) \right\|_\infty \\
 &\leq \frac{4H^2(1 + \lambda_k)}{T} \quad \text{(Using Thm. 17)}
 \end{aligned}$$

with probability at least $1 - \delta$. For term (ii),

$$\begin{aligned}
 &\left\| \mathcal{H}_{\pi_k^*} \left(\pi_k^* \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right) \right\|_\infty \\
 &\leq \left\| \mathcal{H}_{\pi_k^*} \right\|_{1,\infty} \left\| \left(\pi_k^* \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right) \right\|_\infty \quad \text{(By Holder's inequality)} \\
 &\leq H \left\| \left(\pi_k^* \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right) \right\|_\infty \quad \text{(Since } \|\mathcal{H}_{\pi_k^*}\|_{1,\infty} \leq H) \\
 &\leq H \left\| \phi^\top W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\|_\infty \quad \text{(By definition of the } \pi \text{ operator)}
 \end{aligned}$$

$$\leq \tilde{O} \left(H^2(1 + \lambda_k) \sqrt{\frac{d}{TM}} \right) \quad (\text{By Lemma 18})$$

Combining the above relations,

$$\left\| V_{\square}^{\pi_k^*} - \bar{V}_{\square}^T \right\|_{\infty} \leq \frac{4H^2(1 + \lambda_k)}{T} + \tilde{O} \left(H^2(1 + \lambda_k) \sqrt{\frac{d}{TM}} \right)$$

Using that for any $|\mathcal{S}|$ -dimensional vector V , $V(\rho) = \mathbb{E}_{s \sim \rho} V(s) \leq \|V\|_{\infty}$, we get that,

$$V_{r+\lambda_k c}^{\pi_k^*}(\rho) - \bar{V}_{\square}^T(\rho) \leq \frac{4H^2(1 + \lambda_k)}{T} + \tilde{O} \left(H^2(1 + \lambda_k) \sqrt{\frac{d}{TM}} \right)$$

with probability at least $1 - \delta$. ■

Lemma 14 *Let \bar{V}_{\square}^T be defined as in Eq. (14). For any $k \in [K]$, with $\square = r + \lambda_k c$ and $M \geq \tilde{O}(dH^2)$, we have*

$$\bar{V}_{\square}^T(\rho) - V_{r+\lambda_k c}^{\pi_T}(\rho) \leq \tilde{O} \left(\frac{H^2(1 + \lambda_k)}{T} + H^2(1 + \lambda_k) \sqrt{\frac{d}{TM}} \right)$$

with probability at least $1 - \delta$.

Proof The proof is similar as for the above lemma. By the value difference lemma, we have that,

$$\bar{V}_{\square}^T - V_{r+\lambda_k c}^{\pi_T} = (I - \gamma P_{\pi_T})^{-1} (\bar{V}_{\square}^T - (\pi_T \square) - \gamma P_{\pi_T} \bar{V}_{\square}^T) \quad (21)$$

Now, we have

$$\begin{aligned} \bar{V}_{\square}^T - V_{r+\lambda_k c}^{\pi_T} &= (I - \gamma P_{\pi_T})^{-1} \left(\frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) - (\pi_T \square) - \gamma P_{\pi_T} \bar{V}_{\square}^T \right) \\ &= (I - \gamma P_{\pi_T})^{-1} \left(\frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) - (\pi_T \square) - \gamma P_{\pi_T} \frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) \right) \\ &= (I - \gamma P_{\pi_T})^{-1} \left[(\pi_T \square) + \gamma P_{\pi_T} \frac{1}{T} (\pi_{T-1} \tilde{Q}_{\square}^{T-1}) + \left(\pi_T \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_{\square}^i - \langle \phi, \theta_{\square}^i \rangle) \right) \right\rangle \right) \right. \\ &\quad \left. - (\pi_T \square) - \gamma P_{\pi_T} \frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) \right] \quad (\text{Using Lemma 21 for } \frac{1}{T} \tilde{Q}^T) \\ &= (I - \gamma P_{\pi_T})^{-1} \left[\gamma P_{\pi_T} \frac{1}{T} (\pi_{T-1} \tilde{Q}_{\square}^{T-1}) - \gamma P_{\pi_T} \frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) \right. \\ &\quad \left. + \left(\pi_T \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_{\square}^i - \langle \phi, \theta_{\square}^i \rangle) \right) \right\rangle \right) \right]. \end{aligned}$$

By defining $\mathcal{H}_{\pi_T} := (I - \gamma P_{\pi_T})^{-1}$, taking the infinity norm and using the triangle inequality, we obtain

$$\begin{aligned} \left\| \bar{V}_{\square}^T - V_{r+\lambda_k c}^{\pi_T} \right\|_{\infty} &\leq \underbrace{\left\| \gamma \mathcal{H}_{\pi_T} P_{\pi_T} \left(\frac{1}{T} (\pi_{T-1} \tilde{Q}_{\square}^{T-1}) - \frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) \right) \right\|_{\infty}}_{\text{Term (i)}} \\ &\quad + \underbrace{\left\| \mathcal{H}_{\pi_T} \left(\pi_T \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_{\square}^i - \langle \phi, \theta_{\square}^i \rangle) \right) \right\rangle \right) \right\|_{\infty}}_{\text{Term (ii)}}. \end{aligned} \quad (22)$$

In order to bound Term (i), we use Holder's inequality and that $\|\mathcal{H}_{\pi_T} P_{\pi_T}\|_{1,\infty} \leq H$,

$$\begin{aligned} \left\| \gamma \mathcal{H}_{\pi_T} P_{\pi_T} \left(\frac{1}{T} (\pi_{T-1} \tilde{Q}_{\square}^{T-1}) - \frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) \right) \right\|_{\infty} &\leq H \left\| \frac{1}{T} (\pi_{T-1} \tilde{Q}_{\square}^{T-1}) - \frac{1}{T} (\pi_T \tilde{Q}_{\square}^T) \right\|_{\infty} \\ &\leq \frac{4H^2(1 + \lambda_k)}{T} \quad (\text{Using Thm. 17}) \end{aligned}$$

with probability at least $1 - \delta$. For term (ii),

$$\begin{aligned} &\left\| \mathcal{H}_{\pi_T} \left(\pi_T \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_{\square}^i - \langle \phi, \theta_{\square}^i \rangle) \right) \right\rangle \right) \right\|_{\infty} \\ &\leq \|\mathcal{H}_{\pi_T}\|_{1,\infty} \left\| \left(\pi_T \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_{\square}^i - \langle \phi, \theta_{\square}^i \rangle) \right) \right\rangle \right) \right\|_{\infty} \quad (\text{By Holder's inequality}) \\ &\leq H \left\| \left(\pi_T \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_{\square}^i - \langle \phi, \theta_{\square}^i \rangle) \right) \right\rangle \right) \right\|_{\infty} \quad (\text{Since } \|\mathcal{H}_{\pi_T}\|_{1,\infty} \leq H) \\ &\leq H \left\| \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_{\square}^i - \langle \phi, \theta_{\square}^i \rangle) \right) \right\rangle \right\|_{\infty} \quad (\text{By definition of the } \pi \text{ operator}) \\ &\leq \tilde{O} \left(H^2(1 + \lambda_k) \sqrt{\frac{d}{TM}} \right) \quad (\text{By Lemma 18}) \end{aligned}$$

Combining the above relations and using that for any $|\mathcal{S}|$ -dimensional vector V , $V(\rho) = \mathbb{E}_{s \sim \rho} V(s) \leq \|V\|_{\infty}$, we get that,

$$\bar{V}_{\square}^T(\rho) - V_{r+\lambda_k c}^{\pi_T}(\rho) \leq \frac{4H^2(1 + \lambda_k)}{T} + \tilde{O} \left(H^2(1 + \lambda_k) \sqrt{\frac{d}{TM}} \right)$$

with probability at least $1 - \delta$. ■

D.2.1. AUXILIARY LEMMAS

Lemma 15 *For any $k \in [K]$ and $t \in [T]$, with $\square = r + \lambda_k c$ and $M \geq \tilde{O}(dH^2)$, we have*

$$\|\hat{V}_{\square}^t\|_{\infty} \leq 2H(1 + \lambda_k)$$

with probability at least $1 - \delta$.

Proof First, we note that from the last line in Algorithm 2,

$$\begin{aligned}
 \hat{V}_\square^t &= (\pi_t \tilde{Q}_\square^t) - (\pi_{t-1} \tilde{Q}_\square^{t-1}) \\
 &= \left(\pi_t \left\langle \phi, \sum_{i=0}^t \theta_\square^i \right\rangle \right) - \left(\pi_{t-1} \left\langle \phi, \sum_{i=0}^{t-1} \theta_\square^i \right\rangle \right) \\
 &\leq \left(\pi_t \left\langle \phi, \sum_{i=0}^t \theta_\square^i \right\rangle \right) - \left(\pi_t \left\langle \phi, \sum_{i=0}^{t-1} \theta_\square^i \right\rangle \right) \quad (\text{By the greediness of } \pi_{t-1}) \\
 &= (\pi_t \langle \phi, \theta_\square^t \rangle). \tag{23}
 \end{aligned}$$

Next, we bound the term $\langle \phi, \theta_\square^t \rangle$. We have

$$\begin{aligned}
 |\langle \phi, \theta_\square^t \rangle| &= \left| \langle \phi, W(\hat{Q}_\square^t) \rangle \right| \quad (\text{By the definition of } W \text{ in Eq. (16)}) \\
 &\leq \left| \langle \phi, W(\langle \phi, \theta_\square^t \rangle) \rangle \right| + \left| \langle \phi, W(\hat{Q}_\square^t) - W(\langle \phi, \theta_\square^t \rangle) \rangle \right| \quad (\text{By triangle inequality}) \\
 &= \left| \langle \phi, \theta_\square^t \rangle \right| + \left| \langle \phi, W(\hat{Q}_\square^t - \langle \phi, \theta_\square^t \rangle) \rangle \right| \quad (\text{Since } W(z) \text{ is linear in } z) \\
 &= \left| \square + \gamma P \hat{V}_\square^{t-1} \right| + \left| \langle \phi, W(\hat{Q}_\square^t - \langle \phi, \theta_\square^t \rangle) \rangle \right| \quad (\text{By the definition of } \theta_\square^t) \\
 &\leq (1 + \lambda_k + \gamma \|\hat{V}_\square^{t-1}\|_\infty) \mathbf{1} + \left| \langle \phi, W(\hat{Q}_\square^t - \langle \phi, \theta_\square^t \rangle) \rangle \right| \quad (\text{Since } \square(s, a) \leq 1 + \lambda_k) \\
 &\leq (1 + \lambda_k + \gamma \|\hat{V}_\square^{t-1}\|_\infty) \mathbf{1} + \sqrt{2d} \max_{(s,a) \in \mathcal{C}} \left| \hat{Q}_\square^t(s, a) - (\langle \phi, \theta_\square^t \rangle)(s, a) \right| \mathbf{1} \\
 &\quad (\text{By Lemma 22 with } z = \hat{Q}_\square^t - \langle \phi, \theta_\square^t \rangle) \\
 &= (1 + \lambda_k + \gamma \|\hat{V}_\square^{t-1}\|_\infty) \mathbf{1} \\
 &\quad + \sqrt{2d} \max_{(s,a) \in \mathcal{C}} \left| \square(s, a) + \gamma(\hat{P}_{t-1} \hat{V}_\square^{t-1})(s, a) - \square(s, a) - \gamma(P \hat{V}_\square^{t-1})(s, a) \right| \mathbf{1} \\
 &\quad (\text{By definition of } \hat{Q}_\square^t \text{ and } \theta_\square^t) \\
 &= (1 + \lambda_k + \gamma \|\hat{V}_\square^{t-1}\|_\infty) \mathbf{1} + \sqrt{2d} \max_{(s,a) \in \mathcal{C}} \left| \gamma(\hat{P}_{t-1} \hat{V}_\square^{t-1})(s, a) - \gamma(P \hat{V}_\square^{t-1})(s, a) \right| \mathbf{1} \\
 \implies |\langle \phi, \theta_\square^t \rangle| &\leq (1 + \lambda_k + \gamma \|\hat{V}_\square^{t-1}\|_\infty) \mathbf{1} + \sqrt{2d} \frac{\|\hat{V}_\square^{t-1}\|_\infty}{\|\hat{V}_\square^{t-1}\|_\infty} \max_{(s,a) \in \mathcal{C}} \left| \gamma(\hat{P}_{t-1} \hat{V}_\square^{t-1})(s, a) - \gamma(P \hat{V}_\square^{t-1})(s, a) \right| \mathbf{1}
 \end{aligned}$$

Next, we bound the term

$$\frac{1}{\|\hat{V}_\square^{t-1}\|_\infty} \max_{(s,a) \in \mathcal{C}} \left| \gamma(\hat{P}_{t-1} \hat{V}_\square^{t-1})(s, a) - \gamma(P \hat{V}_\square^{t-1})(s, a) \right|.$$

We first note that this term is upper bounded by 2. Now, using the Azuma-Hoeffding inequality (Lemma 58) and taking a union bound over $(s, a) \in \mathcal{C}$ and $t \in [T]$, we have

$$\mathbb{P} \left(\exists (s, a, t) \in \mathcal{C} \times [T] \text{ s.t. } \frac{1}{\|\hat{V}_\square^{t-1}\|_\infty} \max_{(s,a) \in \mathcal{C}} \left| \gamma(\hat{P}_{t-1} \hat{V}_\square^{t-1})(s, a) - \gamma(P \hat{V}_\square^{t-1})(s, a) \right| \geq \tilde{O} \left(\gamma \sqrt{\frac{1}{M}} \right) \right) \leq \delta.$$

Therefore, with probability at least $1 - \delta$, we have

$$|\langle \phi, \theta_\square^t \rangle| \leq (1 + \lambda_k + \gamma \|\hat{V}_\square^{t-1}\|_\infty) \mathbf{1} + \|\hat{V}_\square^{t-1}\|_\infty \tilde{O} \left(\gamma \sqrt{\frac{d}{M}} \right) \mathbf{1}. \tag{24}$$

Given the above inequality, we can prove the claim by induction on t . Since $\hat{V}_\square^0 = 0$, the base case is satisfied. We assume that $\|\hat{V}_\square^{t-1}\|_\infty \leq 2H(1 + \lambda_k)$. By combining Eq. (23) and Eq. (24), we have

$$\begin{aligned}
 \|\hat{V}_\square^t\|_\infty &\leq \|(\pi_t \langle \phi, \theta_\square^t \rangle)\|_\infty \\
 &\leq \|\langle \phi, \theta_\square^t \rangle\|_\infty && \text{(By definition of the } \pi \text{ operator)} \\
 &\leq \left(1 + \lambda_k + \gamma \|\hat{V}_\square^{t-1}\|_\infty + \|\hat{V}_\square^{t-1}\|_\infty \tilde{O}\left(\gamma \sqrt{\frac{d}{M}}\right)\right) \\
 &\leq \left(1 + 2H\gamma + 2H\tilde{O}\left(\gamma \sqrt{\frac{d}{M}}\right)\right) (1 + \lambda_k). && \text{(Induction hypothesis)}
 \end{aligned}$$

By taking $M \geq \tilde{O}(dH^2)$, we have

$$\begin{aligned}
 \|\hat{V}_\square^t\|_\infty &\leq (1 + 2H\gamma + 1)(1 + \lambda_k) \\
 &= (2 + 2H\gamma)(1 + \lambda_k) \\
 &\leq 2H(1 + \lambda_k) && \text{(Since } H = 1/(1 - \gamma))
 \end{aligned}$$

which completes the proof. ■

The following corollary is a direct consequence of Eq. (24) and the above lemma.

Corollary 16 *For any $k \in [K]$ and $t \in [T]$, with $\square = r + \lambda_k c$ and $M \geq \tilde{O}(dH^2)$, we have*

$$\|\langle \phi, \theta_\square^t \rangle\|_\infty \leq 2H(1 + \lambda_k)$$

with probability at least $1 - \delta$ respectively.

Lemma 17 *For any $k \in [K]$, with $M \geq \tilde{O}(dH^2)$, we have*

$$\left\| \frac{1}{T}(\pi_T \tilde{Q}_\square^T) - \frac{1}{T}(\pi_{T-1} \tilde{Q}_\square^{T-1}) \right\|_\infty \leq \frac{2H(1 + \lambda_k)}{T}$$

with probability at least $1 - \delta$.

Proof By the definition of \tilde{Q}_\square^t and due to the greediness of π_{T-1} , we have

$$\begin{aligned}
 \frac{1}{T}(\pi_T \tilde{Q}_\square^T) - \frac{1}{T}(\pi_{T-1} \tilde{Q}_\square^{T-1}) &\leq \frac{1}{T}(\pi_T \tilde{Q}_\square^T) - \frac{1}{T}(\pi_T \tilde{Q}_\square^{T-1}) \\
 &= \left(\pi_T \left\langle \phi, \frac{1}{T} \sum_{i=0}^T \theta_\square^i - \frac{1}{T} \sum_{i=0}^{T-1} \theta_\square^i \right\rangle \right) \\
 &= \frac{1}{T}(\pi_T \langle \phi, \theta_\square^T \rangle) \\
 &\leq \frac{1}{T} \|\langle \phi, \theta_\square^T \rangle\|_\infty \mathbf{1} && \text{(By definition of the } \pi \text{ operator)} \\
 &\leq \frac{2H(1 + \lambda_k)}{T} \mathbf{1} && \text{(By Corollary 16)}
 \end{aligned}$$

with probability at least $1 - \delta$. Similarly, by the greediness of π_T , we have

$$\begin{aligned}
 \frac{1}{T}(\pi_{T-1}\tilde{Q}_\square^{T-1}) - \frac{1}{T}(\pi_T\tilde{Q}_\square^T) &\leq \frac{1}{T}(\pi_{T-1}\tilde{Q}_\square^{T-1}) - \frac{1}{T}(\pi_{T-1}\tilde{Q}_\square^T) \\
 &= \left(\pi_{T-1} \left\langle \phi, \frac{1}{T} \sum_{i=0}^{T-1} \theta_\square^i - \frac{1}{T} \sum_{i=0}^T \theta_\square^i \right\rangle \right) \\
 &= -\frac{1}{T}(\pi_T \langle \phi, \theta_\square^T \rangle) \leq \frac{1}{T} \|(\pi_T \langle \phi, \theta_\square^T \rangle)\|_\infty \\
 &\leq \frac{1}{T} \|\langle \phi, \theta_\square^T \rangle\|_\infty \mathbf{1} \quad (\text{By definition of the } \pi \text{ operator}) \\
 &\leq \frac{2H(1 + \lambda_k)}{T} \mathbf{1} \quad (\text{By Corollary 16})
 \end{aligned}$$

with probability at least $1 - \delta$. ■

Lemma 18 For any $k \in [K]$ and $t \in [T]$, with $\square = r + \lambda_k c$ and $M \geq \tilde{O}(dH^2)$, we have

$$\left\| \left\langle \phi, W \left(\frac{1}{t} \sum_{i=1}^t (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right\|_\infty \leq \tilde{O} \left(H(1 + \lambda_k) \sqrt{\frac{d}{tM}} \right)$$

with probability at least $1 - \delta$.

Proof

$$\begin{aligned}
 &\left\| \left\langle \phi, W \left(\frac{1}{t} \sum_{i=1}^t (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right\|_\infty \\
 &\leq \sqrt{2d} \max_{(s,a) \in \mathcal{C}} \left| \frac{1}{t} \sum_{i=0}^{t-1} [\hat{Q}_\square^i(s,a) - (\langle \phi, \theta_\square^i \rangle)(s,a)] \right| \\
 &\quad \quad \quad (\text{By Lemma 22 with } z = \frac{1}{t} \sum_{i=0}^{t-1} [\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle]) \\
 &= \sqrt{2d} \max_{(s,a) \in \mathcal{C}} \left| \frac{1}{t} \sum_{i=0}^{t-1} [\gamma(\hat{P}_i \hat{V}_\square^i)(s,a) - \gamma(P \hat{V}_\square^i)(s,a)] \right| \quad (\text{By definition of } \hat{Q}_\square^i \text{ and } \theta_\square^i)
 \end{aligned}$$

By Lemma 15, we have that, with probability at least $1 - \delta$, the bound $\|\hat{V}_\square^i\|_\infty \leq 2H(1 + \lambda_k)$ holds for all $i \in [T]$. Now, using Lemma 57 and taking the union bound over $(s, a) \in \mathcal{C}$, we have

$$\mathbb{P} \left(\exists (s, a) \in \mathcal{C} \text{ s.t. } \frac{1}{t} \sum_{i=0}^{t-1} [(\hat{P}_i \hat{V}_\square^i)(s, a) - (P \hat{V}_\square^i)(s, a)] \geq \tilde{O} \left(H(1 + \lambda_k) \sqrt{\frac{1}{tM}} \right) \right) \leq \delta.$$

Therefore, by appropriately rescaling δ , we have that with probability at least $1 - \delta$,

$$\left\| \left\langle \phi, W \left(\frac{1}{t} \sum_{i=1}^t (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle \right\|_\infty \leq \tilde{O} \left(H(1 + \lambda_k) \sqrt{\frac{d}{tM}} \right).$$

■

Lemma 19 For any $k \in [K]$ and $i \in [T]$, with $\square = r + \lambda_k c$ and $M \geq \tilde{O}(dH^2)$, we have

$$\left\| \langle \phi, W(\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \rangle \right\|_\infty \leq \tilde{O} \left(H(1 + \lambda_k) \sqrt{\frac{d}{M}} \right)$$

with probability at least $1 - \delta$.

Proof By following a similar proof as that for the above lemma,

$$\left\| \langle \phi, W(\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \rangle \right\|_\infty \leq \sqrt{2d} \max_{(s,a) \in \mathcal{C}} \left| \gamma \hat{P}_i \hat{V}_\square^i(s, a) - \gamma P \hat{V}_\square^i(s, a) \right|. \quad (25)$$

By Lemma 15, we have that, with probability at least $1 - \delta$, $(\hat{P}_i \hat{V}_\square^i)(s, a) \leq 2H(1 + \lambda_k)$ holds for all $i \in [T]$ and all (s, a) . We note that by the definition of \hat{P}_i , $(\hat{P}_i \hat{V}_\square^i)(s, a)$ is the empirical average of M value functions. Now, using Lemma 58 with $N = M$ and taking the union bound over $(s, a) \in \mathcal{C}$ and $i \in [T]$, we have

$$\mathbb{P} \left(\exists (s, a, t) \in \mathcal{C} \times [T] \text{ s.t. } (\hat{P}_i \hat{V}_\square^i)(s, a) - (P \hat{V}_\square^i)(s, a) \geq \tilde{O} \left(H(1 + \lambda_k) \sqrt{\frac{1}{M}} \right) \right) \leq \delta$$

Combining the above inequality with Eq. (25) and appropriately rescaling δ completes the proof. ■

Lemma 20 For any $t \in [T]$, we have

$$\frac{1}{t} \left\langle \phi, \sum_{i=0}^t \theta_\square^i \right\rangle = \square + \gamma \frac{1}{t} P (\pi_{t-1} \tilde{Q}_\square^{t-1}).$$

Proof We first recall that by definition, $\langle \phi, \theta_\square^t \rangle := \square + \gamma P \hat{V}_\square^{t-1}$, $\hat{V}_\square^0 = \mathbf{0}$, and $\theta_\square^0 = \mathbf{0}$. Now we have

$$\begin{aligned} \frac{1}{t} \left\langle \phi, \sum_{i=0}^t \theta_\square^i \right\rangle &:= \frac{1}{t} \sum_{i=0}^{t-1} (\square + \gamma P \hat{V}_\square^i) \\ &= \square + \gamma P \left(\frac{1}{t} \sum_{i=0}^{t-1} \hat{V}_\square^i \right) \\ &= \square + \gamma P \left(\frac{1}{t} \sum_{i=0}^{t-1} [(\pi_i \tilde{Q}_\square^i) - (\pi_{i-1} \tilde{Q}_\square^{i-1})] \right) \\ &\hspace{15em} \text{(From the last line in Algorithm 2)} \\ &= \square + \gamma \frac{1}{t} P (\pi_{t-1} \tilde{Q}_\square^{t-1}). \hspace{10em} \text{(Telescoping Sum)} \end{aligned}$$

■

Lemma 21 *We have*

$$\frac{1}{T} \tilde{Q}_\square^T = \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle + \square + \gamma P \frac{1}{T} (\pi_{T-1} \tilde{Q}_\square^{T-1}).$$

Proof We first recall that by the definition of W , we have $\theta_\square^i = W(\hat{Q}_\square^i)$ and $\theta_\square^i = W(\langle \phi, \theta_\square^i \rangle)$. Thus,

$$\begin{aligned} \frac{1}{T} \tilde{Q}_\square^T &= \frac{1}{T} \tilde{Q}_\square^T - \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i \rangle + \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i \rangle \\ &= \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i \rangle - \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i \rangle + \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i \rangle && \text{(By definition of } \tilde{Q}_\square^T) \\ &= \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i - \theta_\square^i \rangle + \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i \rangle \\ &= \frac{1}{T} \sum_{i=0}^T \left\langle \phi, W(\hat{Q}_\square^i) - W(\langle \phi, \theta_\square^i \rangle) \right\rangle + \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i \rangle \\ & && \text{(Since } \theta_\square^i = W(\hat{Q}_\square^i) \text{ and } \theta_\square^i = W(\langle \phi, \theta_\square^i \rangle)) \\ &= \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle + \frac{1}{T} \sum_{i=0}^T \langle \phi, \theta_\square^i \rangle && \text{(} W(z) \text{ is linear in } z) \\ &= \left\langle \phi, W \left(\frac{1}{T} \sum_{i=0}^T (\hat{Q}_\square^i - \langle \phi, \theta_\square^i \rangle) \right) \right\rangle + \square + \frac{1}{T} \gamma P (\pi_{T-1} \tilde{Q}_\square^{T-1}). \\ & && \text{(By Lemma 20 with } t = T - 1) \end{aligned}$$

■

The following lemma bounds the extrapolation error due to the least-squares regression. It is the unweighted version (i.e., uniform weighting with $f = 1$) of Lemma 4.3 in [Kitamura et al. \(2023\)](#).

Lemma 22 (KW Bound) *Let z be a function defined over \mathcal{C} . Then, there exists $\tilde{\rho} \in \Delta(\mathcal{S} \times \mathcal{A})$ with a finite support $\mathcal{C} := \text{Supp}(\tilde{\rho})$ of size less than or equal to $u_{\mathcal{C}}$ such that*

$$\max_{(s,a) \in \mathcal{S} \times \mathcal{A}} [\langle \phi(s, a), W(z) \rangle] \leq \sqrt{2d} \max_{(x',b') \in \mathcal{C}} |z(x', b')|,$$

where $W(z) := G^{-1} \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x, b) \phi(x, b) z(x, b)$.

D.3. Proof of Lemma 8 (Optimality Guarantees for Algorithm 3 - Linear CMDP)

We define

$$\tilde{\mathcal{V}}_\diamond^T := \frac{1}{T} \sum_{i=1}^T \hat{\mathcal{V}}_\diamond^i \stackrel{(a)}{=} \frac{1}{T} \left(\pi \left\langle \phi, \sum_{i=1}^T \omega_\diamond^i \right\rangle \right) \quad (26)$$

$$\langle \phi, \omega_\diamond^t \rangle := \diamond + \gamma P \hat{\mathcal{V}}_\diamond^{t-1} \quad (27)$$

where (a) is from the last line in Algorithm 3.

Lemma 23 For a fixed $\varepsilon \in (0, 1]$, $\delta \in (0, 1)$, Alg. 3 with $M = \tilde{O}\left(\frac{dH^2}{\varepsilon}\right)$ and $T = O\left(\frac{H^2}{\varepsilon}\right)$, the output \bar{V}_\diamond^T satisfies the following condition with probability $1 - \delta$,

$$|\bar{V}_\diamond^T(\rho) - V_\diamond^\pi(\rho)| \leq O(\varepsilon).$$

Proof Using the value difference lemma,

$$\bar{V}_\diamond^T - V_\diamond^\pi = (I - \gamma P_\pi)^{-1}(\bar{V}_\diamond^T - (\pi \diamond) - \gamma P_\pi \bar{V}_\diamond^T).$$

We now have

$$\begin{aligned} \bar{V}_\diamond^T - V_\diamond^\pi &= (I - \gamma P_\pi)^{-1} \left[\left(\pi \left\langle \phi, \frac{1}{T} \sum_{i=1}^T \omega_\diamond^i \right\rangle \right) - (\pi \diamond) - \gamma P_\pi \bar{V}_\diamond^T \right] \\ &\quad \text{(By definition of } \bar{V}_\diamond^T \text{ in Eq. (26))} \\ &= (I - \gamma P_\pi)^{-1} \left[\left(\pi \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right) + (\pi \diamond) + \gamma P_\pi \left(\pi \left\langle \phi, \frac{1}{T} \sum_{i=1}^{T-1} \omega_\diamond^i \right\rangle \right) \right. \\ &\quad \left. - (\pi \diamond) - \gamma P_\pi \bar{V}_\diamond^T \right] \\ &\quad \text{(By Lemma 25)} \\ &= (I - \gamma P_\pi)^{-1} \left[\left(\pi \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right) + \gamma P_\pi \left(\pi \left\langle \phi, \frac{1}{T} \sum_{i=1}^{T-1} \omega_\diamond^i \right\rangle \right) - \gamma P_\pi \bar{V}_\diamond^T \right] \\ &= (I - \gamma P_\pi)^{-1} \left[\left(\pi \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right) + \gamma P_\pi \left(\pi \left\langle \phi, \frac{1}{T} \sum_{i=1}^{T-1} \omega_\diamond^i \right\rangle \right) \right. \\ &\quad \left. - \gamma P_\pi \left(\pi \left\langle \phi, \frac{1}{T} \sum_{i=1}^T \omega_\diamond^i \right\rangle \right) \right] \\ &\quad \text{(By definition of } \bar{V}_\diamond^T \text{ in Eq. (26))} \\ &= (I - \gamma P_\pi)^{-1} \left[\left(\pi \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right) - \gamma P_\pi \frac{1}{T} (\pi \langle \phi, \omega_\diamond^T \rangle) \right]. \end{aligned}$$

Taking the infinity norm and using the triangle inequality,

$$\begin{aligned} \|\bar{V}_\diamond^T - V_\diamond^\pi\|_\infty &\leq \left\| (I - \gamma P_\pi)^{-1} \left(\pi \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right) \right\|_\infty \\ &\quad + \left\| (I - \gamma P_\pi)^{-1} \gamma P_\pi \frac{1}{T} (\pi \langle \phi, \omega_\diamond^T \rangle) \right\|_\infty \\ &\leq \|(I - \gamma P_\pi)^{-1}\|_{1, \infty} \left\| \pi \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right\|_\infty \\ &\quad + \|(I - \gamma P_\pi)^{-1} \gamma P_\pi\|_{1, \infty} \left\| \frac{1}{T} (\pi \langle \phi, \omega_\diamond^T \rangle) \right\|_\infty \quad \text{(By Holder's inequality)} \\ &\leq H \left[\left\| \pi \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right\|_\infty + \left\| \frac{1}{T} (\pi \langle \phi, \omega_\diamond^T \rangle) \right\|_\infty \right] \\ &\quad \text{(Since } \|(I - \gamma P_\pi)^{-1}\|_{1, \infty} \leq H, \|(I - \gamma P_\pi)^{-1} \gamma P_\pi\|_{1, \infty} \leq H) \end{aligned}$$

$$\begin{aligned}
 &\leq H \left[\left\| \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right\|_\infty + \left\| \frac{1}{T} \langle \phi, \omega_\diamond^T \rangle \right\|_\infty \right] \\
 &\hspace{15em} \text{(By definition of the } \pi \text{ operator)} \\
 &\leq \tilde{O} \left(H^2 \sqrt{\frac{d}{TM}} + \frac{H^2}{T} \right). \hspace{10em} \text{(By Lemma 24 and Lemma 26)}
 \end{aligned}$$

Using that for any $|\mathcal{S}|$ -dimensional vector V , $V(\rho) = \mathbb{E}_{s \sim \rho} |V(s)| \leq \|V\|_\infty$ completes the proof. \blacksquare

D.3.1. AUXILIARY LEMMAS

Since the updates in Algorithm 3 are a special case of those in Algorithm 2, the proofs of the auxiliary lemmas are analogous. We therefore present lemmas analogous to Lemma 15, Lemma 21 and Lemma 18, whose proofs follow by the same reasoning.

Lemma 24 *For any $t \in [T]$, with $\diamond = r$ or c and $M \geq \tilde{O}(dH^2)$, we have*

$$\|\langle \phi, \omega_\diamond^t \rangle\|_\infty \leq 2H \quad \text{and} \quad \|\hat{V}_\diamond^t\|_\infty \leq 2H$$

with probability at least $1 - \delta$.

Lemma 25 *For any $t \in [T]$, $\diamond = r$ or c , we have*

$$\left\langle \phi, \frac{1}{T} \sum_{i=1}^T \omega_\diamond^i \right\rangle = \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle + \diamond + \gamma P \left(\pi \left\langle \phi, \frac{1}{T} \sum_{i=1}^{T-1} \omega_\diamond^i \right\rangle \right).$$

Lemma 26 *With $\diamond = r$ or c and $M \geq \tilde{O}(dH^2)$, we have*

$$\left\| \left\langle \phi, W \left(\frac{1}{T} \sum_{i=1}^T (\hat{Q}_\diamond^i - \langle \phi, \omega_\diamond^i \rangle) \right) \right\rangle \right\|_\infty \leq \tilde{O} \left(H \sqrt{\frac{d}{TM}} \right)$$

with probability at least $1 - \delta$.

D.4. Proof of Corollary 9

Corollary 27 *Using LS-MDVI (Alg. 2) and LS-PE (Alg. 3) as instantiations of the MDP-Solver and PolicyEvaluation in Alg. 1 and using the DataCollection oracle described in Sec. 4.1 has the following guarantee: for a fixed $\varepsilon \in (0, 1]$, $\delta \in (0, 1)$, Alg. 1 with $\tilde{O}\left(\frac{d^2 H^4}{\varepsilon^2}\right)$ samples, $U = O\left(\frac{1}{\zeta(1-\gamma)}\right)$, $\eta = \frac{U(1-\gamma)}{\sqrt{K}}$, $K = O\left(\frac{1}{\varepsilon^2(1-\gamma)^2}\right)$, and $b' = b - O(\varepsilon)$, returns a mixture policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,*

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - O(\varepsilon), \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b - O(\varepsilon).$$

With the same algorithm parameters, but with $b' = b + O(\varepsilon)$ and $\tilde{O}\left(\frac{d^2 H^6}{\zeta^2 \varepsilon^2}\right)$ samples, Alg. 1 returns a mixture policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - O(\varepsilon), \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b.$$

Proof By Lemma 6 and Lemma 8, the sample complexity required to ensure $f(\mathcal{B}) \leq O(\varepsilon)$ is $TM|\mathcal{C}| = \tilde{O}\left(\frac{d^2 H^4}{\varepsilon^2}\right)$. Therefore, the guarantee for the relaxed feasibility setting follows directly from our meta-theorem (Theorem 4). For the strict feasibility setting, we rescale ε by a factor of $O(\zeta(1-\gamma))$. Since $\varepsilon \leq 1$ and $1-\gamma \leq 1$, the condition of $f(\mathcal{B}) \leq \zeta/6$ in Theorem 4 can be satisfied. The rescaling increases the sample complexity by a multiplicative factor of $\frac{1}{\zeta^2(1-\gamma)^2}$, thereby completing the proof. \blacksquare

D.5. Instantiating the MDP-Solver: G-Sampling-and-Stop

Instead of LS-MDVI, we can instantiate the linear MDP-Solver in Algorithm 1 with the GSS algorithm (Taupin et al., 2023). The GSS algorithm begins by computing a G-optimal sampling distribution over state-action pairs that minimizes the worst-case variance of value estimates. It then repeatedly samples transitions and rewards according to this distribution and uses regularized least-squares estimators to learn the reward and transition parameters of the MDP. For an arbitrary distribution $\tilde{\rho}$ over $\mathcal{S} \times \mathcal{A}$, let $G \in \mathbb{R}^{d \times d}$ and $g(\tilde{\rho}) \in \mathbb{R}$ be defined as:

$$G := \sum_{(x,b) \in \mathcal{C}} \tilde{\rho}(x,b) \phi(x,b) \phi(x,b)^\top \quad \text{and} \quad g(\tilde{\rho}) := \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \langle \phi(s,a), G^{-1} \phi(s,a) \rangle,$$

The GSS method samples one state-action pair $(s_t, a_t) \sim \rho^*$ in an iteration t where $\rho^* := \arg \min_{\rho \in \Delta_{\mathcal{S} \times \mathcal{A}}} g(\rho)$. We denote this data collection procedure as DataCollection-GSS. Note that this is different than the sampling scheme used in App. A.

For solving a linear MDP, the GSS algorithm uses a stopping rule based on confidence bounds derived from matrix concentration inequalities, and determines when the estimates are accurate enough to ensure that the returned policy is ε -optimal for the true MDP with high probability. The stopping time is denoted by

$$\tau = \inf\{t \geq 1 : Z(t) \geq \beta(t)\}$$

where $\beta(t)$ is a certain threshold and $Z(t)$ is the quantity we seek to control in order to achieve the desired sample complexity. Their main result in the setting of infinite-horizon γ -discounted linear unconstrained MDPs is stated below.

Theorem 28 (Theorem 2 and Theorem 3 in (Taupin et al., 2023)) *Let $\varepsilon, \delta \in (0, 1)$. The GSS algorithm returns an ε -optimal policy with probability at least $1 - \delta$, and the expected number of samples used is bounded by*

$$O\left(\frac{d}{(1-\gamma)^4 \varepsilon^2} \left(\log\left(\frac{1}{\delta}\right) + d \log\left(\frac{d}{(1-\gamma)^4 \varepsilon^2}\right)\right)\right).$$

Using the GSS algorithm as an alternative instantiation of MDP-Solver($r + \lambda_k c, \mathcal{B}, \phi$), we have that, with $N = \tilde{O}\left(\frac{d^2 H^4}{\varepsilon^2}\right)$, the GSS algorithm satisfies Assumption 2 with $f_{\text{mdp}}(\mathcal{B}) = O(\varepsilon)$. Hence, instantiating the three oracles by DataCollection-GSS, the GSS algorithm and using the same PolicyEvaluation oracle as in Alg. 3, we can use our meta-theorem (Theorem 4) to obtain the same sample complexity bounds as in Thm. 9.

Appendix E. Lower Bound

Theorem 10 *Let $\delta \in (0, 0.08]$, $\gamma \in [7/12, 1)$, $H = 1/(1 - \gamma)$, $\varepsilon \in (0, 0.002)$, $\zeta \in (0, 49/2280)$, $b \in [H/2, H]$, and $d \geq 6$. There exists a class of linear constrained MDPs such that any (ε, δ) -sound algorithm requires $\Omega(d^2 H^5 / \varepsilon^2 \zeta^2)$ samples from the generative model in the worst case.*

Proof We construct a lower bound by using the ideas from [Weisz et al. \(2022\)](#); [Vaswani et al. \(2022\)](#).

Hard Instance. We construct a class of hard linear CMDPs, where each individual CMDP, denoted as \mathcal{M}_β , is parameterized by a vector β . Each MDP in \mathcal{M}_β has four states: $\mathcal{S} = \{o, s_0, s_1, z\}$ with o being the initial state. The action space is $\mathcal{A} = \bar{\mathcal{A}} \times \{0, 1\} = \{\pm 1/\sqrt{d-5}\}^{d-5} \times \{0, 1\}$, where $\bar{\mathcal{A}}$ is a subset of a $(d-5)$ -dimensional hypercube. Thus, each action can be written as $a = (\bar{a}^\top, a')^\top \in \mathcal{A}$ where \bar{a} is a $d-5$ dimensional vector and $a' \in \{0, 1\}$ is a scalar. Recall from Assumption 1 that the reward $r(s, a) = \langle \phi(s, a), \psi_r \rangle$, cost $c(s, a) = \langle \phi(s, a), \psi_c \rangle$, and transition probability $\mathcal{P}(\cdot|s, a) = \langle \phi(s, a), \mu \rangle$ are all defined as linear functions. Let $\beta \in \bar{\mathcal{A}}$, $\Delta \in (0, 0.2(1 - \gamma)]$, $u = \frac{1-\gamma^2-\gamma\Delta}{\gamma}(b - \zeta)$. Note that the parameter ζ is the Slater constant for this CMDP instance. As can be verified from Eq. (30), there exists a policy that achieves a constraint value of $b + \zeta$, satisfying the definition. Now, we define the following parameters.

$$\begin{aligned} \phi(s_0, a) &= (1, 0, 0, 0, 0, \bar{a}^\top)^\top, & \phi(s_1, a) &= (0, 1, 0, 0, 0, \dots, 0)^\top, \\ \phi(o, a) &= (0, 0, a', 1 - a', 0, \dots, 0)^\top, & \phi(z, a) &= (0, 0, 0, 1, 1, 0, \dots, 0)^\top \\ \mu(s_0) &= (\gamma, 0, 1, 0, 0, \Delta\beta^\top)^\top, & \mu(s_1) &= (1 - \gamma, 1, 0, 0, 0, -\Delta\beta^\top)^\top, \\ \mu(o) &= (0, \dots, 0)^\top, & \mu(z) &= (0, 0, 0, 1, 0, \dots, 0)^\top, \\ \psi_r &= (1, 0, \dots, 0)^\top, & \psi_c &= (u, 0, 0, 0, (b + \zeta)(1 - \gamma)/\gamma, 0, \dots, 0)^\top. \end{aligned}$$

Under the above parameter settings, the transition dynamics is defined as follows.

$$\begin{aligned} \mathcal{P}_\beta(s_0|s_0, a) &= \gamma + \Delta\beta^\top \bar{a}, & \mathcal{P}_\beta(s_1|s_0, a) &= 1 - \gamma - \Delta\beta^\top \bar{a}, \\ \mathcal{P}_\beta(s_0|o, a) &= a', & \mathcal{P}_\beta(z|o, a) &= 1 - a', \end{aligned}$$

States s_1 and z are absorbing. The reward and constraint functions are then specified as follows. For any action a ,

$$\begin{aligned} r(s_0, a) &= 1, & r(z, a) &= 0, \\ c(s_0, a) &= u, & c(z, a) &= (b + \zeta)(1 - \gamma)/\gamma. \end{aligned}$$

Rewards and costs for other states are zero. By solving the Bellman equation $V_{r,\beta}^\pi(s_0) = 1 + \gamma(\gamma + \Delta\mathbb{E}_{\bar{a} \sim \pi}[\bar{a}^\top \beta])V_{r,\beta}^\pi(s_0)$, we can define the value functions for a policy π in the MDP \mathcal{M}_β . The notation $V_{r,\beta}^\pi$ represents the reward value function for policy π within the specific MDP instance defined by β . Similarly, $V_{c,\beta}^\pi$ is the cost value function. The value functions starting from the initial state o are

$$\begin{aligned} V_{r,\beta}^\pi(o) &= \gamma \sum_{a \in \mathcal{A}} a' \pi(a|o) V_{r,\beta}^\pi(s_0) = \frac{\gamma \sum_{a \in \mathcal{A}} a' \pi(a|o)}{1 - \gamma^2 - \gamma \Delta \mathbb{E}_{\bar{a} \sim \pi}[\bar{a}^\top \beta]}, & (28) \\ V_{c,\beta}^\pi(o) &= \gamma \sum_{a \in \mathcal{A}} a' \pi(a|o) V_{c,\beta}^\pi(s_0) + \gamma \sum_{a \in \mathcal{A}} (1 - a') \pi(a|o) V_{c,\beta}^\pi(z) \end{aligned}$$

$$= \frac{\gamma u \sum_{a \in \mathcal{A}} a' \pi(a|o)}{1 - \gamma^2 - \gamma \Delta \mathbb{E}_{\bar{a} \sim \pi}[\bar{a}^\top \beta]} + \sum_{a \in \mathcal{A}} \pi(a|o)(1 - a')(b + \zeta). \quad (29)$$

Note that $V_{r,\beta}^\pi(s_1) = 0$ and $V_{c,\beta}^\pi(z) = \frac{1}{1-\gamma} \cdot \frac{(b+\zeta)(1-\gamma)}{\gamma} = \frac{b+\zeta}{\gamma}$. We now define the probability of a policy π transitioning from o to s_0 as $p := \mathcal{P}_\beta^\pi[o \rightarrow s_0]$. According to the hard instance construction, $\mathcal{P}_\beta(s_0|o, a) = a'$. Thus, p is the marginal probability under π that $a' = 1$ when in state o . This is equivalent to stating that the policy π must satisfy:

$$\sum_{\bar{a} \in \bar{\mathcal{A}}} \pi((\bar{a}^\top, 1)^\top | o) = p \quad \text{and} \quad \sum_{\bar{a} \in \bar{\mathcal{A}}} \pi((\bar{a}^\top, 0)^\top | o) = 1 - p.$$

Since we are in the strict feasibility setting, the constraint has to be satisfied as $V_{c,\beta}^\pi(o) \geq b$ for a feasible policy π . We denote the optimal feasible policy for the CMDP \mathcal{M}_β as π_β^* . Next, we show that the policy π_β^* must satisfy:

$$\sum_{\bar{a} \in \bar{\mathcal{A}}} \pi_\beta^*((\bar{a}^\top, 1)^\top | o) = \frac{1}{2} \quad \text{and} \quad \sum_{a' \in \{0,1\}} \pi_\beta^*((\beta^\top, a')^\top | s_0) = 1.$$

Since states z and s_1 are absorbing, the choice of policies on them is irrelevant. It is clear from Eq. (28) that for any fixed p , $V_{r,\beta}^\pi(o)$ is maximized when $\mathbb{E}_{\bar{a} \sim \pi}[\bar{a}^\top \beta] = \beta^\top \beta = 1$. In this case, $V_{r,\beta}^\pi(o) = \frac{\gamma p}{1 - \gamma^2 - \gamma \Delta}$. Next, when $\mathbb{E}_{\bar{a} \sim \pi}[\bar{a}^\top \beta] = \beta^\top \beta$, we observe that

$$\begin{aligned} V_{c,\beta}^\pi(o) &= \frac{\gamma u p}{1 - \gamma^2 - \gamma \Delta} + (1 - p)(b + \zeta) \\ &= p(b - \zeta) + (1 - p)(b + \zeta) \quad (u = \frac{1 - \gamma^2 - \gamma \Delta}{\gamma} (b - \zeta)) \\ &= b + (1 - 2p)\zeta. \end{aligned} \quad (30)$$

Therefore, for a policy to be feasible, it needs to ensure $p \leq 1/2$. The value function is then maximized when $p = 1/2$. It means that, at the initial state o , the optimal policy should maintain an equal probability of moving into s_0 and z . At state s_0 , the optimal policy should pick an action from the set $\{(\beta^\top, 0)^\top, (\beta^\top, 1)^\top\}$ with probability 1. Note that for an action $a = (\bar{a}^\top, a')^\top$, the constraint satisfaction depends only on the action component a' , and the value function at state s_0 is determined by \bar{a} . Thus, \bar{a} and a' are optimized independently. For this optimal policy π_β^* , we have

$$V_{r,\beta}^{\pi_\beta^*}(o) = \frac{\gamma}{2(1 - \gamma^2 - \gamma \Delta)}, \quad \text{and} \quad V_{c,\beta}^{\pi_\beta^*}(o) = b.$$

Now, consider an algorithm which incorrectly identifies the true parameter β as a different parameter β' . The algorithm then outputs a policy $\hat{\pi}$. According to this misestimated parameter, the policy satisfies $\mathbb{E}_{\bar{a} \sim \hat{\pi}}[\bar{a}^\top \beta] = \beta'^\top \beta$. Hence, in the true CMDP \mathcal{M}_β , its corresponding value function involves $\beta'^\top \beta$. Suppose the probability of transition to s_0 from o with this policy is \hat{p} . We have

$$\begin{aligned} V_{c,\beta}^{\hat{\pi}}(o) &= \frac{\gamma u \hat{p}}{1 - \gamma^2 - \gamma \Delta \beta'^\top \beta} + (1 - \hat{p})(b + \zeta) \\ &= \frac{\gamma u \hat{p}}{1 - \gamma^2 - \gamma \Delta} + (1 - \hat{p})(b + \zeta) - \frac{\gamma u \hat{p}}{1 - \gamma^2 - \gamma \Delta} + \frac{\gamma u \hat{p}}{1 - \gamma^2 - \gamma \Delta \beta'^\top \beta} \end{aligned}$$

$$\begin{aligned}
 &= \hat{p}(b - \zeta) + (1 - \hat{p})(b + \zeta) - \varepsilon_c & (u = \frac{1 - \gamma^2 - \gamma\Delta}{\gamma}(b - \zeta)) \\
 &= b + (1 - 2\hat{p})\zeta - \varepsilon_c
 \end{aligned}$$

where

$$\varepsilon_c = \frac{\gamma u \hat{p}}{1 - \gamma^2 - \gamma\Delta} - \frac{\gamma u \hat{p}}{1 - \gamma^2 - \gamma\Delta\beta'^\top\beta} = \frac{\gamma^2 u \hat{p} \Delta (1 - \beta'^\top\beta)}{(1 - \gamma^2 - \gamma\Delta)(1 - \gamma^2 - \gamma\Delta\beta'^\top\beta)}.$$

Therefore, in order for $\hat{\pi}$ to be feasible in the original CMDP \mathcal{M}_β , we require that

$$(1 - 2\hat{p})\zeta \geq \varepsilon_c \Rightarrow \hat{p} \leq \frac{1}{2} - \frac{\varepsilon_c}{2\zeta}. \quad (31)$$

In the meanwhile, for a policy $\hat{\pi}$ to be (ε, δ) -sound, it must satisfied that, with probability $1 - \delta$,

$$\varepsilon \geq V_{r,\beta}^{\pi_{\hat{\pi}}^*}(o) - V_{r,\beta}^{\hat{\pi}}(o) = \frac{\gamma}{2(1 - \gamma^2 - \gamma\Delta)} - \frac{\hat{p}\gamma}{(1 - \gamma^2 - \gamma\Delta\beta'^\top\beta)}$$

which implies

$$\begin{aligned}
 \hat{p} &\geq \frac{(1 - \gamma^2 - \gamma\Delta\beta'^\top\beta)}{2(1 - \gamma^2 - \gamma\Delta)} - \varepsilon(1 - \gamma^2 - \gamma\Delta\beta'^\top\beta)/\gamma \\
 &\geq \frac{(1 - \gamma^2 - \gamma\Delta)}{2(1 - \gamma^2 - \gamma\Delta)} - \varepsilon(1 - \gamma^2 - \gamma\Delta\beta'^\top\beta)/\gamma & (\beta'^\top\beta \leq 1) \\
 &= \frac{1}{2} - \varepsilon(1 - \gamma^2 - \gamma\Delta\beta'^\top\beta)/\gamma.
 \end{aligned}$$

This lower bound implies that achieving higher rewards requires the policy to transition less frequently into the safe state. Since $\varepsilon \in (0, 0.002)$, $\gamma \in [7/12, 1)$, $\Delta \leq 0.2(1 - \gamma)$, and $\beta'^\top\beta \geq -1$, we have

$$\varepsilon(1 - \gamma^2 - \gamma\Delta\beta'^\top\beta)/\gamma \leq \varepsilon(1 - \gamma^2 + 0.2\gamma(1 - \gamma))/\gamma \leq \frac{3}{2}\varepsilon \leq 0.003$$

and thus

$$\hat{p} \geq 0.5 - 0.003 \geq 0.4. \quad (32)$$

Now we analyze parameter u . Recall that $\zeta \in (0, 49/2280)$, $\Delta \leq 1 - \gamma$, $b \in [H/2, H]$, and $H = \frac{1}{1-\gamma} \in [12/5, \infty)$. We obtain

$$u = \frac{1 - \gamma^2 - \gamma\Delta}{\gamma}(b - \zeta) \geq \frac{1 - \gamma}{\gamma} \left(\frac{1}{2(1 - \gamma)} - \frac{49}{2280} \right) \geq \frac{1}{2} - \frac{49}{2280} \geq \frac{1}{4}.$$

Next, we denote τ as the total number of queries a planner sends to the generative model before it stops its process and outputs a final policy. For any (ε, δ) -sound algorithm outputting a policy $\hat{\pi}$ with query complexity τ , we have (with probability $1 - \delta$)

$$\varepsilon \geq V_{r,\beta}^{\pi_{\hat{\pi}}^*}(o) - V_{r,\beta}^{\hat{\pi}}(o)$$

$$\begin{aligned}
 &= \frac{\gamma}{2(1-\gamma^2-\gamma\Delta)} - \frac{\hat{p}\gamma}{(1-\gamma^2-\gamma\Delta\beta'^\top\beta)} \\
 &\geq \frac{\gamma}{2(1-\gamma^2-\gamma\Delta)} - \frac{\hat{p}\gamma}{(1-\gamma^2-\gamma\Delta)} \quad (\beta'^\top\beta \leq 1) \\
 &= \frac{\gamma}{(1-\gamma^2-\gamma\Delta)} \left(\frac{1}{2} - \hat{p} \right) \\
 &\geq \frac{\gamma}{(1-\gamma^2)} \left(\frac{1}{2} - \hat{p} \right) \quad (\Delta \geq 0) \\
 &= \frac{\gamma}{(1-\gamma)(1+\gamma)} \left(\frac{1}{2} - \hat{p} \right) \\
 &\geq \frac{7}{19(1-\gamma)} \left(\frac{1}{2} - \hat{p} \right) \quad (\gamma \in [7/12, 1)) \\
 &\geq \frac{7\varepsilon_c}{19\zeta(1-\gamma)} \quad (\text{by Eq. (31)}) \\
 &= \frac{7}{19\zeta(1-\gamma)} \left(\frac{\gamma u \hat{p}}{1-\gamma^2-\gamma\Delta} - \frac{\gamma u \hat{p}}{1-\gamma^2-\gamma\Delta\beta'^\top\beta} \right) \\
 &= \frac{7\gamma u \hat{p}}{19\zeta(1-\gamma)} \left(\frac{1}{1-\gamma^2-\gamma\Delta} - \frac{1}{1-\gamma^2-\gamma\Delta\beta'^\top\beta} \right) \\
 &= \frac{7\gamma u \hat{p}}{19\zeta(1-\gamma)} \left(V_{r,\beta}^{\pi_\beta^*}(s_0) - V_{r,\beta}^{\hat{\pi}}(s_0) \right) \\
 &\geq \frac{49}{2280\zeta(1-\gamma)} \left(V_{r,\beta}^{\pi_\beta^*}(s_0) - V_{r,\beta}^{\hat{\pi}}(s_0) \right). \quad (\text{by the lower bound on } \gamma, u, \text{ and } \hat{p})
 \end{aligned}$$

Note that once the agent enters state s_0 , the CMDP framework no longer plays a role. At this point, the problem has been reduced to establishing a lower bound for unconstrained linear MDPs where the starting state distribution is s_0 , a result proven in Weisz et al. (2022). Note that the hard instance constructed in Weisz et al. (2022) is identical to ours after the initial transition from initial state o to s_0 .

Specifically, as shown in their Theorem H.3, under the parameter settings $\delta \in (0, 0.08]$, $\gamma \in [7/12, 1)$, $H = 1/(1-\gamma)$, $\alpha \in (0, 0.005\gamma H/(1+\gamma)^2)$, $\Delta \leq 0.2(1-\gamma)$, and $d \geq 3$, if τ is the number of samples, then,

$$\mathbb{P}_\beta \left(V_{r,\beta}^{\pi_\beta^*}(s_0) - V_{r,\beta}^{\hat{\pi}}(s_0) \geq \alpha \right) \geq \frac{3}{35} - \frac{8}{35} \sqrt{1 - \exp\left(-\frac{5\Delta^2 H \mathbb{E}_\beta[\tau]}{(d-2)^2}\right)}$$

which implies that the algorithm is not (α, δ) sound unless $\mathbb{E}_\beta[\tau] \geq \Omega(d^2 H^3 / \alpha^2)$. Substituting $\alpha = \frac{2280\varepsilon\zeta(1-\gamma)}{49}$ completes our proof. We note that the range of α in Weisz et al. (2022) is $(0, 0.005\gamma H/(1+\gamma)^2)$ where $0.005\gamma H/(1+\gamma)^2 \geq 0.00279$. Replacing $\alpha = \frac{2280\varepsilon\zeta(1-\gamma)}{49} \leq \varepsilon \in (0, 0.002)$ satisfies their choice of range. Except for the range of ε and the lower bound on d , other parameter choices exactly match those in Weisz et al. (2022). \blacksquare

Appendix F. Algorithms for Solving Tabular CMDPs

Algorithm 6 Tabular Mirror Descent Value Iteration (Tabular-MDVI)

Input: T (number of iterations), M (number of next-state samples obtained per state-action pair in each iteration), \square (rewards in MDP), $\mathcal{B} = \mathcal{B}_0 \cup \dots \cup \mathcal{B}_{T-1}$ (Buffer).

Output: π_T where $\forall s \in \mathcal{S} : \pi_T(\cdot|s) \in \arg \max_a \tilde{Q}_\square^T(s, a)$.

Initialize: $\hat{V}_\square^0 = \mathbf{0}$, $\hat{Q}_\square^{-1} = \mathbf{0}$.

procedure *Tabular-MDVI*($T, M, \square, \mathcal{B}$)

```

for  $t = 0, 1, 2, \dots, T - 1$  do
     $\forall (s, a) \in \mathcal{S} \times \mathcal{A} : \text{Access } (s, a, s'_m)_{m=1}^M$  from the buffer  $\mathcal{B}_t$ .
     $\forall (s, a) \in \mathcal{S} \times \mathcal{A} : \hat{Q}_\square^t(s, a) = \square(s, a) + \gamma \frac{1}{M} \sum_{m=1}^M \hat{V}_\square^t(s'_m)$ .
    Define  $\tilde{Q}_\square^t = \sum_{i=0}^t \hat{Q}_\square^i$ ;  $\forall s \in \mathcal{S} : \hat{V}_\square^{t+1}(s) = \max_a \{\tilde{Q}_\square^t(s, a)\} - \max_a \{\tilde{Q}_\square^{t-1}(s, a)\}$ .
end
end

```

Algorithm 7 Tabular Policy Evaluation (Tabular-PE)

Input: T (number of iterations), M (number of next-state samples obtained per state-action pair in each iteration), \diamond (either r or c), $\mathcal{B} = \mathcal{B}_0 \cup \dots \cup \mathcal{B}_{T-1}$ (Buffer), π (policy to be evaluated).

Output: $\bar{V}_\diamond^T(\rho) = \frac{1}{T} \sum_{i=1}^T \hat{V}_\diamond^i(\rho)$.

Initialize: $\hat{V}_\diamond^0 = \mathbf{0}$.

procedure *Tabular-PE*($T, M, \diamond, \mathcal{B}, \pi$)

```

for  $t = 0, 1, 2, \dots, T - 1$  do
     $\forall (s, a) \in \mathcal{S} \times \mathcal{A} : \text{Access } (s, a, s'_m)_{m=1}^M$  from the buffer  $\mathcal{B}_t$ .
     $\forall (s, a) \in \mathcal{S} \times \mathcal{A} : \hat{Q}_\diamond^t(s, a) = \diamond(s, a) + \gamma \frac{1}{M} \sum_{m=1}^M \hat{V}_\diamond^t(s'_m)$ .
     $\hat{V}_\diamond^{t+1} = \pi \hat{Q}_\diamond^t$ .
end
end

```

Appendix G. Instantiating the Framework for Tabular Constrained MDPs

We now instantiate the framework for tabular CMDPs, and prove that the resulting algorithm attains near-optimal sample complexity. In contrast to the linear setting, we set $\mathcal{C} = \mathcal{S} \times \mathcal{A}$ as the input to the `DataCollection` oracle. For the `MDP-Solver` and `PolicyEvaluation`, we adapt Algs. 2 and 3 to the tabular setting. In particular, for both these algorithms, we set the features to be $|\mathcal{S}||\mathcal{A}|$ dimensional one-hot encodings of the state-action space implying that the feature map ϕ is an $|\mathcal{S}||\mathcal{A}|$ -dimensional identity matrix. Consequently, the resulting algorithm does not require linear regression to estimate the Q -function. We provide the pseudo-code for these two instantiations is provided in App. F. Their corresponding optimality guarantees are proved in App. H and stated below.

Lemma 29 *For a fixed $\varepsilon \in (0, 1/H^2]$, $\delta \in (0, 1)$, any $k \in [K]$, and $T \geq 2 \log(T)/\gamma$, when using Alg. 6 at iteration k of Alg. 1 with $\square = r + \lambda_k c$, $M = \tilde{O}\left(\frac{H}{\varepsilon}\right)$ and $T = O\left(\frac{H^2}{\varepsilon}\right)$, the output policy π_T satisfies the following condition with probability $1 - \delta$,*

$$\max_{\pi} V_{r+\lambda_k c}^{\pi}(\rho) - V_{r+\lambda_k c}^{\pi_T}(\rho) \leq O((1 + \lambda_k)\varepsilon),$$

The resulting sample complexity is $N = T M |\mathcal{C}| = \tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^3}{\varepsilon^2}\right)$.

Lemma 30 For a fixed $\varepsilon \in (0, H]$, $\delta \in (0, 1)$, Alg. 7 with $M = \tilde{O}\left(\frac{H}{\varepsilon}\right)$ and $T = O\left(\frac{H^2}{\varepsilon}\right)$, the output \bar{V}_\diamond^T satisfies the following condition with probability $1 - \delta$,

$$|\bar{V}_\diamond^T(\rho) - V_\diamond^\pi(\rho)| \leq O(\varepsilon),$$

The resulting sample complexity is $N = T M |\mathcal{C}| = \tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^3}{\varepsilon^2}\right)$.

The proofs of Thms. 29 and 30 can use the total variance technique and a Bernstein-type concentration argument (Azar et al., 2013; Kozuno et al., 2022) and result in near-optimal bounds in the tabular setting. Moreover, the corresponding algorithms do not require constructing coresets or using (variance-weighted) linear regression. Consequently, unlike the linear setting in Sec. 4, the same buffer \mathcal{B} can be reused across all iterations of Alg. 1. This allows the near-optimal sample complexities of both Algs. 6 and 7 to be preserved for tabular CMDPs. In particular, we prove the following result in App. H.4.

Corollary 31 Let Alg. 6 and Alg. 7 be the instantiations of the `MDP-Solver` and `PolicyEvaluation` in Alg. 1. For a fixed $\varepsilon \in (0, 1/H^2]$, $\delta \in (0, 1)$, Alg. 1 with $\tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^3}{\varepsilon^2}\right)$ samples, $U = O\left(\frac{1}{\zeta(1-\gamma)}\right)$, $\eta = \frac{U(1-\gamma)}{\sqrt{K}}$, $K = O\left(\frac{1}{\varepsilon^2(1-\gamma)^2}\right)$, and $b' = b - O(\varepsilon)$, returns a policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - O(\varepsilon), \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b - O(\varepsilon).$$

Under the same conditions, but with $b' = b + O(\varepsilon)$ and $\tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^5}{\zeta^2\varepsilon^2}\right)$ samples, Alg. 1 returns a policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - O(\varepsilon), \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b.$$

The above result matches the near-optimal sample complexity bounds attained by the model-based algorithm in Vaswani et al. (2022). Furthermore, instantiating the `MDP-Solver` to be the model-based algorithm (Agarwal et al., 2020; Li et al., 2020) and using Alg. 1 will result in a near-optimal sample complexity for solving tabular CMDPs (see App. H.5 for details). Note that the `MDP-Solver` can also be instantiated by a range of model-free algorithms for solving unconstrained MDPs with access to a generative model (Azar et al., 2013; Sidford et al., 2023, 2018; Wang, 2017; Jin et al., 2024). Consequently, our framework can be interpreted as a generalization of the the primal-dual approach in (Vaswani et al., 2022) to handle model-free algorithms and linear function approximation.

Appendix H. Proofs for Section G

Throughout, we treat π as an operator that returns an $|\mathcal{S}|$ -dimensional vector s.t. for an arbitrary $|\mathcal{S}||\mathcal{A}|$ -dimensional vector u such that $(\pi u)(s) := \sum_{a \in \mathcal{A}} \pi(a|s) u(s, a)$. Furthermore, we define $P_\pi := \pi P$ where $P_\pi \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ and denotes the transition probability matrix induced by policy π .

We also recall that $\pi_k^* := \arg \max_{\pi} V_{r+\lambda_k c}^{\pi}$ and define $\bar{V}_{\square}^T := \frac{1}{T} \sum_{i=1}^T \hat{V}_{\square}^i$. We define $y_{t,m,s,a}$ to be the m -th next-state sample s'_m corresponding to the state-action pair (s, a) at iteration t . For a value function V , $\text{Var}(V)$ denote the function

$$\text{Var}(V) : (s, a) \mapsto (PV^2)(s, a) - (PV)^2(s, a)$$

and $\sigma(V) := \sqrt{\text{Var}(V)}$.

H.1. Proof of Lemma 29 (Optimality Guarantees for Algorithm 6 - Tabular CMDP)

Lemma 32 For a fixed $\varepsilon \in (0, 1/H^2]$, $\delta \in (0, 1)$, any $k \in [K]$, and $T \geq 2 \log(T)/\gamma$, when using Alg. 6 at iteration k of Alg. 1 with $\square = r + \lambda_k c$, $M = \tilde{O}\left(\frac{H}{\varepsilon}\right)$ and $T = O\left(\frac{H^2}{\varepsilon}\right)$, the output policy π_T satisfies the following condition with probability $1 - \delta$,

$$\max_{\pi} V_{r+\lambda_k c}^{\pi}(\rho) - V_{r+\lambda_k c}^{\pi_T}(\rho) \leq O((1 + \lambda_k)\varepsilon),$$

The resulting sample complexity is $N = T M |\mathcal{C}| = \tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^3}{\varepsilon^2}\right)$.

Proof By Lemma 35 and Lemma 36, we have

$$\begin{aligned} V_{\square}^{\pi_k^*}(\rho) - V_{\square}^{\pi_T}(\rho) &= V_{\square}^{\pi_k^*}(\rho) - \bar{V}_{\square}^T(\rho) + \bar{V}_{\square}^T(\rho) - V_{\square}^{\pi_T}(\rho) \\ &\leq \frac{7H^2(1 + \lambda_k)}{T} + \sqrt{\frac{6H^3}{TM}} + \sqrt{\frac{6H^6(1 + \lambda_k)^2}{TM} \left(\frac{50H^2}{T^2} + \frac{4\iota^2}{M} \right)} \end{aligned}$$

with probability at least $1 - \delta$. By letting $M = \frac{6H\iota^2}{\varepsilon}$, $T = \frac{82H^2}{\varepsilon}$, and $\varepsilon \in (0, 1/H^2]$, we have

$$\begin{aligned} V_{\square}^{\pi_k^*}(\rho) - V_{\square}^{\pi_T}(\rho) &\leq (1 + \lambda_k)\varepsilon/12 + \varepsilon/9\iota + (1 + \lambda_k)H^{3/2}\varepsilon \left(\frac{\varepsilon}{9H} + \sqrt{\varepsilon/81H} \right) \\ &= (1 + \lambda_k)\varepsilon/12 + \varepsilon/9\iota + (1 + \lambda_k)\varepsilon^2\sqrt{H}/9 + (1 + \lambda_k)H\varepsilon^{3/2}/9 \\ &\leq (1 + \lambda_k)\varepsilon/12 + \varepsilon/9 + (1 + \lambda_k)\varepsilon/9 + (1 + \lambda_k)\varepsilon/9 \\ &\quad (\varepsilon \in (0, 1/H^2] \text{ and } \iota = \log(2|\mathcal{S}||\mathcal{A}|/\delta) \geq 1) \\ &\leq (1 + \lambda_k)\varepsilon \end{aligned}$$

which completes the proof. ■

H.1.1. PROOF OF LEMMA 33 AND LEMMA 34 (PROOFS WITH Hoeffding's Inequality)

Lemma 33 Let π_k^* be defined as in Eq. (13), and let \bar{V}_{\square}^T denote the averaged empirical value function in Algorithm 6 when run with λ_k . For any $k \in [K]$, we have

$$V_r^{\pi_k^*}(\rho) + \lambda_k V_c^{\pi_k^*}(\rho) - \bar{V}_r^T(\rho) - \lambda_k \bar{V}_c^T(\rho) \leq \frac{3H^2(1 + \lambda_k)}{T} + 2H(1 + \lambda_k) \sqrt{\frac{\log(2|\mathcal{S}||\mathcal{A}|/\delta)}{TM}}$$

with probability at least $1 - \delta$.

Proof Since $(I - \gamma P_{\pi_k^*})V_{\square}^{\pi_k^*} = (\pi_k^* \square)$, we have

$$\begin{aligned} (I - \gamma P_{\pi_k^*})(V_{\square}^{\pi_k^*} - \bar{V}_{\square}^T) &= (\pi_k^* \square) - (\bar{V}_{\square}^T - \gamma P_{\pi_k^*} \bar{V}_{\square}^T) \\ &= (\pi_k^* \square) + \gamma P_{\pi_k^*} \bar{V}_{\square}^T - \bar{V}_{\square}^T \\ \implies V_{\square}^{\pi_k^*} - \bar{V}_{\square}^T &= (I - \gamma P_{\pi_k^*})^{-1}((\pi_k^* \square) + \gamma P_{\pi_k^*} \bar{V}_{\square}^T - \bar{V}_{\square}^T) \end{aligned} \quad (33)$$

By Lemma 38 and due to the greediness of π_t , for all $t \in [T]$, we have

$$\begin{aligned} \bar{V}_{\square}^t &= \frac{1}{t} \sum_{i=0}^{t-1} (\pi_t \hat{Q}_{\square}^i) \\ &\geq \frac{1}{t} \sum_{i=0}^{t-1} (\pi_k^* \hat{Q}_{\square}^i). \end{aligned} \quad (34)$$

Now, we have

$$\begin{aligned} V_{\square}^{\pi_k^*} - \bar{V}_{\square}^T &= (I - \gamma P_{\pi_k^*})^{-1}((\pi_k^* \square) + \gamma P_{\pi_k^*} \bar{V}_{\square}^T - \bar{V}_{\square}^T) && \text{(By Eq. (33))} \\ &\leq (I - \gamma P_{\pi_k^*})^{-1} \left((\pi_k^* \square) + \gamma P_{\pi_k^*} \bar{V}_{\square}^T - \frac{1}{T} \sum_{i=0}^{T-1} (\pi_k^* \hat{Q}_{\square}^i) \right) && \text{(By Eq. (34))} \\ &= (I - \gamma P_{\pi_k^*})^{-1} \left((\pi_k^* \square) + \gamma P_{\pi_k^*} \bar{V}_{\square}^T - (\pi_k^* \square) - \gamma P_{\pi_k^*} \frac{1}{T} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_{\square}^i) - \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_{\pi_k^*}^i \hat{V}_{\square}^i - \gamma P_{\pi_k^*} \hat{V}_{\square}^i] \right) && \text{(By Lemma 39)} \\ &= (I - \gamma P_{\pi_k^*})^{-1} \left(\gamma P_{\pi_k^*} \bar{V}_{\square}^T - \gamma P_{\pi_k^*} \frac{1}{T} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_{\square}^i) - \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_{\pi_k^*}^i \hat{V}_{\square}^i - \gamma P_{\pi_k^*} \hat{V}_{\square}^i] \right) \\ &= (I - \gamma P_{\pi_k^*})^{-1} \left(\gamma P_{\pi_k^*} \bar{V}_{\square}^T - \gamma P_{\pi_k^*} \frac{1}{T-1} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_{\square}^i) - \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_{\pi_k^*}^i \hat{V}_{\square}^i - \gamma P_{\pi_k^*} \hat{V}_{\square}^i] \right) \\ &\quad + (I - \gamma P_{\pi_k^*})^{-1} \left(\frac{1}{T(T-1)} \gamma P_{\pi_k^*} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_{\square}^i) \right). \end{aligned}$$

We note that $\frac{1}{T-1} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_{\square}^i) = \bar{V}_{\square}^{T-1}$ by Lemma 38. By defining $\mathcal{H}_{\pi_k^*} := \gamma(I - \gamma P_{\pi_k^*})^{-1} \pi_k^* \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$, we obtain

$$V_{\square}^{\pi_k^*} - \bar{V}_{\square}^T \leq \underbrace{\mathcal{H}_{\pi_k^*} P(\bar{V}_{\square}^T - \bar{V}_{\square}^{T-1})}_{\text{Term (i)}} + \underbrace{\mathcal{H}_{\pi_k^*} \frac{1}{T} \sum_{i=0}^{T-1} [P \hat{V}_{\square}^i - \hat{P}_i \hat{V}_{\square}^i]}_{\text{Term (ii)}} + \underbrace{\mathcal{H}_{\pi_k^*} \left(\frac{1}{T(T-1)} P \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_{\square}^i) \right)}_{\text{Term (iii)}}. \quad (35)$$

Note that for any vector $Q \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$,

$$\|\mathcal{H}_{\pi_k^*} Q\|_{\infty} = \|\gamma(I - \gamma P_{\pi_k^*})^{-1} \pi_k^* Q\|_{\infty}$$

$$\begin{aligned}
 &\leq \|\gamma(I - \gamma P_{\pi_k^*})^{-1}\|_1 \|\pi_k^* Q\|_\infty && \text{(By Holders inequality)} \\
 &\leq H \|\pi_k^* Q\|_\infty && \text{(Since } \|\gamma(I - \gamma P_{\pi_k^*})^{-1}\|_1 \leq H) \\
 &\leq H \|Q\|_\infty. && \text{(By definition of the } \pi \text{ operator)}
 \end{aligned}$$

In order to bound Term (i), using Thm. 40, we have

$$\left\| \mathcal{H}_{\pi_k^*} P(\bar{V}_\square^T - \bar{V}_\square^{T-1}) \right\|_\infty \leq \frac{2H^2(1 + \lambda_k)}{T}.$$

For bounding Term (ii), letting $t = T$ in Lemma 57 and invoking it twice for r and c , we have

$$\left\| \mathcal{H}_{\pi_k^*} \frac{1}{T} \sum_{i=0}^{T-1} \left[P\hat{V}_\square^i - \hat{P}_i \hat{V}_\square^i \right] \right\|_\infty \leq 2H^2(1 + \lambda_k) \sqrt{\frac{\iota}{TM}}$$

with probability at least $1 - \delta$.

Finally, we bound Term (iii) by noting that $\left\| \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_\square^i) \right\|_\infty \leq (T-1)H(1 + \lambda_k)$ due to Lemma 37. Hence,

$$\left\| \mathcal{H}_{\pi_k^*} \left(\frac{1}{T(T-1)} P \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_\square^i) \right) \right\|_\infty \leq \frac{H^2(1 + \lambda_k)}{T}.$$

Note that for any vector V , $V(\rho) \leq \|V\|_\infty$. Putting everything together, we have

$$V_r^{\pi_k^*}(\rho) + \lambda_k V_c^{\pi_k^*}(\rho) - \bar{V}_r^T(\rho) - \lambda_k \bar{V}_c^T(\rho) \leq \frac{3H^2(1 + \lambda_k)}{T} + 2H^2(1 + \lambda_k) \sqrt{\frac{\iota}{TM}}$$

with probability at least $1 - \delta$. ■

Lemma 34 *Let π_T be the output policy, and let \bar{V}_\diamond^T denote the averaged empirical value function in Algorithm 6 when run with λ_k . For any $k \in [K]$, we have*

$$\bar{V}_r^T(\rho) + \lambda_k \bar{V}_c^T(\rho) - V_r^{\pi_T}(\rho) - \lambda_k V_c^{\pi_T}(\rho) \leq \frac{2H^2(1 + \lambda_k)}{T} + 2H^2(1 + \lambda_k) \sqrt{\frac{\iota}{TM}}$$

with probability at least $1 - \delta$.

Proof The proof follows similar steps as before. Since $(I - \gamma P_{\pi_T})V_\square^{\pi_T} = \pi_T \square$, we have

$$\begin{aligned}
 (I - \gamma P_{\pi_T})(\bar{V}_\square^T - V_\square^{\pi_T}) &= (\bar{V}_\square^T - \gamma P_{\pi_T} \bar{V}_\square^T) - \pi_T \square \\
 &= \bar{V}_\square^T - (\pi_T \square) + \gamma P_{\pi_T} \bar{V}_\square^T \\
 \implies \bar{V}_\square^T - V_\square^{\pi_T} &= (I - \gamma P_{\pi_T})^{-1}(\bar{V}_\square^T - (\pi_T \square) + \gamma P_{\pi_T} \bar{V}_\square^T)
 \end{aligned} \tag{36}$$

Recall that for all $t \in [T]$, we have

$$\bar{V}_\diamond^t = \frac{1}{t} \sum_{i=1}^t \hat{V}_\diamond^i = \frac{1}{t} \sum_{i=0}^{t-1} (\pi_t \hat{Q}_\diamond^i). \tag{37}$$

Now, we have

$$\bar{V}_\square^T - V_\square^{\pi_T} = (I - \gamma P_{\pi_T})^{-1} (\bar{V}_\square^T - (\pi_T \square) + \gamma P_{\pi_T} \bar{V}_\square^T) \quad (\text{By Eq. (36)})$$

$$= (I - \gamma P_{\pi_T})^{-1} \left(\frac{1}{T} \sum_{i=0}^{T-1} (\pi_T \hat{Q}_\square^i) - (\pi_T \square) + \gamma P_{\pi_T} \bar{V}_\square^T \right) \quad (\text{By Eq. (37)})$$

$$= (I - \gamma P_{\pi_T})^{-1} \left(\frac{1}{T} \sum_{i=0}^{T-1} (\pi_T \hat{Q}_\square^i) - (\pi_T \square) - \gamma P_{\pi_T} \bar{V}_\square^T \right)$$

$$= (I - \gamma P_{\pi_T})^{-1} \left((\pi_T \square) + \gamma P_{\pi_T} \frac{1}{T} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_\square^i) + \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_{\pi_T}^i \hat{V}_\square^i - \gamma P_{\pi_T} \hat{V}_\square^i] \right. \\ \left. - (\pi_T \square) - \gamma P_{\pi_T} \bar{V}_\square^T \right) \quad (\text{By Lemma 39})$$

$$= (I - \gamma P_{\pi_T})^{-1} \left(\gamma P_{\pi_T} \frac{1}{T} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_\square^i) + \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_{\pi_T}^i \hat{V}_\square^i - \gamma P_{\pi_T} \hat{V}_\square^i] - \gamma P_{\pi_T} \bar{V}_\square^T \right)$$

$$\leq (I - \gamma P_{\pi_T})^{-1} \left(\gamma P_{\pi_T} \frac{1}{T-1} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_\square^i) + \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_{\pi_T}^i \hat{V}_\square^i - \gamma P_{\pi_T} \hat{V}_\square^i] - \gamma P_{\pi_T} \bar{V}_\square^T \right).$$

We note that $\frac{1}{T-1} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_\square^i) = \bar{V}_\square^{T-1}$. By letting $\mathcal{H}_{\pi_T} = \gamma(I - \gamma P_{\pi_T})^{-1} \pi_T$, we obtain

$$\bar{V}_\square^T - V_\square^{\pi_T} \leq \mathcal{H}_{\pi_T} P (\bar{V}_\square^{T-1} - \bar{V}_\square^T) + \mathcal{H}_{\pi_T} \frac{1}{T} \sum_{i=0}^{T-1} [\hat{P}_i \hat{V}_\square^i - P \hat{V}_\square^i]. \quad (38)$$

Note that for any vector $Q \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$,

$$\begin{aligned} \|\mathcal{H}_{\pi_T} Q\|_\infty &= \|\gamma(I - \gamma P_{\pi_T})^{-1} \pi_T Q\|_\infty \\ &\leq \|\gamma(I - \gamma P_{\pi_T})^{-1}\|_1 \|\pi_T Q\|_\infty && (\text{By Holders inequality}) \\ &\leq H \|\pi_T Q\|_\infty && (\text{Since } \|\gamma(I - \gamma P_{\pi_T})^{-1}\|_1 \leq H) \\ &\leq H \|Q\|_\infty. && (\text{By definition of the } \pi \text{ operator}) \end{aligned}$$

Thus, letting $t = T$ in Lemma 57, we have

$$\left\| \mathcal{H}_{\pi_T} \frac{1}{T} \sum_{i=0}^{T-1} [\hat{P}_i \hat{V}_\square^i - P \hat{V}_\square^i] \right\|_\infty \leq 2H^2(1 + \lambda_k) \sqrt{\frac{\iota}{TM}}$$

with probability at least $1 - \delta$. By Lemma 40,

$$\|\mathcal{H}_{\pi_T} P (\bar{V}_\square^{T-1} - \bar{V}_\square^T)\|_\infty \leq \frac{2H^2(1 + \lambda_k)}{T}.$$

Note that for any vector V , $V(\rho) \leq \|V\|_\infty$. Putting everything together, we have

$$\bar{V}_r^T(\rho) + \lambda_k \bar{V}_c^T(\rho) - V_r^{\pi_T}(\rho) - \lambda_k V_c^{\pi_T}(\rho) \leq \frac{2H^2(1 + \lambda_k)}{T} + 2H^2(1 + \lambda_k) \sqrt{\frac{\iota}{TM}}$$

with probability at least $1 - \delta$. ■

H.1.2. PROOF OF LEMMA 35 AND LEMMA 36 (PROOFS WITH BERNSTEIN'S INEQUALITY)

Lemma 35 *Let π_k^* be defined as in Eq. (13), and let \bar{V}_\diamond^T denote the averaged empirical value function in Algorithm 6 when run with λ_k . For any $k \in [K]$ and $T \geq 2 \log(T)/\gamma$, we have*

$$V_r^{\pi_k^*}(\rho) + \lambda_k V_c^{\pi_k^*}(\rho) - \bar{V}_r^T(\rho) - \lambda_k \bar{V}_c^T(\rho) \leq \sqrt{\frac{3H^4(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{16H^2\iota^2}{M} \right)} + \sqrt{\frac{3H^3}{TM}} + \frac{4H^2(1+\lambda_k)}{T}.$$

with probability at least $1 - \delta$.

Proof From Eq. (35), we have

$$V_\square^{\pi_k^*} - \bar{V}_\square^T \leq \underbrace{\mathcal{H}_{\pi_k^*} P(\bar{V}_\square^T - \bar{V}_\square^{T-1})}_{\text{Term (i)}} + \underbrace{\mathcal{H}_{\pi_k^*} \frac{1}{T} \sum_{i=0}^{T-1} [P\hat{V}_\square^i - \hat{P}_i\hat{V}_\square^i]}_{\text{Term (ii)}} + \underbrace{\mathcal{H}_{\pi_k^*} \left(\frac{1}{T(T-1)} P \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_\square^i) \right)}_{\text{Term (iii)}}$$

We bound Term (i) and Term (iii) the same way as before. Thus, we only have Term (ii) remains. By Lemma 46, we know

$$\frac{1}{t} \sum_{i=1}^t [\hat{P}_i \hat{V}_\square^i - P\hat{V}_\square^i](s, a) \leq \frac{H(1+\lambda_k)\iota}{tM} + \sqrt{Z}$$

where

$$Z := \frac{3H^2(1+\lambda_k)^2}{tM} \left(\frac{4}{T^2} + \frac{16H^2\iota^2}{M} \right) + \frac{3\text{Var}(V_\square^{\pi_k^*}(s, a))}{tM}.$$

Therefore,

$$\begin{aligned} \mathcal{H}_{\pi_k^*} \frac{1}{T} \sum_{i=0}^{T-1} [P\hat{V}_\square^i - \hat{P}_i\hat{V}_\square^i] &\leq \mathcal{H}_{\pi_k^*} \sqrt{\frac{3H^2(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{16H^2\iota^2}{M} \right)} \mathbf{1} + \mathcal{H}_{\pi_k^*} \frac{H(1+\lambda_k)\iota}{TM} \mathbf{1} + \mathcal{H}_{\pi_k^*} \sqrt{\frac{3}{TM}} \sigma(V_\square^{\pi_k^*}) \\ &\leq \mathcal{H}_{\pi_k^*} \sqrt{\frac{3H^2(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{16H^2\iota^2}{M} \right)} \mathbf{1} + \mathcal{H}_{\pi_k^*} \frac{H(1+\lambda_k)\iota}{TM} \mathbf{1} + \sqrt{\frac{3H^3}{TM}} \mathbf{1} \\ &\hspace{15em} \text{(By Lemma 63)} \\ &\leq \sqrt{\frac{3H^4(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{16H^2\iota^2}{M} \right)} \mathbf{1} + \frac{H^2(1+\lambda_k)\iota}{TM} \mathbf{1} + \sqrt{\frac{3H^3}{TM}} \mathbf{1}. \end{aligned}$$

Lastly, combining the upper bounds for Term (i) and Term (iii), we have

$$\begin{aligned} V_\square^{\pi_k^*} - \bar{V}_\square^T &\leq \sqrt{\frac{3H^4(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{16H^2\iota^2}{M} \right)} \mathbf{1} + \frac{H^2(1+\lambda_k)\iota}{TM} \mathbf{1} + \sqrt{\frac{3H^3}{TM}} \mathbf{1} + \frac{3H^2(1+\lambda_k)}{T} \mathbf{1} \\ &\leq \sqrt{\frac{3H^4(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{16H^2\iota^2}{M} \right)} \mathbf{1} + \sqrt{\frac{3H^3}{TM}} \mathbf{1} + \frac{4H^2(1+\lambda_k)}{T} \mathbf{1}. \end{aligned}$$

■

Lemma 36 *Let π_T be the output policy, and let \bar{V}_\diamond^T denote the averaged empirical value function in Algorithm 6 when run with λ_k . For any $k \in [K]$ and $T \geq 2 \log(T)/\gamma$, we have*

$$\bar{V}_r^T(\rho) + \lambda_k \bar{V}_c^T(\rho) - V_r^{\pi_T}(\rho) - \lambda_k V_c^{\pi_T}(\rho) \leq \frac{3H^2(1+\lambda_k)}{T} + \sqrt{\frac{3H^3}{TM}} + \sqrt{\frac{3H^6(1+\lambda_k)^2}{TM} \left(\frac{50}{T^2} + \frac{4\ell^2}{M} \right)}$$

with probability at least $1 - \delta$.

Proof Similarly as before, we have

$$\begin{aligned} \bar{V}_\square^T - V_\square^{\pi_T} &\leq \mathcal{H}_{\pi_T} P(\bar{V}_\square^{T-1} - \bar{V}_\square^T) + \mathcal{H}_{\pi_T} \frac{1}{T} \sum_{i=0}^{T-1} [\hat{P}_i \hat{V}_\square^i - P \hat{V}_\square^i] && \text{(By Eq. (38))} \\ &\leq \frac{2H^2(1+\lambda_k)}{T} \mathbf{1} + \mathcal{H}_{\pi_T} \frac{1}{T} \sum_{i=0}^{T-1} [\hat{P}_i \hat{V}_\square^i - P \hat{V}_\square^i] && \text{(By Lemma 40)} \\ &\leq \frac{2H^2(1+\lambda_k)}{T} \mathbf{1} + \mathcal{H}_{\pi_T} \sqrt{\frac{3H^2(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{4H^2\ell^2}{M} \right)} \mathbf{1} \\ &\quad + \mathcal{H}_{\pi_T} \frac{H(1+\lambda_k)\ell}{TM} \mathbf{1} + \mathcal{H}_{\pi_T} \sqrt{\frac{3}{TM}} \sigma(V_\square^{\pi_k^*}) && \text{(By Lemma 46)} \\ &\leq \frac{2H^2(1+\lambda_k)}{T} \mathbf{1} + \sqrt{\frac{3H^4(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{4H^2\ell^2}{M} \right)} \mathbf{1} \\ &\quad + \frac{H^2(1+\lambda_k)\ell}{TM} \mathbf{1} + \mathcal{H}_{\pi_T} \sqrt{\frac{3}{TM}} \sigma(V_\square^{\pi_k^*}) \\ &\leq \frac{3H^2(1+\lambda_k)}{T} \mathbf{1} + \sqrt{\frac{3H^4(1+\lambda_k)^2}{TM} \left(\frac{4}{T^2} + \frac{4H^2\ell^2}{M} \right)} \mathbf{1} + \mathcal{H}_{\pi_T} \sqrt{\frac{3}{TM}} \sigma(V_\square^{\pi_k^*}). \end{aligned}$$

Now, it remains to bound the last term. We first observe that

$$\begin{aligned} \sigma(V_\square^{\pi_k^*}) &\leq (V_\square^{\pi_k^*} - V_\square^{\pi_T}) + \sigma(V_\square^{\pi_T}) && \text{(By Lemma 62)} \\ &\leq |V_\square^{\pi_k^*} - V_\square^{\pi_T}| + \sigma(V_\square^{\pi_T}) && \text{(By Lemma 61)} \\ &\leq \frac{5H^2(1+\lambda_k)}{T} \mathbf{1} + 4H^2(1+\lambda_k) \sqrt{\frac{\ell}{TM}} \mathbf{1} + \sigma(V_\square^{\pi_T}) \\ &&& \text{(By combining Lemma 33 and Lemma 34)} \end{aligned}$$

Therefore,

$$\begin{aligned} \mathcal{H}_{\pi_T} \sqrt{\frac{3}{TM}} \sigma(V_\square^{\pi_k^*}) &\leq \mathcal{H}_{\pi_T} \sqrt{\frac{3}{TM}} \left(\frac{5H^2(1+\lambda_k)}{T} + 4H^2(1+\lambda_k) \sqrt{\frac{\ell}{TM}} \right) \mathbf{1} + \sqrt{\frac{3}{TM}} \mathcal{H}_{\pi_T} \sigma(V_\square^{\pi_T}) \\ &\leq \sqrt{\frac{3H^2}{TM}} \left(\frac{5H^2(1+\lambda_k)}{T} + 4H^2(1+\lambda_k) \sqrt{\frac{\ell}{TM}} \right) \mathbf{1} + \sqrt{\frac{3}{TM}} \mathcal{H}_{\pi_T} \sigma(V_\square^{\pi_T}) \\ &\leq \sqrt{\frac{3H^2}{TM}} \left(\frac{5H^2(1+\lambda_k)}{T} + 4H^2(1+\lambda_k) \sqrt{\frac{\ell}{TM}} \right) \mathbf{1} + \sqrt{\frac{3H^3}{TM}} \mathbf{1} \\ &&& \text{(By Lemma 63)} \end{aligned}$$

$$= \sqrt{\frac{3H^6(1+\lambda_k)^2}{TM}} \left(\frac{5}{T} + \sqrt{\frac{16\ell}{TM}} \right) \mathbf{1} + \sqrt{\frac{3H^3}{TM}} \mathbf{1}.$$

By combining the above results and consolidating like terms, we conclude the proof. \blacksquare

H.1.3. AUXILIARY LEMMAS

Lemma 37 Denote $\square = r + \lambda_k c$. For any $k \in [K]$ and any $t \in [T]$, $\hat{Q}_\square^t(s, a)$ and $\hat{V}_\square^t(s)$ are bounded by $(1 + \lambda_k)H$.

Proof We prove it by induction. By initialization, $\hat{Q}_\square^1(s, a) = r(s, a) + \lambda_k c(s, a) \leq 1 + \lambda_k$ and $\hat{V}_\square^1(s) = \hat{Q}_\square^1(s, \pi_1(\cdot|s)) \leq 1 + \lambda_k \leq (1 + \lambda_k)H$. Now, suppose $\hat{V}_\square^{t-1}(s)$ is bounded by $(1 + \lambda_k)H$ for some $t \geq 1$. We have

$$\begin{aligned} \hat{V}_\square^t &= \sum_{i=0}^t (\pi_t \hat{Q}_\square^i) - \sum_{i=0}^{t-1} (\pi_{t-1} \hat{Q}_\square^i) && \text{(From the last line in Algorithm 6)} \\ &\leq \sum_{i=0}^t (\pi_t \hat{Q}_\square^i) - \sum_{i=0}^{t-1} (\pi_t \hat{Q}_\square^i) && \text{(By the greediness of } \pi_{t-1}) \\ &= (\pi_t \hat{Q}_\square^{t-1}) \\ &\leq (\pi_t \square) + \gamma (\hat{P}_{\pi_t}^{t-1} \hat{V}_\square^{t-1}) && \text{(From the second last line in Algorithm 6)} \\ &\leq (1 + \lambda_k + \gamma(1 + \lambda_k)H) \mathbf{1} && \text{(Induction hypothesis)} \\ &= (1 + \lambda_k)H \mathbf{1}. && (1 + \frac{\gamma}{1-\gamma} = \frac{1}{1-\gamma}) \end{aligned}$$

Therefore, $\hat{V}_\square^t(s)$ is bounded by $(1 + \lambda_k)H$. As a consequence, $\hat{Q}_\square^t(s, a)$ is also bounded by $(1 + \lambda_k)H$. \blacksquare

Lemma 38 For any $k \in [K]$ and $t \in [T]$, we have

$$\bar{V}_\diamond^t := \frac{1}{t} \sum_{i=1}^t \hat{V}_\diamond^i = \frac{1}{t} \sum_{i=0}^{t-1} (\pi_t \hat{Q}_\diamond^i).$$

Proof

$$\begin{aligned} \bar{V}_\diamond^t &:= \frac{1}{t} \sum_{i=1}^t \hat{V}_\diamond^i \\ &= \frac{1}{t} \sum_{i=0}^{t-1} \left(\sum_{j=0}^i (\pi_{i+1} \hat{Q}_\diamond^j) - \sum_{j=0}^{i-1} (\pi_i \hat{Q}_\diamond^j) \right) && \text{(From the last line in Algorithm 6)} \\ &= \frac{1}{t} \sum_{i=0}^{t-1} (\pi_t \hat{Q}_\diamond^i). && \text{(Due to telescoping sum)} \end{aligned}$$

\blacksquare

Lemma 39 For any $k \in [K]$ and $t \in [T]$, we have

$$\sum_{i=0}^{t-1} \hat{Q}_\diamond^i = t \diamond + \gamma P \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\diamond^i) + \sum_{i=0}^{t-1} [\gamma \hat{P}_i \hat{V}_\diamond^i - \gamma P \hat{V}_\diamond^i],$$

and

$$\sum_{i=0}^{t-1} \hat{Q}_\square^i = t(r + \lambda_k c) + \gamma P \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) + \sum_{i=0}^{t-1} [\gamma \hat{P}_{\pi_t}^i \hat{V}_\square^i - \gamma P_{\pi_t} \hat{V}_\square^i].$$

Proof We prove the first equality. The second equality follows by linearity.

$$\begin{aligned} \sum_{i=0}^{t-1} \hat{Q}_\diamond^i &= \sum_{i=0}^{t-1} [\diamond + \gamma \hat{P}_i \hat{V}_\diamond^i] && \text{(From the second last line in Algorithm 6)} \\ &= \sum_{i=0}^{t-1} [\diamond + \gamma P \hat{V}_\diamond^i - \gamma P \hat{V}_\diamond^i + \gamma \hat{P}_i \hat{V}_\diamond^i] \\ &= t \diamond + \gamma P \sum_{i=0}^{t-1} \hat{V}_\diamond^i + \sum_{i=0}^{t-1} [\gamma \hat{P}_i \hat{V}_\diamond^i - \gamma P \hat{V}_\diamond^i] \\ &= t \diamond + \gamma P \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\diamond^i) + \sum_{i=0}^{t-1} [\gamma \hat{P}_i \hat{V}_\diamond^i - \gamma P \hat{V}_\diamond^i]. && \text{(From Lemma 38)} \end{aligned}$$

■

Lemma 40 For any $k \in [K]$,

$$\|\bar{V}_\square^T - \bar{V}_\square^{T-1}\|_\infty \leq \frac{2H(1 + \lambda_k)}{T}.$$

Proof We present the proof for the case of $\bar{V}_\square^T - \bar{V}_\square^{T-1}$. The proof for the another case is similar. By the definition of \bar{V}_\square^t and due to the greediness of π_{T-1} , we have

$$\begin{aligned} \bar{V}_\square^T - \bar{V}_\square^{T-1} &= \frac{1}{T} \sum_{i=0}^{T-1} (\pi_T \hat{Q}_\square^i) - \frac{1}{T-1} \sum_{i=0}^{T-2} (\pi_{T-1} \hat{Q}_\square^i) \\ &\leq \frac{1}{T} \sum_{i=0}^{T-1} (\pi_T \hat{Q}_\square^i) - \frac{1}{T-1} \sum_{i=0}^{T-2} (\pi_T \hat{Q}_\square^i) \\ &\leq \frac{1}{T} \sum_{i=0}^{T-1} (\pi_T \hat{Q}_\square^i) - \frac{1}{T} \sum_{i=0}^{T-2} (\pi_T \hat{Q}_\square^i) \\ &\leq \frac{1}{T} (\pi_T \hat{Q}_\square^{T-1}) \\ &\leq \frac{2H(1 + \lambda_k)}{T} \mathbf{1}. && \text{(By Lemma 37)} \end{aligned}$$

■

H.2. Proof of Lemma 30 (Optimality Guarantees for Algorithm 7 - Tabular CMDP)

Lemma 41 For a fixed $\varepsilon \in (0, H]$, $\delta \in (0, 1)$, Alg. 7 with $M = \tilde{O}\left(\frac{H}{\varepsilon}\right)$ and $T = O\left(\frac{H^2}{\varepsilon}\right)$, the output $\bar{\mathcal{V}}_\diamond^T$ satisfies the following condition with probability $1 - \delta$,

$$|\bar{\mathcal{V}}_\diamond^T(\rho) - V_\diamond^\pi(\rho)| \leq O(\varepsilon),$$

The resulting sample complexity is $N = TM|\mathcal{C}| = \tilde{O}\left(\frac{|S||A|H^3}{\varepsilon^2}\right)$.

Proof By Lemma 44, we have

$$\|\bar{\mathcal{V}}_\diamond^T - V_\diamond^\pi\|_\infty \leq \tilde{O}\left(\frac{H^2}{T} + \frac{H}{tM} + \sqrt{\frac{H^4}{tM^2}} + \sqrt{\frac{H^3}{tM}}\right)$$

with probability at least $1 - \delta$. By letting $M = \tilde{O}\left(\frac{H}{\varepsilon}\right)$, $T = \tilde{O}\left(\frac{H^2}{\varepsilon}\right)$, and $\varepsilon \in (0, H]$, we have

$$\|\bar{\mathcal{V}}_\diamond^T - V_\diamond^\pi\|_\infty \leq O\left(\varepsilon + \frac{\varepsilon^2}{H^2} + \varepsilon + \varepsilon\right) \leq O(\varepsilon)$$

with total sample complexity $N = TM|\mathcal{C}| = \tilde{O}\left(\frac{|S||A|H^3}{\varepsilon^2}\right)$ and probability at least $1 - \delta$. \blacksquare

H.2.1. AUXILIARY LEMMAS

Since Algorithm 7 is equivalent to running Algorithm 6 with a fixed policy, the following lemma follows directly from Lemma 37.

Lemma 42 For any $t \in [T]$, $\hat{\mathcal{Q}}_\diamond^t(s, a)$ and $\hat{\mathcal{V}}_\diamond^t(s)$ are bounded by H .

Lemma 43 $\bar{\mathcal{V}}_\diamond^T - V_\diamond^\pi \leq \mathcal{H}_\pi P(\bar{\mathcal{V}}_\diamond^{T-1} - \bar{\mathcal{V}}_\diamond^T) + \mathcal{H}_\pi \frac{1}{T} \sum_{i=0}^{T-1} [\hat{P}_i \hat{\mathcal{V}}_\diamond^i - P \hat{\mathcal{V}}_\diamond^i]$.

Proof First, we notice that

$$\begin{aligned} \sum_{i=0}^{t-1} \hat{\mathcal{Q}}_\diamond^i &= \sum_{i=0}^{t-1} [\diamond + \gamma \hat{P}_i \hat{\mathcal{V}}_\diamond^i] \\ &= \sum_{i=0}^{t-1} [\diamond + \gamma P \hat{\mathcal{V}}_\diamond^i - \gamma P \hat{\mathcal{V}}_\diamond^i + \gamma \hat{P}_i \hat{\mathcal{V}}_\diamond^i] \\ &= t \diamond + \gamma P \sum_{i=0}^{t-1} \hat{\mathcal{V}}_\diamond^i + \sum_{i=0}^{t-1} [\gamma \hat{P}_i \hat{\mathcal{V}}_\diamond^i - \gamma P \hat{\mathcal{V}}_\diamond^i] \\ &= t \diamond + \gamma P \sum_{i=0}^{t-2} (\pi \hat{\mathcal{Q}}_\diamond^i) + \sum_{i=0}^{t-1} [\gamma \hat{P}_i \hat{\mathcal{V}}_\diamond^i - \gamma P \hat{\mathcal{V}}_\diamond^i]. \end{aligned} \tag{39}$$

It is different from Lemma 39 because the policy is now fixed at each iteration. The rest of the proof follows the same set of steps as in the proof of Lemma 33. Denote $\diamond = r$ or c . Since $(I - \gamma P_\pi)V_\diamond^\pi = \pi_\diamond$, we have

$$\begin{aligned} (I - \gamma P_\pi)(\bar{V}_\diamond^T - V_\diamond^\pi) &= (\bar{V}_\diamond^T - \gamma P_\pi \bar{V}_\diamond^T) - \pi_\diamond \\ &= \bar{V}_\diamond^T - (\pi_\diamond) + \gamma P_\pi \bar{V}_\diamond^T \\ \implies \bar{V}_\diamond^T - V_\diamond^\pi &= (I - \gamma P_\pi)^{-1}(\bar{V}_\diamond^T - (\pi_\diamond) + \gamma P_\pi \bar{V}_\diamond^T) \end{aligned} \quad (40)$$

Recall that $\bar{V}_\diamond^t = \frac{1}{t} \sum_{i=1}^t \hat{V}_\diamond^i = \pi \frac{1}{t} \sum_{i=0}^{t-1} \hat{Q}_\diamond^i$ for all $t \in [T]$. Now, we have

$$\begin{aligned} \bar{V}_\diamond^T - V_\diamond^\pi &= (I - \gamma P_\pi)^{-1}(\bar{V}_\diamond^T - (\pi_\diamond) + \gamma P_\pi \bar{V}_\diamond^T) && \text{(By Eq. (40))} \\ &= (I - \gamma P_\pi)^{-1} \left(\frac{1}{T} \sum_{i=0}^{T-1} (\pi \hat{Q}_\diamond^i) - (\pi_\diamond) + \gamma P_\pi \bar{V}_\diamond^T \right) \\ &= (I - \gamma P_\pi)^{-1} \left((\pi_\diamond) + \gamma P_\pi \frac{1}{T} \sum_{i=0}^{T-2} (\pi \hat{Q}_\diamond^i) + \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_\pi^i \hat{V}_\diamond^i - \gamma P_\pi \hat{V}_\diamond^i] - (\pi_\diamond) - \gamma P_\pi \bar{V}_\diamond^T \right) \\ &&& \text{(By Eq. (39))} \\ &= (I - \gamma P_\pi)^{-1} \left(\gamma P_\pi \frac{1}{T} \sum_{i=0}^{T-2} (\pi \hat{Q}_\diamond^i) + \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_\pi^i \hat{V}_\diamond^i - \gamma P_\pi \hat{V}_\diamond^i] - \gamma P_\pi \bar{V}_\diamond^T \right) \\ &\leq (I - \gamma P_\pi)^{-1} \left(\gamma P_\pi \frac{1}{T-1} \sum_{i=0}^{T-2} (\pi \hat{Q}_\diamond^i) + \frac{1}{T} \sum_{i=0}^{T-1} [\gamma \hat{P}_\pi^i \hat{V}_\diamond^i - \gamma P_\pi \hat{V}_\diamond^i] - \gamma P_\pi \bar{V}_\diamond^T \right). \end{aligned}$$

We note that $\frac{1}{T-1} \sum_{i=0}^{T-2} (\pi \hat{Q}_\diamond^i) = \bar{V}_\diamond^{T-1}$, as it is an equivalent result of Lemma 38 with a fixed policy. By letting $\mathcal{H}_\pi = \gamma(I - \gamma P_\pi)^{-1}\pi$, we obtain

$$\bar{V}_\diamond^T - V_\diamond^\pi \leq \mathcal{H}_\pi P(\bar{V}_\diamond^{T-1} - \bar{V}_\diamond^T) + \mathcal{H}_\pi \frac{1}{T} \sum_{i=0}^{T-1} [\hat{P}_i \hat{V}_\diamond^i - P \hat{V}_\diamond^i]. \quad (41)$$

■

Lemma 44 *We have*

$$\|\bar{V}_\diamond^T - V_\diamond^\pi\|_\infty \leq \tilde{O} \left(\frac{H^2}{T} + \frac{H}{tM} + \sqrt{\frac{H^4}{tM^2}} + \sqrt{\frac{H^3}{tM}} \right)$$

with probability at least $1 - \delta$.

Proof By Lemma 43, we have

$$\bar{V}_\diamond^T - V_\diamond^\pi \leq \mathcal{H}_\pi P(\bar{V}_\diamond^{T-1} - \bar{V}_\diamond^T) + \mathcal{H}_\pi \frac{1}{T} \sum_{i=0}^{T-1} [\hat{P}_i \hat{V}_\diamond^i - P \hat{V}_\diamond^i]$$

$$\leq \tilde{O}\left(\frac{H^2}{T}\right) \mathbf{1} + \mathcal{H}_\pi \frac{1}{T} \sum_{i=0}^{T-1} \left[\hat{P}_i \hat{V}_\diamond^i - P \hat{V}_\diamond^i \right]. \quad (\text{By Lemma 45})$$

Thus, it remains to bound the second term. By Lemma 47 we have

$$\frac{1}{t} \sum_{i=0}^{t-1} \left[\hat{P}_i \hat{V}_\diamond^i - P \hat{V}_\diamond^i \right] (s, a) \leq \frac{H\iota}{tM} + \sqrt{Z}$$

where

$$Z := \frac{1}{tM} \left(\frac{H^4}{M} + \text{Var}(V_\diamond^\pi(s, a)) \right)$$

with probability at least $1 - \delta$. Therefore,

$$\begin{aligned} \bar{V}_\diamond^T - V_\diamond^\pi &\leq \tilde{O}\left(\frac{H^2}{T}\right) \mathbf{1} + \frac{H\iota}{tM} \mathbf{1} + \sqrt{\frac{H^4}{tM^2}} \mathbf{1} + \sqrt{\frac{1}{tM}} \mathcal{H}_\pi \sigma(V_\diamond^\pi) \\ &\leq \tilde{O}\left(\frac{H^2}{T}\right) \mathbf{1} + \frac{H\iota}{tM} \mathbf{1} + \sqrt{\frac{H^4}{tM^2}} \mathbf{1} + \sqrt{\frac{H^3}{tM}} \mathbf{1} \\ &\leq \tilde{O}\left(\frac{H^2}{T} + \frac{H\iota}{tM} + \sqrt{\frac{H^4}{tM^2}} + \sqrt{\frac{H^3}{tM}}\right) \mathbf{1} \end{aligned} \quad (\text{By Lemma 63})$$

which completes the proof. \blacksquare

Lemma 45 $\|\bar{V}_\diamond^T - \bar{V}_\diamond^{T-1}\|_\infty \leq \frac{H}{T}$.

Proof Similar to the proof of Lemma 40, we have

$$\begin{aligned} \bar{V}_\diamond^T - \bar{V}_\diamond^{T-1} &= \frac{1}{T} \sum_{i=0}^{T-1} (\pi \hat{Q}_\diamond^i) - \frac{1}{T-1} \sum_{i=0}^{T-2} (\pi \hat{Q}_\diamond^i) \\ &= \frac{1}{T} \sum_{i=0}^{T-1} (\pi \hat{Q}_\diamond^i) - \frac{1}{T-1} \sum_{i=0}^{T-2} (\pi \hat{Q}_\diamond^i) \\ &\leq \frac{1}{T} \sum_{i=0}^{T-1} (\pi \hat{Q}_\diamond^i) - \frac{1}{T} \sum_{i=0}^{T-2} (\pi \hat{Q}_\diamond^i) \\ &\leq \frac{1}{T} (\pi \hat{Q}_\diamond^{T-1}) \\ &\leq \frac{H}{T} \mathbf{1}. \end{aligned} \quad (\text{By Lemma 42})$$

H.3. Proof of Lemma 46 and Lemma 47 (Concentration Error Bounds with Bernstein's Inequality - Tabular CMDP)

All the proofs presented in this section are adapted from the proofs for Lemmas 5 to 8 in [Kozuno et al. \(2022\)](#), with substantial modifications to suit our setting.

Lemma 46 For any $t \geq 2 \log(t)/\gamma$ and $k \in [K]$, we have

$$\frac{1}{t} \sum_{i=1}^t \left[\hat{P}_i \hat{V}_\square^i - P \hat{V}_\square^i \right] (s, a) \leq \frac{H(1 + \lambda_k)t}{tM} + \sqrt{Z}$$

where

$$Z := \frac{3H^2(1 + \lambda_k)^2}{tM} \left(\frac{4}{t^2} + \frac{16H^2\iota^2}{M} \right) + \frac{3\text{Var}(V_\square^{\pi_k^*}(s, a))}{tM}$$

with probability at least $1 - \delta$.

Proof We have

$$\begin{aligned} \frac{1}{t} \sum_{i=1}^t \left[\hat{P}_i \hat{V}_\square^i - P \hat{V}_\square^i \right] (s, a) &= \frac{1}{t} \sum_{i=1}^t \frac{1}{M} \sum_{m=1}^M \left[\hat{V}_\square^i(y_{i,m,s,a}) - (P \hat{V}_\square^i)(s, a) \right] \\ &= \frac{1}{tM} \sum_{i=1}^t \sum_{m=1}^M \left[\hat{V}_\square^i(y_{i,m,s,a}) - (P \hat{V}_\square^i)(s, a) \right]. \end{aligned}$$

The above is a sum of bounded martingale differences with respect to the filtration $(\mathcal{F})_{i=1, m=1}^{t, M}$. Let $X_{i,m} = \frac{1}{tM} \left(\hat{V}_\square^i(y_{i,m,s,a}) - (P \hat{V}_\square^i)(s, a) \right)$. It can be noted that $X_{i,m} \leq \frac{H(1+\lambda_k)}{tM}$ (by Lemma 37) and $\mathbb{E}[X_{i,m}] = 0$. Next, we bound Z' as defined in Lemma 60

$$\begin{aligned} Z' &= \sum_{i=1}^t \sum_{m=1}^M \mathbb{E} [X_{i,m}^2] \\ &= \sum_{i=1}^t \sum_{m=1}^M \mathbb{E} \left[\frac{1}{t^2 M^2} \left(\hat{V}_\square^i(y_{i,m,s,a}) - (P \hat{V}_\square^i)(s, a) \right)^2 \right] \\ &= \frac{1}{t^2 M^2} \sum_{i=1}^t \sum_{m=1}^M \text{Var}(\hat{V}_\square^i(y_{i,m,s,a})) \\ &\leq \frac{3}{t^2 M^2} \sum_{i=1}^t \sum_{m=1}^M \left(\frac{4H^2(1 + \lambda_k)^2}{t^2} + \frac{16H^4(1 + \lambda_k)^2 \iota^2}{M} + \text{Var}(V_\square^{\pi_k^*}(s, a)) \right) \quad (\text{By Lemma 51}) \\ &= \frac{3}{tM} \left(\frac{4H^2(1 + \lambda_k)^2}{t^2} + \frac{16H^4(1 + \lambda_k)^2 \iota^2}{M} + \text{Var}(V_\square^{\pi_k^*}(s, a)) \right) \\ &:= Z. \end{aligned}$$

By letting $U = H(1 + \lambda_k)$ in Lemma 60, we have

$$\frac{1}{t} \sum_{i=1}^t \left[\hat{P}_i \hat{V}_\square^i - P \hat{V}_\square^i \right] (s, a) \leq \frac{H(1 + \lambda_k)t}{tM} + \sqrt{Z}$$

with probability at least $1 - \delta$. ■

Lemma 47 For any $t \geq 2 \log(t)/\gamma$, we have

$$\frac{1}{t} \sum_{i=0}^{t-1} \left[\hat{P}_i \hat{V}_\diamond^i - P \hat{V}_\diamond^i \right] (s, a) \leq \frac{Ht}{tM} + \sqrt{Z}$$

where

$$Z := \frac{1}{tM} \left(\frac{H^4}{M} + \text{Var}(V_\diamond^\pi(s, a)) \right)$$

with probability at least $1 - \delta$.

Proof We have

$$\begin{aligned} \frac{1}{t} \sum_{i=0}^{t-1} \left[\hat{P}_i \hat{V}_\diamond^i - P \hat{V}_\diamond^i \right] (s, a) &= \frac{1}{t} \sum_{i=1}^t \frac{1}{M} \sum_{m=1}^M \left[\hat{V}_\diamond^i(y_{i,m,s,a}) - (P \hat{V}_\diamond^i)(s, a) \right] \\ &= \frac{1}{tM} \sum_{i=1}^t \sum_{m=1}^M \left[\hat{V}_\diamond^i(y_{i,m,s,a}) - (P \hat{V}_\diamond^i)(s, a) \right]. \end{aligned}$$

The above is a sum of bounded martingale differences with respect to the filtration $(\mathcal{F})_{i=1, m=1}^{t, M}$. Let $X_{i,m} = \frac{1}{tM} \left(\hat{V}_\diamond^i(y_{i,m,s,a}) - (P \hat{V}_\diamond^i)(s, a) \right)$. It can be noted that $X_{i,m} \leq \frac{H}{tM}$ and $\mathbb{E}[X_{i,m}] = 0$. Next, we bound Z' as defined in Lemma 60

$$\begin{aligned} Z' &= \sum_{i=1}^t \sum_{m=1}^M \mathbb{E} [X_{i,m}^2] \\ &= \sum_{i=1}^t \sum_{m=1}^M \mathbb{E} \left[\frac{1}{t^2 M^2} \left(\hat{V}_\diamond^i(y_{i,m,s,a}) - (P \hat{V}_\diamond^i)(s, a) \right)^2 \right] \\ &= \frac{1}{t^2 M^2} \sum_{i=1}^t \sum_{m=1}^M \text{Var}(\hat{V}_\diamond^i(y_{i,m,s,a})) \\ &\leq \frac{1}{t^2 M^2} \sum_{i=1}^t \sum_{m=1}^M \left(\frac{H^4}{M} + \text{Var}(V_\diamond^\pi(s, a)) \right) \quad (\text{By Lemma 54}) \\ &= \frac{1}{tM} \left(\frac{H^4}{M} + \text{Var}(V_\diamond^\pi(s, a)) \right) \\ &:= Z \end{aligned}$$

with probability at least $1 - \delta$. Taking the union bound over $(s, a, i) \in \mathcal{S} \times \mathcal{A} \times [t]$ and by Lemma 60, we have

$$\frac{1}{t} \sum_{i=0}^{t-1} \left[\hat{P}_i \hat{V}_\diamond^i - P \hat{V}_\diamond^i \right] (s, a) \leq \frac{H\ell}{tM} + \sqrt{Z}$$

with probability at least $1 - \delta$. ■

H.3.1. AUXILIARY LEMMAS FOR LEMMA 46

Lemma 48 For any $t \in [T]$ and $k \in [K]$,

$$0 \leq V_\square^{\pi_k^*} - V_\square^{\pi_t'} \leq \sum_{i=1}^t \left(\prod_{j=1}^{t-i} [\gamma P_{\pi_{t-j}}] \pi_i - (\gamma P_{\pi_k^*})^{t-i} \pi_k^* \right) \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_j \hat{V}_\square^j - \gamma P \hat{V}_\square^j] + \frac{H(1 + \lambda_k)}{t(t-1)} \mathbf{1}.$$

Proof The first inequality is due to the definition of π_k^* . For the second inequality, since we have

$$V_\square^{\pi_k^*} - V_\square^{\pi_t'} = \underbrace{V_\square^{\pi_k^*} - \frac{1}{t} \sum_{i=1}^t (\pi_t \hat{Q}_\square^i)}_{\text{Term (i)}} + \underbrace{\frac{1}{t} \sum_{i=1}^t (\pi_t \hat{Q}_\square^i) - V_\square^{\pi_t'}}_{\text{Term (ii)}}$$

$$\begin{aligned} V_\square^{\pi_k^*} - \pi_t \frac{1}{t} \sum_{i=1}^t \hat{Q}_\square^i &\leq (\pi_k^* Q_\square^{\pi_k^*}) - \frac{1}{t} \sum_{i=1}^t (\pi_k^* \hat{Q}_\square^i) && \text{(By the greediness of } \pi_t) \\ &= (\pi_k^* \square) + \gamma P_{\pi_k^*} V_\square^{\pi_k^*} - \frac{1}{t} \sum_{i=1}^t (\pi_k^* \hat{Q}_\square^i) \\ &= (\pi_k^* \square) + \gamma P_{\pi_k^*} V_\square^{\pi_k^*} - (\pi_k^* \square) - \gamma P_{\pi_k^*} \frac{1}{t} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) - \frac{1}{t} \sum_{i=0}^{t-1} [\gamma \hat{P}_{\pi_k^*}^i \hat{V}_\square^i - \gamma P_{\pi_k^*} \hat{V}_\square^i] \\ &&& \text{(By Lemma 39)} \\ &= \gamma P_{\pi_k^*} \left(V_\square^{\pi_k^*} - \frac{1}{t} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) \right) - \frac{1}{t} \sum_{i=0}^{t-1} [\gamma \hat{P}_{\pi_k^*}^i \hat{V}_\square^i - \gamma P_{\pi_k^*} \hat{V}_\square^i] \\ &= \gamma P_{\pi_k^*} \left(V_\square^{\pi_k^*} - \frac{1}{t} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) + \frac{1}{t-1} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) - \frac{1}{t-1} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) \right) \\ &\quad - \frac{1}{t} \sum_{i=0}^{t-1} [\gamma \hat{P}_{\pi_k^*}^i \hat{V}_\square^i - \gamma P_{\pi_k^*} \hat{V}_\square^i] \\ &\leq \gamma P_{\pi_k^*} \left(V_\square^{\pi_k^*} - \frac{1}{t-1} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) \right) - \frac{1}{t} \sum_{i=0}^{t-1} [\gamma \hat{P}_{\pi_k^*}^i \hat{V}_\square^i - \gamma P_{\pi_k^*} \hat{V}_\square^i] + \frac{H(1 + \lambda_k)}{t(t-1)} \mathbf{1} \\ &\quad (\|\gamma \pi_k^* P\|_1 \|\pi_{t-1} \hat{Q}_\square^i\|_\infty \leq H(1 + \lambda_k) \text{ for all } k \text{ and } i) \\ &\leq - \sum_{i=1}^t (\gamma P_{\pi_k^*})^{t-i} \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_{\pi_k^*}^j \hat{V}_\square^j - \gamma P_{\pi_k^*} \hat{V}_\square^j] + \frac{H(1 + \lambda_k)}{t(t-1)} \mathbf{1}. \\ &&& \text{(By induction (Lemma 52) and } \hat{V}_\square^0 = \mathbf{0}) \end{aligned}$$

Next, we bound term (ii). We define Q^π the Q-value function for a policy π being its unique fixed point.

$$\begin{aligned}
 \frac{1}{t} \sum_{i=1}^t (\pi_t \hat{Q}_\square^i) - V_\square^{\pi_t} &\leq \frac{1}{t} \sum_{i=1}^t (\pi_t \hat{Q}_\square^i) - \pi_t \prod_{i=1}^{t-1} \mathcal{T}^{\pi_i} Q_\square^{\pi_0} && \text{(From the definition of } \pi_t) \\
 &= \frac{1}{t} \sum_{i=1}^t (\pi_t \hat{Q}_\square^i) - (\pi_t \square) - \gamma P_{\pi_t} \pi_{t-1} \prod_{i=1}^{t-2} \mathcal{T}^{\pi_i} Q_\square^{\pi_0} \\
 &= (\pi_t \square) + \gamma P_{\pi_t} \frac{1}{t} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) + \frac{1}{t} \sum_{i=0}^{t-1} [\gamma \hat{P}_{\pi_t}^i \hat{V}_\square^i - \gamma P_{\pi_t} \hat{V}_\square^i] \\
 &\quad - (\pi_t \square) - \gamma P_{\pi_t} \pi_{t-1} \prod_{i=1}^{t-2} \mathcal{T}^{\pi_i} Q_\square^{\pi_0} && \text{(By Lemma 39)} \\
 &= \gamma P_{\pi_t} \left(\frac{1}{t} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) - V_\square^{\pi_{t-1}} \right) + \frac{1}{t} \sum_{i=0}^{t-1} [\gamma \hat{P}_{\pi_t}^i \hat{V}_\square^i - \gamma P_{\pi_t} \hat{V}_\square^i] \\
 &\leq \gamma P_{\pi_t} \left(\frac{1}{t-1} \sum_{i=0}^{t-2} (\pi_{t-1} \hat{Q}_\square^i) - V_\square^{\pi_{t-1}} \right) + \frac{1}{t} \sum_{i=0}^{t-1} [\gamma \hat{P}_{\pi_t}^i \hat{V}_\square^i - \gamma P_{\pi_t} \hat{V}_\square^i] \\
 &\leq \sum_{i=1}^t \prod_{j=1}^{t-i} [\gamma P_{\pi_{t-j}}] \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_{\pi_i}^j \hat{V}_\square^j - \gamma P_{\pi_i} \hat{V}_\square^j]. \\
 &&& \text{(By induction (Lemma 52) and } \hat{V}_\square^0 = 0)
 \end{aligned}$$

Thus, we obtain the second inequality. \blacksquare

Lemma 49 For any $t \in [T-1]$ and $k \in [K]$,

$$\hat{V}_\square^{t+1} \leq V_\square^{\pi_{t+1}} + \gamma^{t+1} t H (1 + \lambda_k) \mathbf{1} + \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (\hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} - P_{\pi_{t-i}} \hat{V}_\square^{t-i})$$

and

$$\hat{V}_\square^{t+1} \geq V_\square^{\pi_t} - \gamma^{t+1} t H (1 + \lambda_k) \mathbf{1} + \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (\hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} - P_{\pi_{t-i}} \hat{V}_\square^{t-i}).$$

Proof We first note that

$$\begin{aligned}
 \hat{V}_\square^{t+1} &= \sum_{i=0}^t (\pi_{t+1} \hat{Q}_\square^i) - \sum_{i=0}^{t-1} (\pi_t \hat{Q}_\square^i) && \text{(From the last line in Algorithm 6)} \\
 &\leq \sum_{i=0}^t (\pi_{t+1} \hat{Q}_\square^i) - \sum_{i=0}^{t-1} (\pi_{t+1} \hat{Q}_\square^i) && \text{(By the greediness of } \pi_t)
 \end{aligned}$$

$$\begin{aligned}
 &= (\pi_{t+1}\hat{Q}_\square^t) \\
 &= (\pi_{t+1}\square) + \gamma\hat{P}_{\pi_{t+1}}^t\hat{V}_\square^t \\
 &= (\pi_{t+1}\square) + \gamma P_{\pi_{t+1}}\hat{V}_\square^t + \gamma(\hat{P}_{\pi_{t+1}}^t\hat{V}_\square^t - P_{\pi_{t+1}}\hat{V}_\square^t) \\
 &\leq \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}}(\square + \gamma(\hat{P}_{\pi_{t-i}}^{t-i}\hat{V}_\square^{t-i} - P_{\pi_{t-i}}\hat{V}_\square^{t-i})). \quad (\text{By induction on } t)
 \end{aligned}$$

Let \mathcal{T}^π denote the Bellman operator with policy π , we have

$$\begin{aligned}
 \pi_{t+1} \prod_{i=0}^t \mathcal{T}^{\pi_i} Q_\square^{\pi_0} &= \sum_{i=1}^t \gamma^i \prod_{j=t-i-1}^t P_{\pi_{t-j}}(\pi_i \square) + \gamma^{t+1} \prod_{j=0}^t P_{\pi_{t-j}}(\pi_0 Q_\square^{\pi_0}) \\
 \implies \sum_{i=1}^t \gamma^i \prod_{j=t-i-1}^t P_{\pi_{t-j}}(\pi_i \square) &\leq \pi_{t+1} \prod_{i=0}^t \mathcal{T}^{\pi_i} Q_\square^{\pi_0} + \gamma^{t+1} tH(1 + \lambda_k)\mathbf{1}. \\
 & \quad (\text{Since } \prod_{j=0}^t P_{\pi_j} Q_\square^{\pi_0} \leq tH(1 + \lambda_k)\mathbf{1})
 \end{aligned}$$

Combining all above, we obtain

$$\hat{V}_\square^{t+1} \leq \pi_{t+1} \prod_{i=0}^t \mathcal{T}^{\pi_i} Q_\square^{\pi_0} + \gamma^{t+1} tH(1 + \lambda_k)\mathbf{1} + \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (\hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} - P_{\pi_{t-i}} \hat{V}_\square^{t-i})$$

Denoting $\pi'_{k,t}$ a non-stationary policy that follows $\pi_{t+1}, \pi_t, \pi_{t-1}, \dots$ sequentially, we simplify the above inequality as

$$\hat{V}_\square^{t+1} \leq V_\square^{\pi'_{k,t+1}} + \gamma^{t+1} tH(1 + \lambda_k)\mathbf{1} + \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (\hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} - P_{\pi_{t-i}} \hat{V}_\square^{t-i}).$$

Similarly,

$$\begin{aligned}
 \hat{V}_\square^{t+1} &= \sum_{i=1}^t (\pi_{t+1}\hat{Q}_\square^i) - \sum_{i=1}^{t-1} (\pi_t\hat{Q}_\square^i) \\
 &\geq \sum_{i=1}^t (\pi_t\hat{Q}_\square^i) - \sum_{i=1}^{t-1} (\pi_t\hat{Q}_\square^i) \quad (\text{By the greediness of } \pi_{t+1}) \\
 &= (\pi_t\hat{Q}_\square^t) \\
 &= (\pi_t\square) + \gamma\hat{P}_{\pi_t}^t\hat{V}_\square^t \\
 &= (\pi_t\square) + \gamma P_{\pi_t}\hat{V}_\square^t + \gamma(\hat{P}_{\pi_t}^t\hat{V}_\square^t - P_{\pi_t}\hat{V}_\square^t) \\
 &\geq \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}}(\square + \gamma(\hat{P}_{\pi_{t-i}}^{t-i}\hat{V}_\square^{t-i} - P_{\pi_{t-i}}\hat{V}_\square^{t-i})) \quad (\text{By induction on } t)
 \end{aligned}$$

and

$$\pi_{t+1} \prod_{i=1}^{t+1} \mathcal{T}^{\pi_{i-1}} Q_\square^{\pi_0} = \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}}(\pi_i \square) + \gamma^{t+1} \prod_{j=1}^t P_{\pi_{t-j+1}}(\pi_0 Q_\square^{\pi_0})$$

$$\begin{aligned} \implies \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}}(\pi_i \square) &\geq \pi_{t+1} \prod_{i=1}^{t+1} \mathcal{T}^{\pi_{i-1}} Q_{\square}^{\pi_0} - \gamma^{t+1} t H(1 + \lambda_k) \mathbf{1}. \\ &\text{(Since } \prod_{j=1}^t P_{\pi_{j-1}} Q_{\square}^{\pi_0} \leq \gamma^{t+1} t H(1 + \lambda_k) \mathbf{1}) \end{aligned}$$

Combining the above, we obtain

$$\begin{aligned} \hat{V}_{\square}^{t+1} &\geq \pi_t \prod_{i=1}^{t+1} \mathcal{T}^{\pi_{i-1}} Q_{\square}^{\pi_0} - \gamma^{t+1} t H(1 + \lambda_k) \mathbf{1} + \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (\hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_{\square}^{t-i} - P_{\pi_{t-i}} \hat{V}_{\square}^{t-i}) \\ &= V_{\square}^{\pi_t'} - \gamma^{t+1} t H(1 + \lambda_k) \mathbf{1} + \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (\hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_{\square}^{t-i} - P_{\pi_{t-i}} \hat{V}_{\square}^{t-i}) \end{aligned}$$

■

Lemma 50 For any $t \in [T]$ and $k \in [K]$,

$$\begin{aligned} V_{\square}^{\pi_k^*} - \hat{V}_{\square}^t &\leq \left(\gamma^t t + \frac{1 + \lambda_k}{t(t-1)} \right) H \mathbf{1} - \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (P_{\pi_{t-i}} \hat{V}_{\square}^{t-i} - \hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_{\square}^{t-i}) \\ &\quad + \sum_{i=1}^t \left((\gamma P_{\pi_k^*})^{t-i} \pi_k^* - \prod_{j=1}^{t-i} [\gamma P_{\pi_{t-j}}] \pi_i \right) \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_j \hat{V}_{\square}^j - \gamma P \hat{V}_{\square}^j] \end{aligned}$$

and

$$V_{\square}^{\pi_k^*} - \hat{V}_{\square}^t \geq -\gamma^t t H(1 + \lambda_k) \mathbf{1} - \sum_{i=1}^{t-1} \gamma^i \prod_{j=t-i}^{t-1} P_{\pi_{t-j}} \gamma (P_{\pi_{t-i-1}} \hat{V}_{\square}^{t-i-1} - \hat{P}_{\pi_{t-i-1}}^{t-i-1} \hat{V}_{\square}^{t-i-1}).$$

Proof From Lemma 49, we know

$$\begin{aligned} \hat{V}_{\square}^t &\leq V_{\square}^{\pi_t'} + \gamma^t t H(1 + \lambda_k) \mathbf{1} + \sum_{i=1}^{t-1} \gamma^i \prod_{j=t-i}^{t-1} P_{\pi_{t-j}} \gamma (P_{\pi_{t-i-1}} \hat{V}_{\square}^{t-i-1} - \hat{P}_{\pi_{t-i-1}}^{t-i-1} \hat{V}_{\square}^{t-i-1}) \\ &\leq V_{\square}^{\pi_k^*} + \gamma^t t H(1 + \lambda_k) \mathbf{1} + \sum_{i=1}^{t-1} \gamma^i \prod_{j=t-i}^{t-1} P_{\pi_{t-j}} \gamma (P_{\pi_{t-i-1}} \hat{V}_{\square}^{t-i-1} - \hat{P}_{\pi_{t-i-1}}^{t-i-1} \hat{V}_{\square}^{t-i-1}) \end{aligned}$$

which gives us the second inequality. From Lemma 49 and Lemma 48 we have,

$$V_{\square}^{\pi_t'} - \hat{V}_{\square}^{t+1} \leq \gamma^t t H(1 + \lambda_k) \mathbf{1} - \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (P_{\pi_{t-i}} \hat{V}_{\square}^{t-i} - \hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_{\square}^{t-i})$$

and

$$V_{\square}^{\pi_k^*} - V_{\square}^{\pi_t'} \leq \sum_{i=1}^t \left(\prod_{j=1}^{t-i} [\gamma P_{\pi_{t-j}}] \pi_i - (\gamma P_{\pi_k^*})^{t-i} \pi_k^* \right) \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_j \hat{V}_{\square}^j - \gamma P \hat{V}_{\square}^j] + \frac{H(1 + \lambda_k)}{t(t-1)} \mathbf{1}.$$

Combining them gives us the upper bound. ■

Lemma 51 For any $t \geq 2 \log(t)/\gamma$ and $k \in [K]$,

$$\sigma\left(\hat{V}_\square^t\right) \leq \left(\frac{2}{t} + 4H\sqrt{\frac{t}{M}}\right) H(1 + \lambda_k)\mathbf{1} + \sigma(V_\square^{\pi_k^*})$$

with probability at least $1 - \delta$.

Proof We denote $\iota = \log(2|\mathcal{S}||\mathcal{A}|/\delta)$ throughout the proof. By Lemma 50,

$$\begin{aligned} V_\square^{\pi_k^*} - \hat{V}_\square^t &\leq \underbrace{\left(\gamma^t t + \frac{1}{t(t-1)}\right) H(1 + \lambda_k)\mathbf{1}}_{\text{Term (i)}} - \underbrace{\sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (\hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} - P_{\pi_{t-i}} \hat{V}_\square^{t-i})}_{\text{Term (ii)}} \\ &\quad + \underbrace{\sum_{i=1}^t \left((\gamma P_{\pi_k^*})^{t-i} \pi_k^* - \prod_{j=1}^{t-i} [\gamma P_{\pi_{t-j}}] \pi_i \right) \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_j \hat{V}_\square^j - \gamma P \hat{V}_\square^j]}_{\text{Term (iii)}} \end{aligned} \quad (42)$$

We first bound Term (ii) and Term (iii). By Azuma-Hoeffding's inequality (Lemma 58), we have

$$\left\| P_{\pi_{t-i}} \hat{V}_\square^{t-i} - \hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} \right\|_\infty \leq 2H(1 + \lambda_k) \sqrt{\frac{t}{M}},$$

and by Lemma 57 with $t = i$, we have

$$\left\| \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_{\pi_i}^j \hat{V}_\square^j - \gamma P_{\pi_i} \hat{V}_\square^j] \right\|_\infty \leq 2H(1 + \lambda_k) \sqrt{\frac{t}{iM}}$$

each with probability at least $1 - \delta$. Thus, to bound Term (ii), we have

$$\left\| \sum_{i=1}^t \gamma^i \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma (\hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} - P_{\pi_{t-i}} \hat{V}_\square^{t-i}) \right\|_\infty \quad (43)$$

$$\begin{aligned} &\leq \sum_{i=1}^t \gamma^i \left\| \prod_{j=t-i+1}^t P_{\pi_{t-j}} \gamma \right\|_1 \left\| \hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} - P_{\pi_{t-i}} \hat{V}_\square^{t-i} \right\|_\infty \\ &\leq \sum_{i=1}^t \gamma^i \left\| \hat{P}_{\pi_{t-i}}^{t-i} \hat{V}_\square^{t-i} - P_{\pi_{t-i}} \hat{V}_\square^{t-i} \right\|_\infty \end{aligned}$$

$$\leq 2H^2(1 + \lambda_k) \sqrt{\frac{t}{M}} \quad (44)$$

and to bound Term (iii) we have

$$\left\| \sum_{i=1}^t \left((\gamma P_{\pi_k^*})^{t-i} \pi_k^* - \prod_{j=1}^{t-i} [\gamma P_{\pi_{t-j}}] \pi_i \right) \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_j \hat{V}_\square^j - \gamma P \hat{V}_\square^j] \right\|_\infty$$

$$\begin{aligned}
 &\leq \sum_{i=1}^t \gamma^{t-i} \left\| \prod_{j=1}^{t-i} [P_{\pi_{t-j}}] \pi_i - (P_{\pi_k^*})^{t-i} \pi_k^* \right\|_1 \left\| \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_j \hat{V}_\square^j - \gamma P \hat{V}_\square^j] \right\|_\infty \\
 &\leq \sum_{i=1}^t \gamma^{t-i} \left\| \frac{1}{i} \sum_{j=0}^{i-1} [\gamma \hat{P}_j \hat{V}_\square^j - \gamma P \hat{V}_\square^j] \right\|_\infty \\
 &\leq \sum_{i=1}^t \gamma^{t-i} \sqrt{\frac{4H^2(1+\lambda_k)^2 \iota}{iM}} \\
 &\leq \sum_{i=1}^t \gamma^{t-i} \sqrt{\frac{4H^2(1+\lambda_k)^2 \iota}{M}} \\
 &\leq 2H^2(1+\lambda_k) \sqrt{\frac{\iota}{M}}
 \end{aligned} \tag{45}$$

each with probability at least $1 - \delta$. Lastly we bound Term (i). For $t \geq 2 \log(t)/\gamma$, we have

$$\gamma^t \leq \frac{1}{t^2}$$

and thus

$$\gamma^t t + \frac{1}{t(t-1)} \leq \frac{1}{t} + \frac{1}{t} = \frac{2}{t}. \tag{46}$$

Combining Eqs. (42) and (44) to (46), we have

$$|V_{\square}^{\pi_k^*} - \hat{V}_{\square}^t| \leq \left(\frac{2}{t} + 4H \sqrt{\frac{\iota}{M}} \right) (1 + \lambda_k) H \mathbf{1}. \tag{47}$$

Finally, we have

$$\begin{aligned}
 \sigma(\hat{V}_{\square}^t) &\leq \sigma(V_{\square}^{\pi_k^*} - \hat{V}_{\square}^t) + \sigma(V_{\square}^{\pi_k^*}) && \text{(By Lemma 62)} \\
 &\leq |V_{\square}^{\pi_k^*} - \hat{V}_{\square}^t| + \sigma(V_{\square}^{\pi_k^*}) && \text{(By Lemma 61)} \\
 &\leq \left(\frac{2}{t} + 4H \sqrt{\frac{\iota}{M}} \right) (1 + \lambda_k) H \mathbf{1} + \sigma(V_{\square}^{\pi_k^*})
 \end{aligned}$$

with probability at least $1 - \delta$, which completes the proof. \blacksquare

Lemma 52 (Induction Lemma) *Assume $X_k, A_k, B_k \geq 0$, $k = 1, \dots$, and $X_{k+1} \leq A_k X_k + B_k$, then we have $X_{k+1} \leq \prod_{i=1}^k A_i X_1 + \sum_{i=1}^k \prod_{j=i+1}^k A_j B_i$.*

H.3.2. AUXILIARY LEMMAS FOR LEMMA 47

Lemma 53 *For any $t \in [T]$,*

$$\|\hat{V}_{\diamond}^t - V_{\diamond}^{\pi}\|_{\infty} \leq \tilde{O} \left(\frac{H^2}{\sqrt{M}} + \gamma^t H \right)$$

with probability at least $1 - \delta$.

Proof

$$\begin{aligned}
 \hat{\mathcal{V}}_\diamond^t &= (\pi \hat{\mathcal{Q}}_\diamond^{t-1}) \\
 &= (\pi \diamond) + \gamma \hat{P}_\pi^{t-1} \hat{\mathcal{V}}_\diamond^{t-1} \\
 &= (\pi \diamond) + \gamma P_\pi \hat{\mathcal{V}}_\diamond^{t-1} + \gamma (\hat{P}_\pi^{t-1} \hat{\mathcal{V}}_\diamond^{t-1} - P_\pi \hat{\mathcal{V}}_\diamond^{t-1}) \\
 &= \sum_{i=0}^t \gamma^i (P_\pi)^i ((\pi \diamond) + \gamma (\hat{P}_\pi^{t-i-1} \hat{\mathcal{V}}_\diamond^{t-i-1} - P_\pi \hat{\mathcal{V}}_\diamond^{t-i-1})) \quad (\text{By induction on } t) \\
 &= \sum_{i=0}^t \gamma^i (P_\pi)^i (\pi \diamond) + \sum_{i=0}^t \gamma^{i+1} (P_\pi)^i (\hat{P}_\pi^{t-i-1} \hat{\mathcal{V}}_\diamond^{t-i-1} - P_\pi \hat{\mathcal{V}}_\diamond^{t-i-1}) \\
 &= V_\diamond^\pi - \sum_{i=t+1}^{\infty} \gamma^i (P_\pi)^i (\pi \diamond) + \sum_{i=0}^t \gamma^{i+1} (P_\pi)^i (\hat{P}_\pi^{t-i-1} \hat{\mathcal{V}}_\diamond^{t-i-1} - P_\pi \hat{\mathcal{V}}_\diamond^{t-i-1}).
 \end{aligned}$$

Note that with probability at least $1 - \delta$

$$\left\| \sum_{i=0}^t \gamma^{i+1} (P_\pi)^i (\hat{P}_\pi^{t-i-1} \hat{\mathcal{V}}_\diamond^{t-i-1} - P_\pi \hat{\mathcal{V}}_\diamond^{t-i-1}) \right\|_\infty \leq 2H^2 \sqrt{\frac{t}{M}}$$

by a similar argument as in Eq. (44), and

$$\left\| \sum_{i=t+1}^{\infty} \gamma^i (P_\pi)^i \pi \diamond \right\|_\infty \leq \gamma^t H.$$

We conclude that

$$\left\| \hat{\mathcal{V}}_\diamond^t - V_\diamond^\pi \right\|_\infty \leq \tilde{O} \left(\frac{H^2}{\sqrt{M}} + \gamma^t H \right)$$

with probability at least $1 - \delta$. ■

Lemma 54 For any $i \in [T]$, we have

$$\sigma(\hat{\mathcal{V}}_\diamond^i) \leq \tilde{O} \left(\frac{H^2}{\sqrt{M}} + \gamma^t H \right) \mathbf{1} + \sigma(V_\diamond^\pi)$$

with probability at least $1 - \delta$.

Proof We have

$$\begin{aligned}
 \sigma(\hat{\mathcal{V}}_\diamond^i) &\leq \sigma(V_\diamond^\pi - \hat{\mathcal{V}}_\diamond^i) + \sigma(V_\diamond^\pi) && (\text{By Lemma 62}) \\
 &\leq |V_\diamond^\pi - \hat{\mathcal{V}}_\diamond^i| + \sigma(V_\diamond^\pi) && (\text{By Lemma 61}) \\
 &\leq \tilde{O} \left(\frac{H^2}{\sqrt{M}} + \gamma^t H \right) \mathbf{1} + \sigma(V_\diamond^\pi) && (\text{By Lemma 53})
 \end{aligned}$$

with probability at least $1 - \delta$. ■

H.4. Proof of Corollary 31

Corollary 55 *Let Alg. 6 and Alg. 7 be the instantiations of the MDP-Solver and PolicyEvaluation in Alg. 1. For a fixed $\varepsilon \in (0, 1/H^2]$, $\delta \in (0, 1)$, Alg. 1 with $\tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^3}{\varepsilon^2}\right)$ samples, $U = O\left(\frac{1}{\zeta(1-\gamma)}\right)$, $\eta = \frac{U(1-\gamma)}{\sqrt{K}}$, $K = O\left(\frac{1}{\varepsilon^2(1-\gamma)^2}\right)$, and $b' = b - O(\varepsilon)$, returns a policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,*

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - O(\varepsilon), \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b - O(\varepsilon).$$

Under the same conditions, but with $b' = b + O(\varepsilon)$ and $\tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^5}{\zeta^2\varepsilon^2}\right)$ samples, Alg. 1 returns a policy $\bar{\pi}$ satisfying the following condition with probability $1 - \delta$,

$$V_r^{\bar{\pi}}(\rho) \geq V_r^{\pi^*}(\rho) - O(\varepsilon), \quad \text{and} \quad V_c^{\bar{\pi}}(\rho) \geq b.$$

Proof By Lemma 29 and Lemma 30, the sample complexity required to ensure $f(\mathcal{B}) \leq O(\varepsilon)$ is $TM|\mathcal{C}| = \tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|H^3}{\varepsilon^2}\right)$. Therefore, the guarantee for the relaxed feasibility setting follows directly from our meta-theorem (Theorem 4). For the strict feasibility setting, we rescale ε by a factor of $O(\zeta(1-\gamma))$. Since $\varepsilon \leq 1$ and $1-\gamma \leq 1$, the condition of $f(\mathcal{B}) \leq \zeta/6$ in Theorem 4 can be satisfied. The rescaling increases the sample complexity by a multiplicative factor of $\frac{1}{\zeta^2(1-\gamma)^2}$, thereby completing the proof. \blacksquare

H.5. Instantiating the MDP-Solver: Model-based algorithm (Li et al., 2020)

Instead of using MDVI-Tabular, the tabular MDP-Solver subroutine in Algorithm 1 can be instantiated with any model-based method that computes an optimal policy with respect to the estimated model. In this subsection, we adapt the framework analyzed in (Li et al., 2020) to show that, when combined with our overall framework, certain model-based MDP-Solver algorithms can recover the near-optimal sample complexity for solving tabular constrained MDPs.

Since we are using model-based methods, we denote \hat{P} as the probability transition kernel form by

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

where $(s_{s,a}^i)_{i=1}^N$ are the next-state samples from $\mathcal{B} = \text{DataCollection}(\text{Gen}, \mathcal{S} \times \mathcal{A}, N)$. Denote the perturbed reward by

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

where $\text{Unif}(0, \xi)$ denotes the uniform distribution. For any policy π , denote \hat{V}_p^π the corresponding value function of the perturbed empirical MDP $\widehat{\mathcal{M}}_p = (\mathcal{S}, \mathcal{A}, \hat{P}, r_p, \gamma)$. Denote $\hat{\pi}_p^*$ the optimal policy w.r.t. $\widehat{\mathcal{M}}_p$ (i.e. $\hat{\pi}_p^* = \arg \max_\pi \hat{V}_p^\pi$). Their main result is stated as follows.

Theorem 56 (Theorem 1 in (Li et al., 2020)) *There exist some universal constants $c_0, c_1 > 0$ such that: for any $\delta > 0$ and any $0 < \varepsilon \leq \frac{1}{1-\gamma}$, the policy $\hat{\pi}_p^*$ defined in (9) obeys*

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

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

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

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

Therefore, let $\mathcal{B} = \text{DataCollection}(\text{Gen}, \mathcal{S} \times \mathcal{A}, N)$. Then, by instantiating $\text{MDP-Solver}(r + \lambda_k c, \mathcal{B}, \phi)$ with any model-based algorithm that returns an optimal policy with respect to the perturbed empirical MDP constructed from \mathcal{B} , Assumption 2 can be satisfied with $f_{\text{mdp}}(\mathcal{B}) = O(\varepsilon)$. As a consequence, we recover the near-optimal sample complexity bounds for solving tabular constrained MDPs via our meta-theorem (Theorem 4). Furthermore, the limited range of ε (i.e. $(0, 1/H^2]$) in Corollary 31 will be improved to a full range (i.e. $(0, H]$).

Appendix I. Supporting Lemmas

I.1. Concentration Inequalities

The following lemma is used throughout the paper. In the linear setting, we take \mathcal{C} to be the core set and set $R = (1 + \lambda_k)H$. In the tabular setting, we let $\mathcal{C} = \mathcal{S} \times \mathcal{A}$ and set $R = H$.

Lemma 57 *Let \hat{V}^i be an empirical value function with entries bounded in $[0, R]$, and let $\mathcal{C} \subseteq \mathcal{S} \times \mathcal{A}$. Then, for any $t \in 1, \dots, T$, the following holds:*

$$\mathbb{P} \left(\exists (s, a) \in \mathcal{C} \text{ s.t. } \frac{1}{t} \sum_{i=0}^{t-1} [(\hat{P}_i \hat{V}^i)(s, a) - (P \hat{V}^i)(s, a)] \geq 2R \sqrt{\log(2|\mathcal{C}|/\delta)/tM} \right) \leq \delta$$

Proof Consider a fixed $t \in \{1, \dots, T\}$ and $(s, a) \in \mathcal{C}$. Denote $y_{t,m,s,a}$ as the m 'th next-state sample we collect for state-action pair (s, a) at iteration t . Since

$$\begin{aligned} \frac{1}{t} \sum_{i=0}^{t-1} [(\hat{P}_i \hat{V}^i)(s, a) - (P \hat{V}^i)(s, a)] &= \frac{1}{t} \sum_{i=0}^{t-1} \frac{1}{M} \sum_{m=1}^M [\hat{V}^i(y_{t,m,s,a}) - (P \hat{V}^i)(s, a)] \\ &= \frac{1}{t} \sum_{i=0}^{t-1} \frac{1}{M} \sum_{m=1}^M [\hat{V}^i(y_{t,m,s,a}) - (P \hat{V}^i)(s, a)] \end{aligned}$$

is a sum of bounded martingale differences with respect to the filtration $(\mathcal{F}_{i,m})_{i=0,m=1}^{t-1,M}$. Thus, using the Azuma-Hoeffding inequality (Lemma 58),

$$\mathbb{P} \left(\frac{1}{t} \sum_{i=0}^{t-1} \frac{1}{M} \sum_{m=1}^M [\hat{V}^i(y_{t,m,s,a}) - (P \hat{V}^i)(s, a)] \geq 2R \sqrt{\frac{\log(2|\mathcal{C}|/\delta)}{tM}} \right) \leq \frac{\delta}{|\mathcal{C}|}.$$

Taking the union bound over $(s, a) \in \mathcal{C}$

$$\begin{aligned} & \mathbb{P} \left(\max_{(s,a) \in \mathcal{C}} \frac{1}{t} \sum_{i=0}^{t-1} \frac{1}{M} \sum_{m=1}^M \left[\hat{V}^i(y_{t,m,s,a}) - (P\hat{V}^i)(s,a) \right] \leq 2R \sqrt{\frac{\log(2|\mathcal{C}|/\delta)}{tM}} \right) \\ & \geq 1 - \sum_{(s,a) \in \mathcal{C}} \mathbb{P} \left(\frac{1}{t} \sum_{i=0}^{t-1} \frac{1}{M} \sum_{m=1}^M \left[\hat{V}^i(y_{t,m,s,a}) - (P\hat{V}^i)(s,a) \right] \geq 2R \sqrt{\frac{\log(2|\mathcal{C}|/\delta)}{tM}} \right) \\ & \geq 1 - \delta, \end{aligned}$$

which implies the desired result. \blacksquare

Lemma 58 (Azuma-Hoeffding Inequality) Consider a real-valued stochastic process $(X_n)_{n=1}^N$ adapted to a filtration $(\mathcal{F}_n)_{n=1}^N$. Assume that $X_n \in [l_n, u_n]$ and $\mathbb{E}_n[X_n] = 0$ almost surely, for all n . Then,

$$\mathbb{P} \left(\sum_{n=1}^N X_n \geq \sqrt{\sum_{n=1}^N \frac{(u_n - l_n)^2}{2} \log \frac{1}{\delta}} \right) \leq \delta$$

for any $\delta \in (0, 1)$.

Lemma 59 (Bernstein's Inequality) Consider a real-valued stochastic process $(X_n)_{n=1}^N$ adapted to a filtration $(\mathcal{F}_n)_{n=1}^N$. Suppose that $X_n \leq U$ and $\mathbb{E}_n[X_n] = 0$ almost surely, for all n . Then, letting $Z' := \sum_{n=1}^N \mathbb{E}_n[X_n^2]$,

$$\mathbb{P} \left(\sum_{n=1}^N X_n \geq \frac{2U}{3} \log \frac{1}{\delta} + \sqrt{2Z' \log \frac{1}{\delta}} \text{ and } Z' \leq Z \right) \leq \delta$$

for any $Z \in [0, \infty)$ and $\delta \in (0, 1)$.

Lemma 60 (Conditional Bernstein's Inequality) Consider the same notations and assumptions in Lemma 59. Furthermore, let \mathcal{E} be an event that implies $Z' \leq Z$ for some $Z \in [0, \infty)$ with $\mathbb{P}(\mathcal{E}) \geq 1 - \delta'$ for some $\delta' \in (0, 1)$. Then,

$$\mathbb{P} \left(\sum_{n=1}^N X_n \geq \frac{2U}{3} \log \frac{1}{\delta(1-\delta')} + \sqrt{2Z \log \frac{1}{\delta(1-\delta')}} \mid \mathcal{E} \right) \leq \delta$$

for any $\delta \in (0, 1)$.

I.2. Lemmas for Variances

Lemma 61 (Popoviciu's Inequality for Variances) The variance of any random variable bounded by x is bounded by x^2 .

Lemma 62 (Azar et al. (2013)) Suppose two real-valued random variables X, Y whose variances, $\mathbb{V}X$ and $\mathbb{V}Y$, exist and are finite. Then, $\sqrt{\mathbb{V}X} \leq \sqrt{\mathbb{V}[X - Y]} + \sqrt{\mathbb{V}Y}$.

Lemma 63 (Total variance lemma Azar et al. (2013)) For any policy π , $\|(I - P_\pi)^{-1} \sigma(V^\pi)\|_\infty \leq \sqrt{2H^3}$.

I.3. Lemmas for Constrained MDPs

Lemma 64 (Constraint violation bound, Lemma B.2 in Jain et al. (2022)) For any $C \geq \lambda^*$ and any π s.t. $V_r^*(\rho) - V_r^\pi(\rho) + C[b - V_c^\pi(\rho)]_+ \leq \beta$, we have $[b - V_c^\pi(\rho)]_+ \leq \frac{\beta}{C - \lambda^*}$.

Lemma 65 (Bounding the dual variable, Lemma 4.1 in Jain et al. (2022)) The objective Eq. (1) satisfies strong duality, and the optimal dual variables are bounded as

$$\lambda^* \leq \frac{1}{(1 - \gamma)\zeta}, \quad \text{where } \zeta := \max_{\pi} V_c^\pi(\rho) - b > 0.$$

Lemma 66 (Bounding the sensitivity error, Lemma 13 in Vaswani et al. (2022)) If we have

$$\hat{\pi}^* \in \arg \max_{\pi} V_r^\pi(\rho) \text{ s.t. } V_c^\pi(\rho) \geq b + \Delta$$

$$\tilde{\pi}^* \in \arg \max_{\pi} V_r^\pi(\rho) \text{ s.t. } V_c^\pi(\rho) \geq b - \Delta,$$

then the sensitivity error term can be bounded by:

$$\left| V_r^{\hat{\pi}^*}(\rho) - V_r^{\tilde{\pi}^*}(\rho) \right| \leq 2\Delta\lambda^*$$

where λ^* is the optimal Lagrange multiplier (i.e., the solution to Eq. (4)).