
Provably Efficient Model-Free Constrained RL with Linear Function Approximation

Arnob Ghosh

Electrical and Computer Engineering
The Ohio State University
Columbus, OH, USA
ghosh.244@osu.edu

Xingyu Zhou

Electrical and Computer Engineering
Wayne State University
Detroit, MI, USA
xingyu.zhou@wayne.edu

Ness Shroff

Electrical and Computer Engineering
The Ohio State University
Columbus, OH, USA
shroff.11@osu.edu

Abstract

We study the constrained reinforcement learning problem, in which an agent aims to maximize the expected cumulative reward subject to a constraint on the expected total value of a utility function. In contrast to existing model-based approaches or model-free methods accompanied with a ‘simulator’, we aim to develop the *first model-free, simulator-free* algorithm that achieves a sublinear regret and a sublinear constraint violation even in *large-scale* systems. To this end, we consider the episodic constrained Markov decision processes with linear function approximation, where the transition dynamics and the reward function can be represented as a linear function of some known feature mapping. We show that $\tilde{O}(\sqrt{d^3 H^3 T})$ regret and $\tilde{O}(\sqrt{d^3 H^3 T})$ constraint violation bounds can be achieved, where d is the dimension of the feature mapping, H is the length of the episode, and T is the total number of steps. Our bounds are attained without explicitly estimating the unknown transition model or requiring a simulator, and they depend on the state space only through the dimension of the feature mapping. Hence our bounds hold even when the number of states goes to infinity. Our main results are achieved via novel adaptations of the standard LSVI-UCB algorithms. In particular, we first introduce primal-dual optimization into the LSVI-UCB algorithm to balance between regret and constraint violation. More importantly, we replace the standard greedy selection with respect to the state-action function in LSVI-UCB with a soft-max policy. This turns out to be key in establishing uniform concentration for the constrained case via its approximation-smoothness trade-off. Finally, we also show that one can achieve an even zero constraint violation for large enough T by trading the regret a little bit but still maintaining the same order with respect to T .

1 Introduction

In many practical applications of online reinforcement learning (RL) (e.g., financial regulations, safety), there exist additional constraints on the learned policy in the sense that it also needs to ensure that the expected total utility (cost, resp.) exceeds a given threshold (is below a threshold, resp.). Constrained RL problem is formulated as a constrained Markov Decision Process (CMDP), in which the celebrated exploration-exploitation trade-off in online RL becomes more challenging due to the additional need to find a balance between reward regret and constraint violations.

To develop online sample-efficient algorithms for CMDPs, prior works have largely resorted to model-based approaches, where the control policy is constructed based on a learned model [1–7]. However, due to the explicit estimation and storage of the unknown transition model, model-based approaches often lead to large time and space complexities. These issues are exacerbated in the large state space. Hence, there are several recent works starting to investigate model-free algorithms for CMDPs, which directly update the value function or the policy without first estimating the model [8–10]. However, all of these works consider an easier setting compared to standard RL in that they assume access to a simulator [11] (a.k.a. a generative model [12]), which is a strong oracle that allows the agent to query arbitrary state-action pairs and return the reward and the next state, hence greatly alleviating the intrinsic difficulty of exploration in RL. To the best of our knowledge, [13] were the first to study model-free and simulator-free algorithms for CMDPs. The above mentioned work on model-free algorithm considers the finite-state tabular setting and the regret scales polynomially with the number of states. Thus, the result *would not be useful* for large-scale RL applications where the number of states could even be infinite. To address this curse of dimensionality, modern RL has adopted *function approximation* techniques to approximate the (action-)value function of a policy, which greatly expands the potential reach of RL, especially via deep neural networks. However, little is known for the performance guarantee of *model-free algorithms* in CMDPs beyond tabular settings, even in the case of linear function approximation. Motivated by this, we are interested in the following question:

Can we achieve provably sample-efficient and model-free exploration for CMDPs beyond tabular settings (without a simulator)?

Contribution. To answer the above question, we consider the episodic CMDPs with linear function approximation, where the transition dynamics and the reward function can be represented as a linear function of some known feature mapping. Our main contributions are as follows.

- We show that with a proper parameter choice, our proposed algorithm achieves $\tilde{O}(\sqrt{d^3 H^3 T})$ regret and $\tilde{O}(\sqrt{d^3 H^3 T})$ constraint violation bounds with a high probability, where d is the dimension of the feature mapping, H is the length of the episode, and T is the total number of steps. We also show that it is in fact possible to achieve *zero* constraint violation by trading the regret a little bit while maintaining the same order with respect to T .
- Our bounds are attained without explicitly estimating the unknown transition model or *requiring a simulator*, and they depend on the state space only through the dimension of the feature mapping. To the best of knowledge, these sub-linear bounds are the first results for model-free, simulator-free online RL algorithms for CMDPs with function approximations. We even improve the bound ($\tilde{O}(T^{0.8})$) proposed in the model-free finite state tabular setting by [13] (Table 1).
- We combine the primal-dual algorithm with the classic model-free, simulator-free LSVI-UCB algorithm [14] to balance between regret and constraint violations. This naturally leads to the construction of a new composite state-action function (i.e., Q -function), which is the sum of the Q -function for the reward and the Q -function for the utility, weighted by the dual variable. Due to this new type of Q -function in CMDPs, a key challenge arises when establishing the value-aware uniform concentration, which lies at the heart of the performance analysis of model-free exploration. *More specifically, the standard greedy selection with respect to this composite Q -function fails in finding non-trivial covering number for the function class of individual value functions (i.e., V -function) for the reward and the utility respectively.* To address this fundamental issue, we instead adopt a soft-max policy by utilizing its nice property of approximation-smoothness trade-off via its parameter, i.e., temperature coefficient.

1.1 Related Work

Model-based RL algorithms have been proposed for the CMDP [1–3, 5–7]. Apart from [7], the rest considered tabular set-up. In the tabular model-based set-up, the best known regret and constraint violations achieved are $\tilde{O}(\sqrt{|\mathcal{S}|^2 |\mathcal{A}| T})$ where $|\mathcal{S}|$ and $|\mathcal{A}|$ are the dimensions of the state and action spaces respectively. Hence, such results can not cope up with the large state space. [7] considered linear kernel MDP whereas we consider linear MDP. These two are not the same in general. We describe the differences with the algorithm proposed in [7] in Section 4.1.

Model-free RL algorithms have also been proposed [8–10] to solve CMDP. However, all of the above require a generator model, which simulates from any state and action. [13] proposed a ‘triple-Q’

Table 1: Regret and Constraint Violations on Episodic MDP for algorithms which do not use simulators. (LFA: Linear Function Approximation, MOD-FREE: Model-Free)

ALGORITHM	LFA?	MOD-FREE?	REGRET	VIOLATIONS
OPDOP [7]	YES	×	$\tilde{O}(\sqrt{d^2 H^5 T})$	$\tilde{O}(\sqrt{d^2 H^5 T})$
OPT-PRIMALDUAL CMDP †[1]	NO	×	$\tilde{O}(\sqrt{H^5 S ^2 \mathcal{A} T})$	$\tilde{O}(\sqrt{H^5 S ^2 \mathcal{A} T})$
OPTDUAL-CMDP †[1]	NO	×	$\tilde{O}(\sqrt{H^3 S ^2 \mathcal{A} T})$	$\tilde{O}(\sqrt{H^3 S ^2 \mathcal{A} T})$
OPTPRESS-PRIMALDUAL [6]	NO	×	$\tilde{O}(\sqrt{H^5 S ^3 \mathcal{A} T})$	$O(1)$
TRIPLE-Q [13]	NO	✓	$\tilde{O}(H^{3.2} T^{0.8} \sqrt{ S \mathcal{A} })$	0
OUR APPROACH	YES	✓	$\tilde{O}(\sqrt{d^3 H^3 T})$	$\tilde{O}(\sqrt{d^3 H^3 T})^*$

* WE CAN REDUCE THE VIOLATION TO 0 FOR LARGE ENOUGH T (FINITE) WHILE MAINTAINING THE SAME ORDER OF REGRET WITH RESPECT TO T (APPENDIX H).

† WE REPLACE ρ (ξ IN OUR PAPER) BY $\mathcal{O}(H)$ SIMILAR TO OUR PAPER AND \mathcal{N} BY $|S|$.

algorithm which does not require a ‘simulator’. However, it only considered tabular setting. The regret bound shown in [13] is $\tilde{O}(T^{0.8})$ which is far from the optimal for the model-based case $\tilde{O}(\sqrt{T})$. Please see Table 1 to see our contribution compared to the state-of-the-art approaches. [15] proposed a RL algorithm for the scenario where a constraint needs to be satisfied at each step of an episode. We consider a constraint where the cumulative utility over the length of the episodes must exceed a threshold. Hence, the set of constraints is fundamentally different. The authors in [15] also assumed that a safe-action is known for each state which we do not assume in our setting.

2 Problem Formulation

We consider an episodic constrained MDP, denoted by $(\mathcal{S}, \mathcal{A}, \mathbb{P}, H, r, g)$ where \mathcal{S} is the state space, \mathcal{A} is the action space, H is the fixed length of each episode, $\mathbb{P} = \{\mathbb{P}_h\}_{h=1}^H$ is a collection of transition probability measures, $r = \{r_h\}_{h=1}^H$ is a collection of reward functions, and $g = \{g_h\}_{h=1}^H$ is a collection of utility functions. We assume that \mathcal{S} is a measurable space with possibly infinite number of elements, \mathcal{A} is a finite action set. $\mathbb{P}_h(\cdot|x, a)$ is the transition probability kernel which denotes the probability to reach a state when action a is taken at state x . $r_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, and $g_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ and are assumed to be deterministic. However, we can readily extend to settings when r_h and g_h are random.

Each episode $k \in [K]$ starts with the fixed state x_1 . It can be readily generalized to the setting where x_1 is drawn from a distribution. Then at each step $h \in [H]$ in episode k , the agent observes state $x_h^k \in \mathcal{S}$, picks an action $a_h^k \in \mathcal{A}$, receives a reward $r_h(x_h^k, a_h^k)$, and a utility $g_h(x_h^k, a_h^k)$. The MDP evolves to x_{h+1}^k that is drawn from $\mathbb{P}_h(\cdot|x_h^k, a_h^k)$. The episode terminates at step $H + 1$. Without loss of generality, we assume that $r_{H+1} = g_{H+1} = 0$. In this paper, we consider the *challenging* scenario where the agent only observes the bandit information $r_h(x_h^k, a_h^k)$ and $g_h(x_h^k, a_h^k)$ at the visited state-action pair (x_h^k, a_h^k) . The policy-space of an agent is $\Delta(\mathcal{A}|\mathcal{S}, H); \{\{\pi_h(\cdot|\cdot)\}_{h=1}^H : \pi_h(\cdot|x) \in \Delta(\mathcal{A}), \forall x \in \mathcal{S}, \text{ and } h \in [H]\}$. Here $\Delta(\mathcal{A})$ is the probability simplex over the action space. For any $x_h^k \in \mathcal{S}$, $k \in [K]$, and $h \in [H]$, $\pi_{h,k}(a_h^k|x_h^k)$ denotes the probability that the action $a_h^k \in \mathcal{A}$ is taken at episode k when the state is x_h^k .

Let $V_{r,h}^\pi(x)$ denote the expected value of the total reward function starting from step h and state x when the agent selects action using the policy $\pi = \{\pi_h\}_{h=1}^H$

$$V_{r,h}^\pi(x) = \mathbb{E}_\pi \left[\sum_{i=h}^H r_i(x_i, a_i) | x_h = x \right], \quad (1)$$

where \mathbb{E} is taken with respect to the policy π and the transition probability kernel \mathbb{P} . Let $Q_{r,h}^\pi(x, a)$ denote the expected value of the total reward starting from step h and the state-action pair (x, a) and follows the policy π as

$$Q_{r,h}^\pi(x, a) = \mathbb{E}_\pi \left[\sum_{i=h}^H r_i(x_i, a_i) | x_h = x, a_h = a \right]. \quad (2)$$

Similarly, we define the value function for the utility $V_{g,h}^\pi(x)$, and the action-value function for the utility $Q_{g,h}^\pi(x, a)$. We denote $V_{j,h}^\pi(x)$, and $Q_{j,h}^\pi(x, a)$ for $j = r, g$.

Definition 1. For brevity, we denote $\mathbb{P}_h V_{j,h+1}^\pi(x, a) = \mathbb{E}_{x' \sim \mathbb{P}_h(\cdot|x, a)} V_{j,h+1}^\pi(x')$ for $j = r, g$.

Using this notation, the Bellman's equation associated with the policy π becomes

$$Q_{j,h}^\pi(x, a) = (r_h + \mathbb{P}_h V_{j,h+1}^\pi)(x, a) \quad (3)$$

Note that $V_{j,h}^\pi(x) = \langle \pi_h(\cdot|x), Q_{j,h}^\pi(x, \cdot) \rangle_{\mathcal{A}}$, where $\langle \pi_h(\cdot|x), Q_{j,h}^\pi(x, \cdot) \rangle_{\mathcal{A}} = \sum_{a \in \mathcal{A}} \pi_h(a|x) Q_{j,h}^\pi(x, a)$.

The objective of the learning agent is to find an optimal solution of the following problem

$$\text{maximize}_{\pi \in \Delta(\mathcal{A}|\mathcal{S}, \mathcal{H})} V_{r,1}^\pi(x_1), \quad \text{subject to } V_{g,1}^\pi(x_1) \geq b. \quad (4)$$

Note that even though we have only once constraint, it can be readily generalized to the scenario with multiple constraints. In order to avoid trivial solutions, we consider $b \in (0, H]$. We denote the optimal policy as π^* which solves the above optimization problem. Since π^* is obtained by having complete information, it is also denoted as *the best policy in the hindsight*.

Without any constraint information a priori, an agent can not know the policies that satisfy the constraint. Instead, we allow the policy to violate the constraint and minimize the regret while minimizing the total constraint violations over the K episodes. Such an approach is also considered in the existing literature [1, 7, 9]. We now define the performance metric which we seek to minimize.

Performance Metric. Let the policy employed by the agent at episode k be $\pi_k = [\pi_{1,k}, \dots, \pi_{h,k}, \dots, \pi_{H,k}]^T$. The performance metric we are considering is the following

$$\begin{aligned} \text{Regret}(K) &= \sum_{k=1}^K V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi_k}(x_1) \\ \text{Violation}(K) &= \left[\sum_{k=1}^K (b - V_{g,1}^{\pi_k}(x_1)) \right]_+, \end{aligned} \quad (5)$$

where $[z]_+ = \max\{z, 0\}$. The regret is defined as the difference between the total reward value by following the optimal policy π^* , and the total reward value obtained by following agent's policy π_k at episode k over K episodes. The constraint violation is defined as the difference between the threshold value Kb and the total utility function attained by following the policies over all the episodes K .

Linear Function Approximation. To handle a possible large number of states, we consider the following linear MDPs.

Assumption 1. The CMDP is a linear MDP with feature map $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$, if for any h , there exists d unknown signed measures $\mu_h = \{\mu_h^1, \dots, \mu_h^d\}$ over \mathcal{S} such that for any $(x, a, x') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$,

$$\mathbb{P}_h(x'|x, a) = \langle \phi(x, a), \mu_h(x') \rangle \quad (6)$$

and there exists vectors $\theta_{r,h}, \theta_{g,h} \in \mathbb{R}^d$ such that for any $(x, a) \in \mathcal{S} \times \mathcal{A}$,

$$r_h(x, a) = \langle \phi(x, a), \theta_{r,h} \rangle \quad g_h(x, a) = \langle \phi(x, a), \theta_{g,h} \rangle$$

Assumption 1 adapts the definition of linear MDP [14, 16] to the constrained case. By the above definition, the transition model, the reward, and the utility functions are linear in terms of feature map ϕ . We remark that despite being linear, $\mathbb{P}_h(\cdot|x, a)$ can still have infinite degrees of freedom since $\mu_h(\cdot)$ is unknown. Note that tabular MDP is part of linear MDP [14].

Note that [7, 17] studied another related concept known as linear kernel MDP. In the linear kernel MDP, the transition probability is given by $\mathbb{P}_h(x'|x, a) = \langle \psi(x', x, a), \theta_h \rangle$. In general, linear MDP and linear kernel MDPs are two different classes of MDP [17].

Similar to Proposition 1 in [14], we can show that for a linear MDP and for any policy π there exists $\{w_{j,h}^\pi\}_{h=1}^H$ such that $Q_{j,h}^\pi(x, a) = \langle w_{j,h}^\pi, \phi(x, a) \rangle$ for any $(x, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$. We, thus, focus on linear action-value function.

Dual problem and Slater's Condition. We first introduce few notations which we will use throughout this paper.

Definition 2. $V_h^{\pi,Y}(\cdot) = V_{h,r}^{\pi}(\cdot) + YV_{h,g}^{\pi}(\cdot)$, and $Q_h^{\pi,Y}(x, a) = Q_{r,h}^{\pi}(x, a) + YQ_{g,h}^{\pi}(x, a)$, where Y is the dual variable.

Thus, $V_h^{\pi,Y}(\cdot)$ and $Q_h^{\pi,Y}(\cdot, \cdot)$ are respectively the composite value function and Q -functions respectively. We can cast the problem (4) as a saddle point problem $\max_{\pi} \min_Y \mathcal{L}(\pi, Y)$ where $\mathcal{L}(\pi, Y) = V_{r,1}^{\pi}(x_1) + Y(V_{g,1}^{\pi}(x_1) - b) = V_1^{\pi,Y} - Yb$, where π is the primal policy and Y is the dual variable. However, the lagrangian is non-concave in π [18] even though it is convex in Y . Nevertheless, the strong duality holds [19]. Hence, there exists optimal dual variable Y^* , such that $\max_{\pi} \mathcal{L}(\pi, Y^*)$ will correspond to the optimal reward value function.

We assume the following slater's condition in this paper.

Assumption 2 (Slater's Condition). *There exists $\gamma > 0$, and $\bar{\pi} \in \Delta(\mathcal{A}|\mathcal{D}, \mathcal{H})$, such that $V_{g,1}^{\bar{\pi}}(x_1) \geq b + \gamma$,*

Lemma 1 (Boundedness of Y^*). *The optimal dual-variable $Y^* \leq \frac{V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\bar{\pi}}(x_1)}{\gamma} \leq \frac{H}{\gamma}$.*

The slater's condition is mild in practice and commonly adopted in previous works [7, 1, 20]. We use the properties of the slater's condition to bound the performance of our proposed algorithm.

Definition 3. *We set $\xi = 2H/\gamma$.*

3 Our Approach

Algorithm 1 Model Free Primal-Dual Algorithm for Linear Function Approximation

- 1: **Initialization:** $Y_1 = 0, w_{j,h} = 0, \xi = 2H/\gamma, \alpha = \frac{\log(|\mathcal{A}|)K}{2(1 + \xi + H)}, \eta = \xi/\sqrt{KH^2}, \beta = C_1 dH \sqrt{\log(4 \log |\mathcal{A}| dT/p)}$
 - 2: **for** episodes $k = 1, \dots, K$ **do**
 - 3: Receive the initial state x_1^k .
 - 4: **for** step $h = H, H-1, \dots, 1$ **do**
 - 5: $\Lambda_h^k \leftarrow \sum_{\tau=1}^{k-1} \phi(x_h^{\tau}, a_h^{\tau}) \phi(x_h^{\tau}, a_h^{\tau})^T + \lambda \mathbf{I}$
 - 6: $w_{r,h}^k \leftarrow (\Lambda_h^k)^{-1} [\sum_{\tau=1}^{k-1} \phi(x_h^{\tau}, a_h^{\tau}) [r_h(x_h^{\tau}, a_h^{\tau}) + V_{r,h+1}^k(x_h^{\tau})]]$
 - 7: $w_{g,h}^k \leftarrow (\Lambda_h^k)^{-1} [\sum_{\tau=1}^{k-1} \phi(x_h^{\tau}, a_h^{\tau}) [g_h(x_h^{\tau}, a_h^{\tau}) + V_{g,h+1}^k(x_h^{\tau})]]$
 - 8: $Q_{r,h}^k(\cdot, \cdot) \leftarrow \min\{\langle w_{r,h}^k, \phi(\cdot, \cdot) \rangle + \beta(\phi(\cdot, \cdot)^T (\Lambda_h^k)^{-1} \phi(\cdot, \cdot))^{1/2}, H\}$
 - 9: $Q_{g,h}^k(\cdot, \cdot) \leftarrow \min\{\langle w_{g,h}^k, \phi(\cdot, \cdot) \rangle + \beta(\phi(\cdot, \cdot)^T (\Lambda_h^k)^{-1} \phi(\cdot, \cdot))^{1/2}, H\}$
 - 10: $\pi_{h,k}(a|\cdot) = \frac{\exp(\alpha(Q_{r,h}^k(\cdot, a) + Y_k Q_{g,h}^k(\cdot, a)))}{\sum_a \exp(\alpha(Q_{r,h}^k(\cdot, a) + Y_k Q_{g,h}^k(\cdot, a)))}$
 - 11: $V_{r,h}^k(\cdot) = \sum_a \pi_{h,k}(a|\cdot) Q_{r,h}^k(\cdot, a)$
 - 12: $V_{g,h}^k(\cdot) = \sum_a \pi_{h,k}(a|\cdot) Q_{g,h}^k(\cdot, a)$
 - 13: **for** step $h = 1, \dots, H$ **do**
 - 14: Compute $Q_{r,h}^k(x_h^k, a), Q_{g,h}^k(x_h^k, a), \pi(a|x_h^k)$ for all a .
 - 15: Take action $a_h^k \sim \pi_{h,k}(\cdot|x_h^k)$ and observe x_{h+1}^k .
 - 16: $Y_{k+1} = \max\{\min\{Y_k + \eta(b - V_{g,1}^k(x_1)), \xi\}, 0\}$
-

We now describe our proposed algorithm in Algorithm 1. This algorithm is based on the primal-dual adaptation of the LSVI-UCB [14]. For a given dual variable, the primal policy is updated, and then the dual value is updated based on the estimated utility value function. At each episode, the algorithm consists of three parts. The first part (Steps 4-12) consists of updating the parameters $w_{r,h}^k, w_{g,h}^k$ and Λ_h^k which are used to update the $Q_{j,h}^k$ and $V_{j,h}^k$ at episode k . Λ_h^k is the Gram-matrix for the regularized least square problem (see Eqn. (8), later). Note that the Steps 8-12 are not evaluated for each state, rather, they are evaluated only for the encountered states till episode $k-1$. Hence, we do not need to iterate over potentially infinite number of states. For the first episode, since $k-1=0$ and $\tau=1$, we have $w_{j,h}^k=0, \forall j$ and $\Lambda_h^k = \lambda \mathbf{I}$. We note that $Q_{j,H+1}^k(\cdot, \cdot) = 0$ for $j=r, g$.

The value functions are updated (Steps 11-12) based on Q function and the policy. The policy is based (Step 10) on a soft-max policy unlike the greedy one in the unconstrained case [14]. Soft-max policy $\text{SOFT-MAX}_\alpha(\mathbf{X}) = \{\text{SOFT-MAX}_\alpha^i(\mathbf{X})\}_{i=1}^{|\mathcal{A}|}$ for any vector $\mathbf{X} \in \mathbb{R}^{|\mathcal{A}|}$ is a $|\mathcal{A}|$ -dimensional vector with parameter α where the i -th component

$$\text{SOFT-MAX}_\alpha^i(\mathbf{X}) = \frac{\exp(\alpha X_i)}{\sum_{n=1}^{|\mathcal{A}|} \exp(\alpha X_n)} \quad (7)$$

At step h , $\pi_{h,k}(a|x)$ is computed based on the soft-max policy on the composite Q -function vector $\{Q_{r,h}^k(x,a) + Y_k Q_{g,h}^k(x,a)\}_{a \in \mathcal{A}}$ where Y_k is the lagrangian multiplier. When $\alpha = \infty$, this becomes equal to the greedy policy. The second part (Steps 13-15) is the execution of the soft-max policy based on the composite Q -value for the encountered state x_h^k .

Q function and Value function Estimation. We need to estimate the value-function and Q -function with respect to the policy π_k . However, there are challenges. We do not know \mathbb{P}_h in Bellman's equation (3), rather $\mathbb{P}_h V_{j,h+1}^{\pi_k}$ should be replaced by the empirical samples. Further, in the large state space, we can not iterate over all (x,a) . Rather, we parameterize $Q_{j,h}^{\pi_k}(\cdot, \cdot)$ by a linear form $\langle w_{j,h}^k(\cdot, \cdot), \phi(\cdot, \cdot) \rangle$. The intuition is to obtain $w_{j,h}^k$ from the Bellman's equation using the regularized least-square regression. We obtain $w_{j,h}^k$ for $j = r, g$ according to the following equation

$$w_{j,h}^k \leftarrow \arg \min_{w \in \mathbb{R}^d} \sum_{\tau=1}^{k-1} [j_h(x_h^\tau, a_h^\tau) + V_{j,h+1}^k(x_{h+1}^\tau) - w^T \phi(x_h^\tau, a_h^\tau)]^2 + \lambda \|w\|_2^2 \quad (8)$$

Then, an additional bonus term $\beta(\phi(\cdot, \cdot)^T (\Lambda_h^k)^{-1} \phi(\cdot, \cdot))^{1/2}$ is added as in [14]. β is constant which we will characterize in the next section. Such an additional term is used for upper confidence bound in LSVI-UCB [14]. The same additional term is used for both $Q_{r,h}^k$ and $Q_{g,h}^k$. Note the difference with the LSVI-UCB, here, we need to estimate the value function corresponding to the soft-max policy π_k where in LSVI-UCB, a greedy policy corresponding to the Q -function is used.

Policy. We update a soft-max policy which selects actions according to the estimated 'composite' Q -function at the k -th episode. The reason behind using the soft-max policy instead of a greedy policy will be apparent in the next section when we state the main results and the proof ideas. Note from the strong duality, for optimal dual variable Y^* , optimal primal policy π^* maximizes the composite value function V^{π^*, Y^*} (Definition 2). Thus, the optimal policy should be a greedy one based on this optimal dual value Y^* . However, the greedy policy is not Lipschitz, hence, it does not provide uniform concentration bound for each individual value function, an essential step in the regret bound (Section 4.2). Hence, compared to the unconstrained scenario, we need more exploration in the policy space where the apparent reason is that we do not know the optimal dual variable beforehand. Since we use the soft-max policy, there is a gap compared to the optimal value even when the lagrangian multiplier Y_k becomes equal to Y^* . However, if α in the soft-max policy also scales with K , then we can bound the gap from the optimal value function (Section 4.2).

Dual Update. To infer the constraint violation, we estimate $V_{g,1}^k$ for $V_{g,1}^{\pi_k}$. We update the lagrangian multiplier Y_k by moving towards minimizing the lagrangian $\mathcal{L}(\pi, Y)$ over $Y \geq 0$ in line 16, where $\eta > 0$ is a step-size and ξ is the upper bound on the dual variable such that optimal dual variable Y^* is contained within $[0, \xi]$. The dual update is similar to the step described in [7, 1].

The dual update works as a trade-off between the reward maximization and the constraint violation reduction. If $b - V_{g,1}^k \geq 0$, that means with a high probability, the constraint will be violated for the policy π_k . Hence, the dual value is increased in order to focus on minimizing the constraint violation. Otherwise, the agent tries to maximize the reward value function.

Space and Time Complexities. We remark that Algorithm 1 only needs to store Y_k , $w_{r,h}^k, w_{g,h}^k, r_h(x_h^k, a_h^k), g_h(x_h^k, a_h^k), \Lambda_h^k$, and $\{\phi(x_h^k, a)\}_{a \in \mathcal{A}}$ for all $(h, k) \in [H] \times [K]$, hence, it takes $\mathcal{O}(d^2 H + dAT)$ space. When we compute $(\Lambda_h^k)^{-1}$ using Sherman-Morrison formula, the computation of $V_{j,h+1}^k$ is dominated by computing $Q_{j,h+1}^k$ and the policy π_k . Hence, it takes $\mathcal{O}(d^2 AT)$ time. Note that since our approach is model-free and we do not need to evaluate integrals as in [7] in order to estimate $\mathbb{P}_h V_{h+1}^k$.

Note that both η and α use the knowledge of K . In case, K is unknown, one can use the "doubling trick" [21] which will only scale the regret and constraint violation by a constant factor.

4 Analysis

We now state the main result. We prove that Algorithm 1 achieves regret and constraint violation which are sublinear in $T = KH$ where T is the total number of steps.

4.1 Main Results

Theorem 1. Fix $p > 0$. If we set $\lambda = 1$, $\beta = C_1 dH\sqrt{\iota}$ in Algorithm 1 where $\iota = \log(\log(|\mathcal{A}|)4dT/p)$ for some absolute constant C_1 . With probability $(1 - p)$,

$$\begin{aligned} \text{Regret}(K) &\leq C\sqrt{d^3 H^3 T \iota^2} + \xi\sqrt{HT} \\ \text{Violation}(K) &\leq \frac{C'2(1 + \xi)}{\xi}\sqrt{d^3 H^3 T \iota^2} \end{aligned}$$

for some absolute constants C , and C' .

We remark the difference with the existing results. Since $\xi = 2H/\gamma$ (by Definition 3), our result indicates that our approach obtains $\tilde{O}(\sqrt{d^3 H^3 T})$ regret and the same order of constraint violation where \tilde{O} absorbs logarithmic factor on T . The regret and constraint violation are sub-linear in T , and similar dependence is observed in [1, 7]. Also note that compared to the unconstrained case [14], there is an additional $\log(|\mathcal{A}|)$ factor in the value of ι which arises because we use soft-max policy instead of the greedy policy which adds to the covering number. The regret and constraint violation do not depend on the dimension of the state space, rather, it depends on the dimension of the feature space. *To the best of our knowledge this is the first result which shows both $\tilde{O}(\sqrt{T})$ regret and constraint violation in the model-free set up (tabular or linear) without requiring a simulator.*

Comparison with [7]: Compared to [7] which also considers linear function approximation (however, it considers linear kernel MDP rather linear MDP) we improve the result in [7] by a factor of H . Second, compared to [7], which is a model-based policy-based algorithm, ours is a model-free value-based algorithm. Due to this, the above uniform concentration challenge does not exist in [7]. Moreover, our model-free algorithm also enjoys an easy implementation and improved computation efficiency since it does not estimate the next step expected value function as in [7] which requires an integration oracle to compute a d -dimensional integration at every step. [7] also needs to store the previous policies and estimated value functions, hence, it needs $O(T)$ additional space complexity. We have an additional \sqrt{d} factor in front of the regret and constraint violation. Similar difference in regret is also observed between the model-based linear kernel unconstrained MDP [22] and model-free linear unconstrained MDP [14] *even in the unconstrained case.*

Similar to the discussion in Section 3.1 on [23], our result directly translates to a sample complexity guarantee (or, PAC guarantee). For example, we can learn a policy π such that $V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi}(x_1) \leq \epsilon$, and $b - V_{g,1}^{\pi}(x_1) \leq \epsilon$ after $\tilde{O}(d^3 H^4 / \epsilon^2)$ number of samples. Here, the policy π is obtained after running Algorithm 1 for $\tilde{O}(d^3 H^3 / \epsilon^2)$ number of episodes, and then selecting policy π_k with probability $1/K$ for any $k \in [K]$.

Recently, [24] proposed an algorithm with provable sample complexity guarantee for linear CMDP. However, the regret and violation guarantees are different from the sample complexity guarantees as the former ones are *any time* guarantee. The proposed algorithms are different since the goal is different. In particular, the uniform concentration bound challenge does not appear there. Note that using the explore-then-commit algorithm [23], one can achieve $\tilde{O}(T^{2/3})$ regret for large T (from $\tilde{O}(1/\epsilon^2)$ sample complexity bound achieved in [24]) which is worse than ours. Additionally, we achieve zero violation (Remark 2) while maintaining the same order of regret with respect to T .

4.2 Outline of the Proof

In this section, we provide an outline of our proof, which is mainly divided into three steps. We first establish a decomposition of the sum of reward regret and constraint violation. Then, we will bound two key terms that are related to optimism and prediction error, respectively. Finally, using standard optimization tools, we can achieve the main results. We highlight that the key challenges lie in the second step where a balance between the optimistic term and prediction error term is handled via the introduced soft-max policy.

Step 1: Bounding the sum of Regret and violation scaled by dual variable Similar to [1], we first establish the following decomposition, which upper bounds the sum of regret and violation. This will serve as the basis when applying optimization tools in Step 3.

Lemma 2 (Decomposition). *For any $Y \in [0, \xi]$, we have*

$$\begin{aligned} & \sum_{k=1}^K (V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi^k}(x_1)) + Y \sum_{k=1}^K (b - V_{g,1}^{\pi^k}(x_1)) \leq \frac{1}{2\eta} Y^2 + \frac{\eta}{2} H^2 K + \\ & \underbrace{\sum_{k=1}^K \left(V_{r,1}^{\pi^*}(x_1) + Y_k V_{g,1}^{\pi^*}(x_1) \right) - \left(V_{r,1}^k(x_1) + Y_k V_{g,1}^k(x_1) \right)}_{\mathcal{T}_1} + \\ & \underbrace{\sum_{k=1}^K \left(V_{r,1}^k(x_1) - V_{r,1}^{\pi^k}(x_1) \right) + Y \sum_{k=1}^K \left(V_{g,1}^k(x_1) - V_{g,1}^{\pi^k}(x_1) \right)}_{\mathcal{T}_2} \end{aligned}$$

Note that \mathcal{T}_1 is similar to the term related to optimism in the unconstrained case with the difference being that we now have two value functions weighted by the dual variable Y_k . Similarly, \mathcal{T}_2 is similar to prediction error term with the additional weight by Y . Since the first term in the above inequality can be easily bounded with a proper choice of η , we are only left to bound \mathcal{T}_1 and \mathcal{T}_2 , respectively.

Step 2: Bounding \mathcal{T}_1 and \mathcal{T}_2 To bound \mathcal{T}_2 and \mathcal{T}_1 , we need to bound the difference between the *individual* estimated value function $V_{j,h}^k$ and the *individual* value function $V_{j,h}^\pi$ corresponding to a given policy π at episode k . As in the unconstrained case, the key step is to control the fluctuations in least-squares value iteration. In particular, we need to show that for all $(k, h) \in [K] \times [H]$ with high probability

$$\left\| \sum_{\tau=1}^{k-1} \phi(x_h^\tau, a_h^\tau) \left[V_{j,h+1}^k(x_{h+1}^\tau) - \mathbb{P}_h V_{j,h+1}^k(x_h^\tau, a_h^\tau) \right] \right\|_{(\Lambda_h^k)^{-1}}$$

is upper bounded by lower order term (e.g., $\mathcal{O}(d\sqrt{\log K})$). To this end, value-aware uniform concentration is required to handle the dependence between $V_{j,h+1}^k$ and samples $\{x_{h+1}^\tau\}_{\tau=1}^{k-1}$, which renders the standard self-normalized inequality infeasible in the model-free setting. The general idea here is to fix a function class $\mathcal{V}_{j,h}$ in advance and then show that each possible value function in our algorithm $V_{j,h}^k$ is within this class which has polynomial log-covering number. *In the following, we fix an $h \in [H]$ and drop the subscript h for notation simplicity.*

Uniform Concentration Bound for class of value function: We first note that this uniform concentration bound is the main motivation for us to choose a soft-max policy as we will see that the standard greedy policy would fail in this case. That is, in order to guarantee that for each possible V_j^k , there is an ϵ -close function in \mathcal{V}_j , it would basically lead to a very large covering number. To address this, we introduce soft-max policy and define the following corresponding function classes. We first define the following class for Q -function for $j = r, g$. $\mathcal{Q}_j = \{Q_j | Q_j(\cdot, \cdot) = \min\{\langle w_j, \phi(\cdot, \cdot) \rangle + \beta \sqrt{\phi(\cdot, \cdot)^T (\Lambda_h)^{-1} \phi(\cdot, \cdot)}, H\}$. Then, we define the following value function class \mathcal{V}_j . $\mathcal{V}_j = \{V_j | V_j(\cdot) = \sum_a \pi(a|\cdot) Q_j(\cdot, a); Q_j \in \mathcal{Q}_j, \pi \in \Pi\}$, where Π is given by the following class $\Pi = \{\pi | \pi(a|\cdot) = \text{SOFT-MAX}_\alpha^\alpha((Q_r(\cdot, \cdot) + Y Q_g(\cdot, \cdot))); \forall a \in \mathcal{A}, Q_r \in \mathcal{Q}_r, Q_g \in \mathcal{Q}_g, Y \in [0, \xi]\}$, where SOFT-MAX is defined in (7).

Why soft-max? At this moment, we can explain why the introduction of soft-max in our algorithm is critical. Suppose we follow the standard greedy selection, which corresponds to $\alpha = \infty$ in above. The key issue in this approach is that one needs a large ϵ -covering for \mathcal{V}_j so that each possible V_j^k can be well-approximated (i.e., ϵ -close) by function in \mathcal{V}_j . This is in sharp contrast to the unconstrained case where \mathcal{V}_j has a polynomial log-covering number. To see this difference, in the unconstrained case, we only have V_r^k that is greedy with respect to Q_r^k . By the fact that \max_a is a contraction map, an ϵ -covering of the Q_r^k implies an ϵ -covering of V_r^k and meanwhile the covering number of \mathcal{Q}_r is reasonably small (Lemma 14) by standard arguments. This no longer holds in the constrained case due to the use of a composite Q -function. In particular, note that if the policy is greedy w.r.t. the composite Q -function, then an ϵ -covering of Q_j^k fails to be an

ϵ -covering of V_j^k in general since the greedy policy is not smooth in that a slight change of the composite Q -function could lead to a substantial change of the output action. This leads to a large distance for *individual* value functions due to the different action choices, even though the Q -function is close (Please see Appendix G for an example). Hence, one can not approximate individual value function within ϵ -bound using greedy policy based on composite Q -function. This fact motivates us to turn to SOFT-MAX_α , which is Lipschitz continuous with a Lipschitz constant at most 2α . Thus, our main idea is as follows. Given Q_r^k, Q_g^k and Y_k , we can first find fixed $\tilde{Q}_r \in \mathcal{Q}_r$, $\tilde{Q}_g \in \mathcal{Q}_g$ and $\tilde{Y} \in [0, \xi]$ such that $\|Q_r^k - \tilde{Q}_r\|_\infty \leq \epsilon_1, \|Q_g^k - \tilde{Q}_g\|_\infty \leq \epsilon_2, \|Y^k - \tilde{Y}\|_\infty \leq \epsilon_3$ and $\|(Q_r^k + Y_k Q_g^k) - (\tilde{Q}_r + \tilde{Y} \tilde{Q}_g)\|_\infty \leq \epsilon$ with a reasonably small covering number. Then, thanks to the smoothness of soft-max function, we have $\|\pi_k - \tilde{\pi}\|_1 \leq 2\alpha\epsilon$ (Lemma 15). Combining this with the closeness of individual Q -function yields the closeness of individual V -function (Lemma 13). Hence, it ensures that the class \mathcal{V}_j in our set-up has log-covering number of $\mathcal{O}(d \log(K))$.

Choosing hyper-parameter α to achieve bound: A larger value of α means that we need a smaller ϵ , hence a larger covering number. Then, one may wonder if we can choose an arbitrarily small value for α . However, the term \mathcal{T}_1 will be enlarged if we choose too small α . Note that in the unconstrained case, \mathcal{T}_1 is upper bounded by zero due to optimism under greedy policy. Now, since we are using soft-max, we need to bound the approximation error between the soft-max and greedy one. As expected, in this case, a larger α leads to a smaller approximation error.

From the discussions above, we can see that the approximation-smoothness trade-off of the soft-max function is well captured by our \mathcal{T}_1 and \mathcal{T}_2 , respectively. Therefore, we need to carefully choose the value of α to balance these two. In particular, with $\alpha = \frac{\log(|\mathcal{A}|)K}{2H}$, we have the following bounds on \mathcal{T}_1 and \mathcal{T}_2 , respectively.

Lemma 3. *With probability $1 - p/2$, we have $\mathcal{T}_1 \leq K \frac{H \log(|\mathcal{A}|)}{\alpha}$. Hence, for $\alpha = \frac{\log(|\mathcal{A}|)K}{2(1 + \xi + H)}$, we have $\mathcal{T}_1 \leq 2H(1 + \xi + H)$ with probability $1 - p/2$.*

Lemma 4. *With probability at least $1 - p/2$, $\mathcal{T}_2 \leq \mathcal{O}((Y + 1)\sqrt{d^3 H^3 T \iota^2})$, where $\iota = \log[\log(|\mathcal{A}|)4dT/p]$*

Remark 1. *The additional $\log|\mathcal{A}|$ factor in ι arises as a trade-off for selecting soft-max policy as it is evident in Lemma 3. When we compensate by making α scaled with $\log|\mathcal{A}|$ in Lemma 3, it increases the covering-number by $\log|\mathcal{A}|$ as well in the ι term.*

Step 3: Final Result by combining all the pieces: By replacing $Y = 0, \eta = \xi/(\sqrt{KH^2})$, and combining Lemma 2,3, and 4, we obtain the regret bound. We also obtain the constraint violation using the idea from [1].

Remark 2. *We can reduce the violation to zero while maintaining the same order on regret with respect to T (Appendix H). We consider a tighter optimization problem where we add ζ in the constraint of (4). In Appendix H, we bound the difference between the optimal value function for the tighter and the original problem as a function of ζ . Since the tighter problem is also CMDP, we attain the regret and violation bound as in Theorem 1 with $b + \zeta$ in place of b . Hence, by choosing ζ , we can show that it is possible to achieve $\tilde{\mathcal{O}}(\sqrt{T})$ regret and zero violation for large enough K (Theorem 3 in Appendix H) albeit with an extra H factor in front of regret bound.*

5 Experiments

We evaluate Algorithm 1 on a simulated model for job scheduling to validate our theoretical results. We consider that the number of jobs belongs to the discrete state $\{0, 1, \dots, 9\}$ where 0 means that there is no job. Total time horizon (H) is divided in 10 steps. After H steps, a new episode begins. At the start of the each episode, the state of the job is assumed to be 9, i.e., the job stack is full. The agent needs to decide whether to send job ($a = 1$) or not ($a = 0$) to a machine. We assume that if $a = 1$, the agent sends 2 jobs to a machine and incurs a cost as the machine spends some resources to process the job. We assume that the rewards are step dependent. In particular, we assume that at time steps from 3 to 6, the reward is $1 - 0.9a$, In other time steps, the reward is $1 - 0.2a$. This mimics the setup where at certain time, it might be more costly to process a job (for example, electricity cost might be higher, or the machine needs to abandon an important job).




Figure 1: The plot for cumulative regret and constraint violation as a function of K . Each plot is an average of 10 trials.

We assume that even if the scheduler schedules a job, the machine might not be able to complete the 2 jobs. In particular,

$$x_{h+1} = \begin{cases} \max\{x_h - 2a, 0\} & \text{w.p. } 0.8 \\ \max\{x_h - a, 0\} & \text{w.p. } 0.1 \\ x_h, & \text{otherwise} \end{cases}$$

Thus, if $a = 0$, the state $x_{h+1} = x_h$. The agent gets an utility of $g(x_h, a_h, x_{h+1}) = (x_h - x_{h+1})/2$. We want that utility to be greater than or equal to 4 at the end of every episode. This will ensure that at most 1 job can remain at the end of each episode.

We run Algorithm 1 for 2×10^5 episodes (K). The parameters we used are the followings: $\alpha = K/(1 + 2H/\gamma + H)$, $\eta = 2H/(\gamma\sqrt{KH^2})$. We set $\gamma = 1$. We have also set $\epsilon = 0.1$ in order to ensure that the violation goes towards 0 as K increases. Note that the setup can be represented in a tabular form. Since linear MDP contains tabular form, the feature space representation becomes simple, in particular, $\phi(s, a) = \mathbf{e}_{s,a}$ where $\mathbf{e}_{s,a}$ is 1 for the state-action pair (s, a) , and 0 otherwise. The dimension of the feature space is $|\mathcal{S}||\mathcal{A}|$. The cumulative regret and constraint violations are shown in Figure 1. As predicted by our theory, the regret scales only as \sqrt{K} . On the other hand, the cumulative violation oscillates. However, the violation approaches 0 as K increases. The oscillation of violation is due to the fact that dual variable also oscillates in order to illicit conservative response when the cumulative violation increases, and illicit aggressive response when the violation decreases.

6 Conclusion and Future Work

We propose a model-free RL-based algorithm for linear MDP. We have achieved $\tilde{\mathcal{O}}(\sqrt{d^3 H^3 T})$ regret and $\tilde{\mathcal{O}}(\sqrt{d^3 H^3 T})$ constraint violation. To the best of our knowledge, this is the first result which shows $\tilde{\mathcal{O}}(\sqrt{T})$ regret and $\tilde{\mathcal{O}}(\sqrt{T})$ constraint violation without requiring a generator for the model-free case. We have extended the LSVI-UCB algorithm in the primal-dual type framework. We have underlined the technical challenges in doing so and explained how the greedy policy fails to achieve an uniform concentration bound for individual value function. Subsequently, we show that a soft-max type algorithm achieves that.

Compared to [7], our regret is off by \sqrt{d} factor which is due to the fact that we need to use uniform concentration bound. The similar gap is also observed for the unconstrained set-up as well. Whether we can tighten this dependence on d remains an important future research direction. Whether we can tighten the dependence on H also constitutes a future research direction. Extending the work to the setup where the feature-space needs to be learnt is also important. Recent works [25–27] on feature-space learning for unconstrained MDPs may provide some insights in this direction. Consideration of non-linear MDP also constitutes a future research direction.

Acknowledgment

XZ is partially supported by the NSF grant CNS-2153220. This work has been supported in part by NSF grants: 2112471 (also partly funded by DHS), CNS-2106933, and CNS-1901057 and a grant from the Army Research Office: W911NF-21-1-0244.

References

- [1] Yonathan Efroni, Shie Mannor, and Matteo Pirotta. Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189*, 2020.
- [2] Rahul Singh, Abhishek Gupta, and Ness B Shroff. Learning in markov decision processes under constraints. *arXiv preprint arXiv:2002.12435*, 2020.
- [3] 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.
- [4] Liyuan Zheng and Lillian Ratliff. Constrained upper confidence reinforcement learning. In *Learning for Dynamics and Control*, pages 620–629. PMLR, 2020.
- [5] Krishna C Kalagarla, Rahul Jain, and Pierluigi Nuzzo. A sample-efficient algorithm for episodic finite-horizon mdp with constraints. *arXiv preprint arXiv:2009.11348*, 2020.
- [6] Tao Liu, Ruida Zhou, Dileep Kalathil, PR Kumar, and Chao Tian. Learning policies with zero or bounded constraint violation for constrained mdps. *arXiv preprint arXiv:2106.02684*, 2021.
- [7] 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.
- [8] Tengyu Xu, Yingbin Liang, and Guanghui Lan. Crpo: A new approach for safe reinforcement learning with convergence guarantee. In *International Conference on Machine Learning*, pages 11480–11491. PMLR, 2021.
- [9] Dongsheng Ding, Kaiqing Zhang, Tamer Basar, and Mihailo R Jovanovic. Natural policy gradient primal-dual method for constrained markov decision processes. In *NeurIPS*, 2020.
- [10] 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.
- [11] Sven Koenig and Reid G Simmons. Complexity analysis of real-time reinforcement learning. In *AAAI*, pages 99–107, 1993.
- [12] Mohammad Gheshlaghi Azar, Rémi Munos, and Bert Kappen. On the sample complexity of reinforcement learning with a generative model. *arXiv preprint arXiv:1206.6461*, 2012.
- [13] Honghao Wei, Xin Liu, and Lei Ying. A provably-efficient model-free algorithm for constrained markov decision processes. *arXiv preprint arXiv:2106.01577*, 2021.
- [14] 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.
- [15] Sanae Amani, Christos Thrampoulidis, and Lin F Yang. Safe reinforcement learning with linear function approximation. *arXiv preprint arXiv:2106.06239*, 2021.
- [16] 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.
- [17] Dongruo Zhou, Jiafan He, and Quanquan Gu. Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*, pages 12793–12802. PMLR, 2021.
- [18] Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98):1–76, 2021.
- [19] Santiago Paternain, Miguel Calvo-Fullana, Luiz FO Chamon, and Alejandro Ribeiro. Safe policies for reinforcement learning via primal-dual methods. *arXiv preprint arXiv:1911.09101*, 2019.

- [20] Shuang Qiu, Xiaohan Wei, Zhuoran Yang, Jieping Ye, and Zhaoran Wang. Upper confidence primal-dual optimization: Stochastically constrained markov decision processes with adversarial losses and unknown transitions. *arXiv preprint arXiv:2003.00660*, 2020.
- [21] Lilian Besson and Emilie Kaufmann. What doubling tricks can and can't do for multi-armed bandits. *arXiv preprint arXiv:1803.06971*, 2018.
- [22] Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pages 463–474. PMLR, 2020.
- [23] Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? *arXiv preprint arXiv:1807.03765*, 2018.
- [24] 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.
- [25] Aditya Modi, Jinglin Chen, Akshay Krishnamurthy, Nan Jiang, and Alekh Agarwal. Model-free representation learning and exploration in low-rank mdps. *arXiv preprint arXiv:2102.07035*, 2021.
- [26] Xuezhou Zhang, Yuda Song, Masatoshi Uehara, Mengdi Wang, Alekh Agarwal, and Wen Sun. Efficient reinforcement learning in block mdps: A model-free representation learning approach. In *International Conference on Machine Learning*, pages 26517–26547. PMLR, 2022.
- [27] Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity and representation learning of low rank mdps. *Advances in neural information processing systems*, 33:20095–20107, 2020.
- [28] Ling Pan, Qingpeng Cai, Qi Meng, Wei Chen, Longbo Huang, and Tie-Yan Liu. Reinforcement learning with dynamic boltzmann softmax updates. *arXiv preprint arXiv:1903.05926*, 2019.
- [29] Alessandro Epasto, Mohammad Mahdian, Vahab Mirrokni, and Manolis Zampetakis. Optimal approximation–smoothness tradeoffs for soft-max functions. *arXiv preprint arXiv:2010.11450*, 2020.
- [30] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.
- [31] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- [32] Sharan Vaswani, Lin F Yang, and Csaba Szepesvári. Near-optimal sample complexity bounds for constrained mdps. *arXiv preprint arXiv:2206.06270*, 2022.
- [33] Bruce Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advances in Applied probability*, 14(3):502–525, 1982.
- [34] Qiaomin Xie, Yudong Chen, Zhaoran Wang, and Zhuoran Yang. Learning zero-sum simultaneous-move markov games using function approximation and correlated equilibrium. *arXiv preprint arXiv:2002.07066*, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] We have clearly compared our results with the existing work. In the Conclusion and Future Work, we identified some important directions where our approach can be extended.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] We clearly state the Assumptions 1 and 2 in Section 2
 - (b) Did you include complete proofs of all theoretical results? [Yes] We have added the proofs in Appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Our main goal is to obtain Theoretical Results. Nevertheless, we have performed experiment on a simulated setup to verify our theoretical results.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We have characterized the parameters.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Organization of Appendix: In Section A, we state some results which we use throughout. In Section B, we prove Lemma 2. In Section C, we prove Lemmas 3 and 4. In Section C.1, we state and prove base results Lemmas 8, 9, and 10 which are necessary to prove the Lemmas 3 and 4. Subsequently, we prove Lemmas 3 and 4 in sequence in Sections C.2 and C.3. In Section D, we prove Theorem 1. In Section E we prove Lemma 8 and highlight that how soft-max policy enables us to obtain a low-covering number for our class of value functions. In Section F, we state some results proved in the existing literature which we have used in proving our results. In Section G, we explain why greedy policy based on the composite Q -function fails using an example. In Section H, we describe the results on how zero violation can be attained.

Notations: Throughout the rest of this paper, we denote $Q_{r,h}^k, Q_{g,h}^k, w_{r,h}^k, w_{g,h}^k, \Lambda_h^k$ as the Q -value and the parameter values estimated at the episode k . $V_{j,h}^k(\cdot) = \langle \pi_{h,k}(\cdot|\cdot), Q_{j,h}^k(\cdot, \cdot) \rangle_{\mathcal{A}}$. $\pi_{h,k}(\cdot|x)$ is the soft-max policy based on the composite Q -function at the k -th episode as $Q_{r,h}^k + Y_k Q_{g,h}^k$. To simplify the presentation, we denote $\phi_h^k = \phi(x_h^k, a_h^k)$.

Without loss of generality, we assume $\|\phi(x, a)\|_2 \leq 1$ for all $(x, a) \in \mathcal{S} \times \mathcal{A}$, $\|\mu_h(\mathcal{S})\|_2 \leq \sqrt{d}$, $\|\theta_{j,h}\|_2 \leq \sqrt{d}$ for $j = r, g$ and all $h \in [H]$.

A Preliminary Results

Lemma 5. *Under Assumption 1, for any fixed policy π , let w_h^π be the corresponding weights such that $Q_{j,h}^\pi = \langle \phi(x, a), w_{j,h}^\pi \rangle$, for $j \in \{r, g\}$, then we have for all $h \in [H]$,*

$$\|w_{j,h}^\pi\| \leq 2H\sqrt{d} \quad (9)$$

Proof. From the linearity of the action-value function, we have

$$\begin{aligned} Q_{j,h}^\pi(x, a) &= j_h(x, a) + \mathbb{P}_h V_{j,h}^\pi(x, a) \\ &= \langle \phi(x, a), \theta_{j,h} \rangle + \int_{\mathcal{S}} V_{j,h+1}^\pi(x') \langle \phi(x, a), d\mu_h(x') \rangle \\ &= \langle \phi(x, a), w_{j,h}^\pi \rangle \end{aligned} \quad (10)$$

where $w_{j,h}^\pi = \theta_{j,h} + \int_{\mathcal{S}} V_{j,h+1}^\pi(x') d\mu_h(x')$.

Now, $\|\theta_{j,h}\| \leq \sqrt{d}$, and $\|\int_{\mathcal{S}} V_{j,h+1}^\pi(x') d\mu_h(x')\| \leq H\sqrt{d}$. Thus, the result follows. \square

Lemma 6. *For any (k, h) , the weight $w_{j,h}^k$ satisfies*

$$\|w_{j,h}^k\| \leq 2H\sqrt{dk/\lambda} \quad (11)$$

Proof. For any vector $v \in \mathcal{R}^d$ we have

$$|v^T w_{j,h}^k| = |v^T (\Lambda_h^k)^{-1} \sum_{\tau=1}^{k-1} \phi_h^\tau(x_h^\tau, a_h^\tau) (j_h(x_h^\tau, a_h^\tau) + \sum_a \pi_{h+1,k}(a|x_{h+1}^\tau) Q_{j,h+1}^k(x_{h+1}^\tau, a))| \quad (12)$$

here $\pi_{h,k}(\cdot|x)$ is the Soft-max policy.

Note that $Q_{j,h+1}^k(x, a) \leq H$ for any (x, a) . Hence, from (12) we have

$$\begin{aligned} |v^T w_{j,h}^k| &\leq \sum_{\tau=1}^{k-1} |v^T (\Lambda_h^k)^{-1} \phi_h^\tau| \cdot 2H \\ &\leq \sqrt{\sum_{\tau=1}^{k-1} v^T (\Lambda_h^k)^{-1} v} \sqrt{\sum_{\tau=1}^{k-1} \phi_h^\tau (\Lambda_h^k)^{-1} \phi_h^\tau} \cdot 2H \\ &\leq 2H \|v\| \frac{\sqrt{dk}}{\sqrt{\lambda}} \end{aligned} \quad (13)$$

Note that $\|w_{j,h}^k\| = \max_{v: \|v\|=1} |v^T w_{j,h}^k|$. Hence, the result follows. \square

B Proof of Lemma 2

We first state and prove the following result which is similar to the one proved in [7].

Lemma 7. For $Y \in [0, \xi]$,

$$\sum_{k=1}^K (Y - Y_k)(b - V_{g,1}^k(x_1)) \leq \frac{Y^2}{2\eta} + \frac{\eta H^2 K}{2} \quad (14)$$

Proof.

$$\begin{aligned} |Y_{k+1} - Y|^2 &= |\text{Proj}_{[0,\xi]}(Y_k + \eta(b - V_{g,1}^k(x_1))) - \text{Proj}_{[0,\xi]}(Y)|^2 \\ &\leq (Y_k + \eta(b - V_{g,1}^k(x_1)) - Y)^2 \\ &\leq (Y_k - Y)^2 + \eta^2 H^2 + 2\eta Y_k(b - V_{g,1}^k(x_1)) \end{aligned} \quad (15)$$

Summing over k , we obtain

$$\begin{aligned} 0 \leq |Y_{K+1} - Y|^2 &\leq |Y_1 - Y|^2 + 2\eta \sum_{k=1}^K (b - V_{g,1}^k(x_1))(Y_k - Y) + \eta^2 H^2 K \\ \sum_{k=1}^K (Y - Y_k)(b - V_{g,1}^k(x_1)) &\leq \frac{|Y_1 - Y|^2}{2\eta} + \frac{\eta H^2 K}{2} \end{aligned} \quad (16)$$

Since $Y_1 = 0$, we have the result. \square

Now, we prove Lemma 2.

Proof. Note that

$$\begin{aligned} Y \sum_{k=1}^K (b - V_{g,1}^{\pi_k}(x_1)) &= \sum_k (Y - Y_k)(b - V_{g,1}^k(x_1)) + Y_k(b - V_{g,1}^k) + Y(V_{g,1}^k(x_1) - V_{g,1}^{\pi_k}(x_1)) \\ &\leq \frac{1}{2\eta} Y^2 + \frac{\eta}{2} H^2 K + \sum_{k=1}^K (Y_k b - Y_k V_{g,1}^k(x_1)) + Y(V_{g,1}^k(x_1) - V_{g,1}^{\pi_k}(x_1)) \\ &\leq \frac{1}{2\eta} Y^2 + \frac{\eta}{2} H^2 K + \sum_{k=1}^K (Y_k V_{g,1}^{\pi_k^*}(x_1) - Y_k V_{g,1}^k(x_1)) + \sum_{k=1}^K Y(V_{g,1}^k(x_1) - Y_k V_{g,1}^{\pi_k}(x_1)) \end{aligned}$$

where the first inequality follows from Lemma 7, and the second inequality follows from the fact that $V_{g,1}^{\pi_k^*}(x_1) \geq b$. Hence, the result simply follows from the above inequality. \square

C Proof of Lemmas 3 and 4

First, we prove some base results which we use to prove both Lemmas 3 and 4 in Section C.1. Subsequently, we prove Lemma 3 in Section C.2 and Lemma 4 in Section C.3.

C.1 Proof of Base Results

We state and prove Lemmas 8,9, and 10.

First, we state the concentration lemma which is essential in controlling the fluctuations in the least square value iteration.

Lemma 8. There exists a constant C_2 such that for any fixed $p \in (0, 1)$, if we let \mathcal{E} be the event that

$$\left\| \sum_{\tau=1}^{k-1} \phi_{j,h}^\tau [V_{j,h+1}^k(x_{h+1}^\tau) - \mathbb{P}_h V_{j,h+1}^k(x_h^\tau, a_h^\tau)] \right\|_{(\Lambda_h^k)^{-1}} \leq C_2 dH \sqrt{\chi} \quad (17)$$

for all $j \in \{r, g\}$, $\chi = \log[4(C_1 + 1) \log(|\mathcal{A}|) dT/p]$, for some constant C_1 , then $\Pr(\mathcal{E}) = 1 - p/2$.

The proof of Lemma 8 is technical and relegated to Appendix E. An extra $\log(|\mathcal{A}|)$ term appears because α appears in the covering number.

We now, recursively bound the difference between the value function maintained in Algorithm 1 (without the bonus term) and the value function for any policy for both the reward and utility value functions. We bound this using the expected difference at the next step plus an error term. This error term can be upper bounded by the bonus term with a high-probability.

Lemma 9. *There exists an absolute constant $\beta = C_1 dH\sqrt{\iota}$, $\iota = \log(\log(|\mathcal{A}|)4dT/p)$, and for any fixed policy π , on the event \mathcal{E} defined in Lemma 8, we have*

$$\langle \phi(x, a), w_{j,h}^k \rangle - Q_{j,h}^\pi(x, a) = \mathbb{P}_h(V_{j,h+1}^k - V_{j,h+1}^\pi)(x, a) + \Delta_h^k(x, a) \quad (18)$$

for some $\Delta_h^k(x, a)$ that satisfies $|\Delta_h^k(x, a)| \leq \beta \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)}$.

Proof. We only prove for $j = r$, the proof for $j = g$ is similar.

Note that $Q_{r,h}^\pi(x, a) = \langle \phi(x, a), w_{r,h}^\pi \rangle = r_h(x, a) + \mathbb{P}_h V_{r,h+1}^\pi(x, a)$.

Hence, we have

$$\begin{aligned} w_{r,h}^k - w_{r,h}^\pi &= (\Lambda_h^k)^{-1} \sum_{\tau=1}^{k-1} \phi_h^\tau [r_h^\tau + V_{r,h+1}^k(x_{h+1}^\tau)] - w_{r,h}^\pi \\ &= -\lambda (\Lambda_h^k)^{-1} (w_{r,h}^\pi) + (\Lambda_h^k)^{-1} \sum_{\tau=1}^{k-1} \phi_h^\tau [V_{r,h+1}^k(x_{h+1}^\tau) - \mathbb{P}_h V_{r,h+1}^k(x_h^\tau, a_h^\tau)] \\ &\quad + (\Lambda_h^k)^{-1} \sum_{\tau=1}^{k-1} \phi_h^\tau [\mathbb{P}_h V_{r,h+1}^k(x_h^\tau, a_h^\tau) - \mathbb{P}_h V_{r,h+1}^\pi(x_h^\tau, a_h^\tau)] \end{aligned} \quad (19)$$

Now, we bound each term in the right hand side of expression in (19). We call those terms as \mathbf{q}_1 , \mathbf{q}_2 , and \mathbf{q}_3 respectively.

First, note that

$$\begin{aligned} |\langle \phi(x, a), \mathbf{q}_1 \rangle| &= |\lambda \langle \phi(x, a), (\Lambda_h^k)^{-1} (w_{r,h}^\pi) \rangle| \\ &\leq \sqrt{\lambda} \|w_{r,h}^\pi\| \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)} \end{aligned} \quad (20)$$

Second, from Lemma 8, for the event in \mathcal{E} , we have

$$|\langle \phi(x, a), \mathbf{q}_2 \rangle| \leq CdH\sqrt{\chi} \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)} \quad (21)$$

where $\chi = \log(4(C_1 + 1) \log(|\mathcal{A}|)dT/p)$. Third,

$$\begin{aligned} \langle \phi(x, a), \mathbf{q}_3 \rangle &= \langle \phi(x, a), (\Lambda_h^k)^{-1} \sum_{\tau=1}^{k-1} \phi_h^\tau [\mathbb{P}_h (V_{r,h+1}^k - V_{r,h+1}^\pi)(x_h^\tau, a_h^\tau)] \rangle \\ &= \langle \phi(x, a), (\Lambda_h^k)^{-1} \sum_{\tau=1}^{k-1} \phi_h^\tau (\phi_h^\tau)^T \int (V_{r,h+1}^k - V_{r,h+1}^\pi)(x') d\mu_h(x') \rangle \\ &= \langle \phi(x, a), \int (V_{r,h+1}^k - V_{r,h+1}^\pi)(x') d\mu_h(x') \rangle - \langle \phi(x, a), \lambda (\Lambda_h^k)^{-1} \int (V_{r,h+1}^k - V_{r,h+1}^\pi)(x') d\mu_h(x') \rangle \end{aligned} \quad (22)$$

The last term in (22) can be bounded as the following

$$|\langle \phi(x, a), \lambda (\Lambda_h^k)^{-1} \int (V_{r,h+1}^k - V_{r,h+1}^\pi)(x') d\mu_h(x') \rangle| \leq 2H\sqrt{d\lambda} \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)} \quad (23)$$

since $\|\int (V_{r,h+1}^k - V_{r,h+1}^\pi)(x') d\mu_h(x')\|_2 \leq 2H\sqrt{d}$ as $\|\mu_h(\mathcal{S})\| \leq \sqrt{d}$. The first term in (22) is equal to

$$\mathbb{P}_h (V_{r,h+1}^k - V_{r,h+1}^\pi)(x, a) \quad (24)$$

Note that $\langle \phi(x, a), w_{r,h}^k \rangle - Q_{r,h}^\pi(x, a) = \langle \phi(x, a), w_{r,h}^k - w_{r,h}^\pi \rangle = \langle \phi(x, a), \mathbf{q}_1 + \mathbf{q}_2 + \mathbf{q}_3 \rangle$. Since $\lambda = 1$, we have from (20), (21),(23), and (24)

$$|\langle \phi(x, a), w_{r,h}^k \rangle - Q_{r,h}^\pi(x, a) - \mathbb{P}_h(V_{r,h+1}^k - V_{r,h+1}^\pi)(x, a)| \leq C_3 dH \sqrt{\chi} \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)} \quad (25)$$

for some constant C_3 which is independent of C_1 . Finally, note that

$$\begin{aligned} C_3 \sqrt{\chi} &= \sqrt{\log(4(C_1 + 1) \log(|\mathcal{A}|) dT/p)} \\ &= C_3 \sqrt{\iota + \log(C_1 + 1)} \\ &\leq C_1 \sqrt{\iota} \end{aligned} \quad (26)$$

where $\iota = \log(4 \log(|\mathcal{A}|) dT/p)$. The last inequality follows from the fact that $\iota \in [\log 4, \infty)$ as $|A| \geq 2$, and C_3 is independent of C_1 . Hence, we can always pick $C_3 \sqrt{\log 4 + \log(C_1 + 1)} \leq C_1 \sqrt{\log 4}$ which satisfies (26) for all values of $\iota \in [\log 4, \infty)$. \square

Next, using the above lemma, we bound the difference between the composite value function maintained by the algorithm and the composite value function for a policy with the Lagrangian Y_k .

Lemma 10. *With prob. $1 - p$, (for the event in \mathcal{E})*

$$Q_{r,h}^\pi(x, a) + Y_k Q_{g,h}^\pi(x, a) \leq Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a) - \mathbb{P}_h(V_{h+1}^k - V_{h+1}^{\pi, Y_k})(x, a) \quad (27)$$

Proof. From Lemma 9 and the fact that $|Q_{r,h}^\pi| \leq H$, we have w.p. $1 - p/2$,

$$\begin{aligned} Q_{r,h}^\pi(x, a) &\leq \min\{\langle \phi(x, a), w_{r,h}^k \rangle + \beta \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)}, H\} \\ &\quad + \mathbb{P}_h(V_{r,h+1}^\pi - V_{r,h+1}^k)(x, a) \\ &= Q_{r,h}^k(x, a) + \mathbb{P}_h(V_{r,h+1}^\pi - V_{r,h+1}^k)(x, a) \end{aligned}$$

where the last equality follows from the definition of $Q_{r,h}^k$.

Similarly, with probability $1 - p/2$,

$$Y_k Q_{g,h}^\pi(x, a) \leq Y_k Q_{g,h}^k(x, a) + Y_k \mathbb{P}_h(V_{g,h+1}^\pi - V_{g,h+1}^k)(x, a)$$

Hence, from union bound, with probability $1 - p$,

$$Q_{r,h}^\pi(x, a) + Y_k Q_{g,h}^\pi(x, a) \leq Q_{r,h}^k + Y_k Q_{g,h}^k(x, a) + \mathbb{P}_h(V_{h+1}^{\pi, Y_k} - V_{h+1}^k)(x, a) \quad \square$$

C.2 Proof of Lemma 3

First, we state and prove a supporting result which bounds the value functions corresponding to the greedy policy and the soft-max policy at a given step. We show that this gap can be controlled by the parameter α .

Lemma 11. *Then, $\bar{V}_h^k(x) - V_h^k(x) \leq \frac{\log|\mathcal{A}|}{\alpha}$*

where

Definition 4. $\bar{V}_h^k(\cdot) = \max_a [Q_{r,h}^k(\cdot, a) + Y_k Q_{g,h}^k(\cdot, a)]$.

$\bar{V}_h^k(\cdot)$ is the value function corresponds to the greedy-policy with respect to the composite Q -function.

Proof. Note that

$$V_h^k(x) = \sum_a \pi_{h,k}(a|x) [Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a)] \quad (28)$$

where

$$\pi_{h,k}(a|x) = \frac{\exp(\alpha[Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a)])}{\sum_a \exp(\alpha[Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a)])} \quad (29)$$

Denote $a_x = \arg \max_a [Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a)]$

Now, recall from Definition 4 that $\bar{V}_h^k(x) = [Q_{r,h}^k(x, a_x) + Y_k Q_{g,h}^k(x, a_x)]$. Then,

$$\begin{aligned} \bar{V}_h^k(x) - V_h^k(x) &= [Q_{r,h}^k(x, a_x) + Y_k Q_{g,h}^k(x, a_x)] \\ &\quad - \sum_a \pi_{h,k}(a|x) [Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a)] \\ &\leq \left(\frac{\log(\sum_a \exp(\alpha(Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a))))}{\alpha} \right) \\ &\quad - \sum_a \pi_{h,k}(a|x) [Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a)] \\ &\leq \frac{\log(|\mathcal{A}|)}{\alpha} \end{aligned} \quad (30)$$

where the last inequality follows from Proposition 1 in [28]. \square

We are now ready to show Lemma 3.

Proof. We prove the lemma by Induction.

First, we prove for the step H .

Note that $Q_{j,H+1}^k = 0 = Q_{j,H+1}^\pi$.

Under the event in \mathcal{E} as described in Lemma 8 and from Lemma 9, we have for $j = r, g$,

$$|\langle \phi(x, a), w_{j,H}^k(x, a) \rangle - Q_{j,H}^\pi(x, a)| \leq \beta \sqrt{\phi(x, a)^T (\Lambda_H^k)^{-1} \phi(x, a)}$$

Hence, for any (x, a) ,

$$\begin{aligned} Q_{j,H}^\pi(x, a) &\leq \min\{\langle \phi(x, a), w_{j,H}^k \rangle + \beta \sqrt{\phi(x, a)^T (\Lambda_H^k)^{-1} \phi(x, a)}, H\} \\ &= Q_{j,H}^k(x, a) \end{aligned} \quad (31)$$

Hence, from the definition of \bar{V}_h^k ,

$$\begin{aligned} \bar{V}_H^k(x) &= \max_a [Q_{r,H}^k(x, a) + Y_k Q_{g,H}^k(x, a)] \geq \sum_a \pi(a|x) [Q_{r,H}^\pi(x, a) + Y_k Q_{g,H}^\pi(x, a)] \\ &= V_H^{\pi, Y_k}(x) \end{aligned} \quad (32)$$

for any policy π . Thus, it also holds for π^* , the optimal policy. Hence, from Lemma 11, we have

$$V_H^{\pi^*, Y_k}(x) - V_H^k(x) \leq \frac{\log(|\mathcal{A}|)}{\alpha}$$

Now, suppose that it is true till the step $h + 1$ and consider the step h .

Since, it is true till step $h + 1$, thus, for any policy π ,

$$\mathbb{P}_h(V_{h+1}^{\pi, Y_k} - V_{h+1}^k)(x, a) \leq \frac{(H - h) \log(|\mathcal{A}|)}{\alpha} \quad (33)$$

From (27) in Lemma 10 and the above result, we have for any (x, a)

$$Q_{r,h}^\pi(x, a) + Y_k Q_{g,h}^\pi(x, a) \leq Q_{r,h}^k(x, a) + Y_k Q_{g,h}^k(x, a) + \frac{(H - h) \log(|\mathcal{A}|)}{\alpha} \quad (34)$$

Hence,

$$V_h^{\pi, Y_k}(x) \leq \bar{V}_h^k(x) + \frac{(H-h)\log(|\mathcal{A}|)}{\alpha} \quad (35)$$

Now, again from Lemma 11, we have $\bar{V}_h^k(x) - V_h^k(x) \leq \frac{\log(|\mathcal{A}|)}{\alpha}$. Thus,

$$V_h^{\pi, Y_k}(x) - V_h^k(x) \leq \frac{(H-h+1)\log(|\mathcal{A}|)}{\alpha} \quad (36)$$

Now, since it is true for any policy π , it will be true for π^* . From the definition of V^{π, Y_k} , we have

$$\left(V_{r,h}^{\pi^*}(x) + Y_k V_{g,h}^{\pi^*}(x) \right) - \left(V_{r,h}^k(x) + Y_k V_{g,h}^k(x) \right) \leq \frac{(H-h+1)\log(|\mathcal{A}|)}{\alpha} \quad (37)$$

Hence, the result follows by summing over K and considering $h = 1$. \square

C.3 Proof of Lemma 4

In order to prove the Lemma 4, we state and prove the following result.

First, we introduce a notation. Let

$$\begin{aligned} D_{j,h,1}^k &= \langle (Q_{j,h}^k(x_h^k, \cdot) - Q_{j,h}^{\pi_k}(x_h^k, \cdot)), \pi_{h,k}(\cdot|x_h^k) \rangle - (Q_{j,h}^k(x_h^k, a_h^k) - Q_{j,h}^{\pi_k}(x_h^k, a_h^k)) \\ D_{j,h,2}^k &= \mathbb{P}_h(V_{j,h+1}^k - V_{j,h+1}^{\pi_k})(x_h^k, a_h^k) - [V_{j,h+1}^k - V_{j,h+1}^{\pi_k}](x_{h+1}^k) \end{aligned} \quad (38)$$

Lemma 12. *On the event defined in \mathcal{E} in Lemma 8, we have*

$$V_{j,1}^k(x_1) - V_{j,1}^{\pi_k}(x_1) \leq \sum_{h=1}^H (D_{j,h,1}^k + D_{j,h,2}^k) + \sum_{h=1}^H 2\beta \sqrt{\phi(x_h^k, a_h^k)^T (\Lambda_h^k)^{-1} \phi(x_h^k, a_h^k)} \quad (39)$$

Proof. By Lemma 9, for any x, h, a, k

$$\begin{aligned} &\langle w_{j,h}^k(x, a), \phi(x, a) \rangle + \beta \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)} - Q_{j,h}^{\pi_k} \\ &\leq \mathbb{P}_h(V_{j,h+1}^k - V_{j,h+1}^{\pi_k})(x, a) + 2\beta \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)} \end{aligned} \quad (40)$$

Thus,

$$\begin{aligned} Q_{j,h}^k(x, a) - Q_{j,h}^{\pi_k}(x, a) &\leq \mathbb{P}_h(V_{j,h+1}^k - V_{j,h+1}^{\pi_k})(x, a) + 2\beta \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)} \\ \mathbb{P}_h(V_{j,h+1}^k - V_{j,h+1}^{\pi_k})(x, a) + 2\beta \sqrt{\phi(x, a)^T (\Lambda_h^k)^{-1} \phi(x, a)} &- (Q_{j,h}^k(x, a) - Q_{j,h}^{\pi_k}(x, a)) \geq 0 \end{aligned} \quad (41)$$

Since $V_{j,h}^k(x) = \sum_a \pi_{h,k}(a|x) Q_{j,h}^k(x, a)$ and $V_{j,h}^{\pi_k}(x) = \sum_a \pi_{h,k}(a|x) Q_{j,h}^{\pi_k}(x, a)$ where $\pi_{h,k}(a|\cdot) = \text{SOFT-MAX}_\alpha^a(Q_{r,h}^k + Y_k Q_{g,h}^k) \forall a$.

Thus, from (41),

$$\begin{aligned} V_{j,h}^k(x_h^k) - V_{j,h}^{\pi_k}(x_h^k) &= \sum_a \pi_{h,k}(a|x_h^k) [Q_{j,h}^k(x_h^k, a) - Q_{j,h}^{\pi_k}(x_h^k, a)] \\ &\leq \sum_a \pi_{h,k}(a|x_h^k) [Q_{j,h}^k(x_h^k, a) - Q_{j,h}^{\pi_k}(x_h^k, a)] \\ &+ 2\beta \sqrt{\phi(x_h^k, a_h^k)^T (\Lambda_h^k)^{-1} \phi(x_h^k, a_h^k)} + \mathbb{P}_h(V_{j,h+1}^k - V_{j,h+1}^{\pi_k})(x_h^k, a_h^k) - (Q_{j,h}^k(x_h^k, a_h^k) - Q_{j,h}^{\pi_k}(x_h^k, a_h^k)) \end{aligned} \quad (42)$$

Thus, from (42), we have

$$V_{j,h}^k(x_h^k) - V_{j,h}^{\pi_k}(x_h^k) \leq D_{j,h,1}^k + D_{j,h,2}^k + [V_{j,h+1}^k - V_{j,h+1}^{\pi_k}](x_{h+1}^k) + 2\beta \sqrt{\phi(x_h^k, a_h^k)^T (\Lambda_h^k)^{-1} \phi(x_h^k, a_h^k)} \quad (43)$$

Hence, by iterating recursively, we have

$$V_{j,1}^k(x_1) - V_{j,1}^{\pi_k}(x_1) \leq \sum_{h=1}^H (D_{j,h,1}^k + D_{j,h,2}^k) + \sum_{h=1}^H 2\beta \sqrt{\phi(x_h^k, a_h^k)^T (\Lambda_h^k)^{-1} \phi(x_h^k, a_h^k)} \quad (44)$$

The result follows. \square

We, are now ready to prove Lemma 4.

Proof. Note from Lemma 12, we have

$$\sum_{k=1}^K V_{j,1}^k(x_1) - V_{j,1}^{\pi_k}(x_1) \leq \sum_{k=1}^K \sum_{h=1}^H (D_{j,h,1}^k + D_{j,h,2}^k) + \sum_{k=1}^K \sum_{h=1}^H 2\beta \sqrt{\phi(x_h^k, a_h^k)^T (\Lambda_h^k)^{-1} \phi(x_h^k, a_h^k)} \quad (45)$$

We, now, bound the individual terms. First, we show that the first term corresponds to a Martingale difference.

For any $(k, h) \in [K] \times [H]$, we define $\mathcal{F}_{h,1}^k$ as σ -algebra generated by the state-action sequences, reward, and constraint values, $\{(x_i^\tau, a_i^\tau)\}_{(\tau,i) \in [k-1] \times [H]} \cup \{(x_i^k, a_i^k)\}_{i \in [h]}$.

Similarly, we define the $\mathcal{F}_{h,2}^k$ as the σ -algebra generated by $\{(x_i^\tau, a_i^\tau)\}_{(\tau,i) \in [k-1] \times [H]} \cup \{(x_i^k, a_i^k)\}_{i \in [h]} \cup \{x_{H+1}^k\}$. x_{H+1}^k is a null state for any $k \in [K]$.

A filtration is a sequence of σ -algebras $\{\mathcal{F}_{h,m}^k\}_{(k,h,m) \in [K] \times [H] \times [2]}$ in terms of time index

$$t(k, h, m) = 2(k-1)H + 2(h-1) + m \quad (46)$$

which holds that $\mathcal{F}_{h,m}^k \subset \mathcal{F}_{h',m'}^{k'}$ for any $t \leq t'$.

Note from the definitions in (38) that $D_{j,h,1}^k \in \mathcal{F}_{h,1}^k$ and $D_{j,h,2}^k \in \mathcal{F}_{h,2}^k$. Thus, for any $(k, h) \in [K] \times [H]$,

$$\mathbb{E}[D_{j,h,1}^k | \mathcal{F}_{h-1,2}^k] = 0, \quad \mathbb{E}[D_{j,h,2}^k | \mathcal{F}_{h,1}^k] = 0 \quad (47)$$

Notice that $t(k, 0, 2) = t(k-1, H, 2) = 2(H-1)k$. Clearly, $\mathcal{F}_{0,2}^k = \mathcal{F}_{H,2}^{k-1}$ for any $k \geq 2$. Let $\mathcal{F}_{0,2}^1$ be empty. We define a Martingale sequence

$$\begin{aligned} M_{j,h,m}^k &= \sum_{\tau=1}^{k-1} \sum_{i=1}^H (D_{j,i,1}^\tau + D_{j,i,2}^\tau) + \sum_{i=1}^{h-1} (D_{j,i,1}^k + D_{j,i,2}^k) + \sum_{l=1}^m D_{j,h,l}^k \\ &= \sum_{(\tau,i,l) \in [K] \times [H] \times [2], t(\tau,i,l) \leq t(k,h,m)} D_{j,i,l}^\tau \end{aligned} \quad (48)$$

where $t(k, h, m) = 2(k-1)H + 2(h-1) + m$ is the time index. Clearly, this martingale is adopted to the filtration $\{\mathcal{F}_{h,m}^k\}_{(k,h,m) \in [K] \times [H] \times [2]}$, and particularly

$$\sum_{k=1}^K \sum_{h=1}^H (D_{j,h,1}^k + D_{j,h,2}^k) = M_{j,H,2}^K \quad (49)$$

Thus, $M_{j,H,2}^K$ is a Martingale difference satisfying $|M_{j,H,2}^K| \leq 4H$ since $|D_{j,h,1}^k|, |D_{j,h,2}^k| \leq 2H$ From the Azuma-Hoeffding inequality, we have

$$\Pr(M_{j,H,2}^K > s) \leq 2 \exp\left(-\frac{s^2}{16TH^2}\right) \quad (50)$$

With probability $1 - p/2$ at least for any $j = r, g$,

$$\sum_k \sum_h M_{j,H,2}^K \leq \sqrt{16TH^2 \log(4/p)} \quad (51)$$

Now, we bound the second term. Note that the minimum eigen value of Λ_h^k is at least $\lambda = 1$ for all $(k, h) \in [K] \times [H]$. By Lemma 17,

$$\sum_{k=1}^K (\phi_h^k)^T (\Lambda_h^k)^{-1} \phi_h^k \leq 2 \log \left[\frac{\det(\Lambda_h^{k+1})}{\det(\Lambda_h^1)} \right] \quad (52)$$

Moreover, note that $\|\Lambda_h^{k+1}\| = \|\sum_{\tau=1}^k \phi_h^\tau (\phi_h^\tau)^T + \lambda \mathbf{I}\| \leq \lambda + k$, hence,

$$\sum_{k=1}^K (\phi_h^k)^T (\Lambda_h^k)^{-1} \phi_h^k \leq 2d \log \left[\frac{\lambda + k}{\lambda} \right] \leq 2d\iota \quad (53)$$

Now, by Cauchy-Schwartz inequality, we have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \sqrt{(\phi_h^k)^T (\Lambda_h^k)^{-1} \phi_h^k} &\leq \sum_{h=1}^H \sqrt{K} \left[\sum_{k=1}^K (\phi_h^k)^T (\Lambda_h^k)^{-1} \phi_h^k \right]^{1/2} \\ &\leq H \sqrt{2dK\iota} \end{aligned} \quad (54)$$

Note that $\beta = C_1 dH \sqrt{\iota}$.

Thus, we have with probability $1 - p/2$,

$$\begin{aligned} \sum_{k=1}^K V_{r,1}^k(x_1) - V_{r,1}^{\pi^k}(x_1) + Y \sum_{k=1}^K (V_{g,1}^k(x_1) - V_{g,1}^{\pi^k}(x_1)) \\ \leq (Y + 1) [\sqrt{2TH^2 \log(4/p)} + C_4 \sqrt{d^3 H^3 T \iota^2}] \end{aligned} \quad (55)$$

Hence, the result follows. \square

D Proof of Theorem 1

We, first, show the regret bound. Note from Lemma 2, Lemma 3, and Lemma 4, we have for $Y \in [0, \xi]$ at least w.p. $1 - p$,

$$\begin{aligned} \sum_{k=1}^K (V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi^k}(x_1)) + Y \sum_{k=1}^K (b - V_{g,1}^{\pi^k}(x_1)) \\ \leq \frac{Y^2}{2\eta} + \frac{\eta}{2} H^2 K + \frac{HK \log |\mathcal{A}|}{\alpha} + \tilde{\mathcal{O}}((Y + 1) \sqrt{d^3 H^3 T \iota^2}) \end{aligned} \quad (56)$$

Replacing Y with 0 in (56), we have

$$\sum_{k=1}^K (V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi^k}(x_1)) \leq \frac{\eta}{2} H^2 K + \frac{HK \log |\mathcal{A}|}{\alpha} + \mathcal{O}(\sqrt{d^3 H^3 T \iota^2}) \quad (57)$$

By noting that $\eta = \frac{\xi}{\sqrt{KH^2}}$, and $\alpha = \frac{\log |\mathcal{A}| K}{2(1 + \xi + H)}$, we have the result.

We, now, show the violation bound. Note from Lemma 2, Lemma 3, and Lemma 4 that w.p. $1 - p$ (at least),

$$\begin{aligned} \sum_{k=1}^K (V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi^k}(x_1) + Y(b - V_{g,1}^{\pi^k}(x_1))) &\leq \frac{1}{2\eta} Y^2 + \frac{\eta}{2} H^2 K + \frac{HK \log |\mathcal{A}|}{\alpha} + \\ &(1 + Y) \mathcal{O}(\sqrt{H^3 T \iota^2}) \end{aligned} \quad (58)$$

Now, put $\eta = \frac{\xi}{\sqrt{KH^2}}$, $\alpha = \frac{\log |\mathcal{A}| K}{2(1 + \xi + H)}$, and $Y \leq \xi$, then, we have

$$\sum_{k=1}^K (V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi^k}(x_1) + Y(b - V_{g,1}^{\pi^k}(x_1))) \leq (1 + \xi) \mathcal{O}(\sqrt{H^3 T \iota^2}) + \xi \sqrt{KH^2} \quad (59)$$

Now, there exists a policy π' such that $V_{r,1}^{\pi'} = \frac{1}{K} \sum_{k=1}^K V_{r,1}^{\pi_k}$, $V_{g,1}^{\pi'} = \frac{1}{K} \sum_{k=1}^K V_{g,1}^{\pi_k}$. By the occupancy measure, $V_{r,1}^{\pi}$ and $V_{g,1}^{\pi}$ are linear in occupancy measure induced by π . Thus, the average of K occupancy measure also produces an occupancy measure which induces policy π' and $V_{r,1}^{\pi'}$, and $V_{g,1}^{\pi'}$. We take $Y = 0$ when $\sum_{k=1}^K (b - V_{g,1}^{\pi_k}(x_1^k)) < 0$, otherwise $Y = \xi$. Hence, we have

$$\begin{aligned} & (V_{r,1}^{\pi^*}(x_1) - \frac{1}{K} \sum_{k=1}^K V_{r,1}^{\pi_k}(x_1) + \xi(b - \frac{1}{K} \sum_{k=1}^K V_{g,1}^{\pi_k}(x_1)))_+ \\ &= (V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi'}(x_1) + \xi[b - V_{g,1}^{\pi'}(x_1)])_+ \\ &\leq \frac{(1 + \xi)\mathcal{O}(\sqrt{d^3 H^3 T \iota^2})}{K} + \xi \frac{\sqrt{KH^2}}{K} \end{aligned}$$

Since $\xi = 2H/\gamma$, and using the result of strong duality (Lemma 19), we have

$$(b - \frac{1}{K} \sum_{k=1}^K V_{g,1}^{\pi_k}(x_1^k))_+ \leq \frac{2(1 + \xi)}{K\xi} \mathcal{O}(\sqrt{d^3 H^3 T \iota^2}) \quad (60)$$

Hence, the result follows.

E Proof of Lemma 8

To simplify the notation, we remove h from the subscript from $Q_{j,h}^k$ and $V_{j,h}^k$ in this Section.

In order to prove the Lemma 8, we first compute the ϵ -covering number for the class of value functions (Lemma 13). In order to compute that we first compute the ϵ -covering number of the individual Q -functions (Lemma 14) which is essential to compute the covering number for composite Q -functions (Corollary 1). Subsequently, we show that if the two Q -functions and the Lagrange multipliers are close, the policies are also close (Lemma 15).

We first introduce the set of Q -functions.

Definition 5. Let $\mathcal{Q}_j = \{Q|Q(\cdot, \cdot) = \min\{w_j^T \phi(\cdot, \cdot) + \beta \sqrt{\phi^T(\cdot, \cdot)^T \Lambda^{-1} \phi(\cdot, \cdot)}, H\}\}$

The set \mathcal{Q} is parameterized by w_j , and Λ . We have $\|w_j\| \leq 2H\sqrt{dk/\lambda}$ (from Lemma 6). The minimum eigen value of Λ satisfies $\lambda_{min} \geq 1$. Hence, the Frobenius norm of Λ^{-1} is bounded. Note that $Q_j^k \in \mathcal{Q}_j$ for $j = r, g$.

We now introduce the class of value function for $j = r, g$.

Definition 6. Let $\mathcal{V}_j = \{V_j|V_j(\cdot) = \sum_a \pi(a|\cdot)Q_j(\cdot, a); Q_r \in \mathcal{Q}_r, Q_g \in \mathcal{Q}_g, Y \in [0, \xi]\}$ for $j = r, g$, where

$$\Pi = \{\pi|\forall a \in \mathcal{A}, \pi(a|\cdot) = \text{SOFT-MAX}_\alpha^a((Q_r(\cdot, \cdot) + YQ_g(\cdot, \cdot))Q_r \in \mathcal{Q}_r, Q_g \in \mathcal{Q}_g, Y \in [0, \xi]).\}$$

The class of value function \mathcal{V}_j is parameterized by w_r, w_g, Λ , and $Y \in [0, \xi]$. Note that even the individual value function depends on the Q -functions for both the reward and utility since the policy depends on the composite Q -function.

First, we need to see whether $V_j^k \in \mathcal{V}_j$. Recall the definition of V_j^k at the k -th episode $V_j^k(\cdot) = \sum_a \pi_k(a|\cdot)Q_j^k(\cdot, a)$ where

$$\pi_k(a|\cdot) = \text{SOFT-MAX}_\alpha^a((Q_r(\cdot, \cdot) + Y_k Q_g(\cdot, \cdot)).$$

Since $Q_j \in \mathcal{Q}_j$ for all j , and $0 \leq Y_k \leq \xi$, thus, $V_j \in \mathcal{V}_j$.

We now bound the ϵ -covering number for the class of value function

Lemma 13. There exists a $\tilde{V}_j \in \mathcal{V}_j$ parameterized by $(\tilde{w}_r, \tilde{w}_g, \tilde{\beta}, \Lambda, \tilde{Y})$ such that $\text{DIST}(V_j, \tilde{V}_j) \leq \epsilon$ where

$$\text{DIST}(V_j, \tilde{V}_j) = \sup_x |V_j(x) - \tilde{V}_j(x)|. \quad (61)$$

Let $N_\epsilon^{V_j}$ be the ϵ -covering number for the set \mathcal{V}_j , then,

$$\log N_\epsilon^{V_j} \leq d \log \left(1 + 8H \frac{\sqrt{dk}}{\sqrt{\lambda\epsilon'}} \right) + d^2 \log \left[1 + 8d^{1/2} \beta^2 / (\lambda(\epsilon')^2) \right] + \log \left(1 + \frac{\xi}{\epsilon'} \right) \quad (62)$$

$$\text{where } \epsilon' = \frac{\epsilon}{H2\alpha(1 + \xi + H) + 1}$$

Note that ϵ -covering number is dependent on ξ . This is because the policy depends on the Lagrange multiplier Y which is upper bounded by ξ . Thus, we also need ϵ -covering for the Lagrange multiplier in order to obtain ϵ -close value function. Note that the ϵ -covering does not depend on sample dependent terms. Rather it only depends on general $w_{j,h}$, Λ , and Y . Since the policy parameter is α , we also have ϵ -covering number is dependent on α .

In order to prove the above lemma, we first state and prove some additional results.

We, first, obtain the $N_\epsilon^{Q_j}$ covering number for the set \mathcal{Q}_j . Towards this end, we first, introduce some notations.

Definition 7. Let \mathcal{C}_w^ϵ be an $\epsilon/2$ -cover of the set $\{w \in \mathbb{R}^d \mid \|w\| \leq 2H\sqrt{dk/\lambda}\}$ with respect to the 2-norm. Let $\mathcal{C}_\mathbf{A}^\epsilon$ be an $\epsilon^2/4$ -cover of the set $\{\mathbf{A} \in \mathbb{R}^{d \times d} \mid \|\mathbf{A}\|_F \leq d^{1/2} \beta^2 \lambda^{-1}\}$ with respect to the Frobenius norm.

Lemma 14.

$$|\mathcal{C}_w^\epsilon| \leq (1 + 8H\sqrt{dk/\lambda}/\epsilon)^d, \quad |\mathcal{C}_\mathbf{A}^\epsilon| \leq [1 + 8d^{1/2} \beta^2 / (\lambda\epsilon^2)]^{d^2} \quad (63)$$

The ϵ -covering number for the set \mathcal{Q}_j , for $j = r, g$, $N_\epsilon^{Q_j}$ of the set \mathcal{Q}_j for $j = r, g$ satisfies the following

$$\log N_\epsilon^{Q_j} \leq d \log \left(1 + \frac{8H\sqrt{dk}}{\sqrt{\lambda\epsilon}} \right) + d^2 \log [1 + 8d^{1/2} \beta^2 / (\lambda\epsilon)^2] \quad (64)$$

The distance metric is the ∞ -norm, i.e., $\text{dist}(Q_1, Q_2) = \sup_{x,a} |Q_1(x, a) - Q_2(x, a)|$.

Proof. For notational simplicity, we represent $\mathbf{A} = \beta^2 \Lambda^{-1}$, and reparameterized the class \mathcal{Q}_j by (w_j, \mathbf{A}) . Now,

$$\begin{aligned} \text{dist}(Q_1, Q_2) &= \sup_{x,a} |[w_1^T \phi(x, a) + \sqrt{\phi^T(x, a) \mathbf{A}_1 \phi(x, a)}] - [w_2^T \phi(x, a) + \sqrt{\phi^T(x, a) \mathbf{A}_2 \phi(x, a)}]| \\ &\leq \sup_{\phi: \|\phi\| \leq 1} |[w_1^T \phi + \sqrt{\phi^T \mathbf{A}_1 \phi}] - [w_2^T \phi + \sqrt{\phi^T \mathbf{A}_2 \phi}]| \\ &\leq \sup_{\phi: \|\phi\| \leq 1} |(w_1 - w_2)^T \phi| + \sup_{\phi: \|\phi\| \leq 1} \sqrt{|\phi^T (\mathbf{A}_1 - \mathbf{A}_2) \phi|} \\ &= \|w_1 - w_2\| + \sqrt{\|\mathbf{A}_1 - \mathbf{A}_2\|} \leq \|w_1 - w_2\| + \sqrt{\|\mathbf{A}_1 - \mathbf{A}_2\|_F} \end{aligned} \quad (65)$$

where the second-last inequality follows from the fact that $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x - y|}$. For matrices $\|\cdot\|$, and $\|\cdot\|_F$ denote matrix operator norm and the Frobenius norm respectively.

Recall that \mathcal{C}_w is an $\epsilon/2$ -cover of the set $\{w \in \mathbb{R}^d \mid \|w\| \leq 2H\sqrt{dk/\lambda}\}$ with respect to the 2-norm. Also recall that $\mathcal{C}_\mathbf{A}$ be an $\epsilon^2/4$ -cover of the set $\{\mathbf{A} \in \mathbb{R}^{d \times d} \mid \|\mathbf{A}\|_F \leq d^{1/2} \beta^2 \lambda^{-1}\}$. Thus, from Lemma 18,

$$|\mathcal{C}_w^\epsilon| \leq (1 + 8H\sqrt{dk/\lambda}/\epsilon)^d, \quad |\mathcal{C}_\mathbf{A}^\epsilon| \leq [1 + 8d^{1/2} \beta^2 / (\lambda\epsilon^2)]^{d^2}$$

For any $Q_j \in \mathcal{Q}_j$, there exists a \tilde{Q}_j parameterized by (w_2, \mathbf{A}_2) where $w_2 \in \mathcal{C}_w^\epsilon$ and $\mathbf{A}_2 \in \mathcal{C}_\mathbf{A}^\epsilon$ such that $\text{dist}(Q_j, \tilde{Q}_j) \leq \epsilon$. Hence, $N_\epsilon^{Q_j} \leq |\mathcal{C}_w^\epsilon| |\mathcal{C}_\mathbf{A}^\epsilon|$, which gives the result since $\log(\cdot)$ is an increasing function. \square

Since the class of Q -function is independent of the policy we do not have ξ and α in the ϵ -covering number.

From the above lemma and since $Y_k \leq \xi$, we have the following,

Corollary 1. *If $\text{dist}(Q_r^k, \tilde{Q}_r) \leq \epsilon'$, $\text{dist}(Q_g^k, \tilde{Q}_g) \leq \epsilon'$, and $|\tilde{Y}_k - Y_k| \leq \epsilon'$, then, $\text{dist}(Q_r^k + Y_k Q_g^k, \tilde{Q}_r + \tilde{Y}_k \tilde{Q}_g) \leq \epsilon'(1 + \xi + H)$.*

Proof. Note that $\tilde{Q}_j \in \mathcal{Q}_j$ belongs to the ϵ' covering of the set \mathcal{Q} .

$$\begin{aligned}
\text{dist}(Q_r^k + Y_k Q_g^k, \tilde{Q}_r + \tilde{Y}_k \tilde{Q}_g) &= \sup_{x,a} |(Q_r^k(x,a) + Y_k Q_g^k(x,a)) - (\tilde{Q}_r(x,a) + \tilde{Y}_k \tilde{Q}_g(x,a))| \\
&\leq \sup_{x,a} |(Q_r^k(x,a) + Y_k Q_g^k(x,a)) - (\tilde{Q}_r(x,a) + Y_k \tilde{Q}_g(x,a))| + \sup_{x,a} |(\tilde{Y}_k - Y_k) \tilde{Q}_g(x,a)| \\
&\leq \sup_{x,a} |Q_r^k(x,a) - \tilde{Q}_r(x,a)| + Y_k \sup_{x,a} |Q_g^k(x,a) - \tilde{Q}_g(x,a)| + \epsilon' H \\
&\leq \epsilon'(1 + Y_k) + \epsilon' H \\
&\leq \epsilon'(1 + H + \xi)
\end{aligned} \tag{66}$$

where the first inequality follows from the property of supremum and the norm. The second inequality follows from the norm, and the fact that $|\tilde{Y}_k - Y_k| \leq \epsilon'$, and $|Q_g^k(x,a)| \leq H$. The third inequality follows from the fact that $\text{dist}(Q_j, \tilde{Q}_j) \leq \epsilon'$. \square

We now show that if there exist \tilde{Q}_j , and \tilde{Y}_k which are close to Q_j and Y_k , then the soft-max policy is also close.

Lemma 15. *Suppose that π is the soft-max policy (temp. coefficient $1/\alpha$) corresponding to the composite Q -functions $(Q_r^k + Y_k Q_g^k)$, i.e., $\forall a \in \mathcal{A}$*

$$\pi(a|\cdot) = \text{SOFT-MAX}_\alpha^a((Q_r(\cdot, \cdot) + Y_k Q_g(\cdot, \cdot))).$$

$\tilde{\pi}$ is the soft-max policy vector with the same temp. coefficient $1/\alpha$ corresponding to the composite Q -function $(\tilde{Q}_r + \tilde{Y}_k \tilde{Q}_g)$, i.e., $\forall a \in \mathcal{A}$,

$$\tilde{\pi}(a|\cdot) = \text{SOFT-MAX}_\alpha^a((\tilde{Q}_r(\cdot, \cdot) + \tilde{Y}_k \tilde{Q}_g(\cdot, \cdot))).$$

then, for any state x ,

$$\|\pi(\cdot|x) - \tilde{\pi}(\cdot|x)\|_1 \leq 2\alpha\epsilon'(1 + \xi + H) \tag{67}$$

where $\pi(\cdot|x) = \{\pi(a|x)\}_{a \in \mathcal{A}}$ and $\tilde{\pi}(\cdot|x) = \{\tilde{\pi}(a|x)\}_{a \in \mathcal{A}}$ when $\text{dist}(Q_{j,h}^k, \tilde{Q}_j) \leq \epsilon'$ for $j = r, g$, and $|\tilde{Y}_k - Y_k| \leq \epsilon'$.

Proof. Let $\text{Exp}^\alpha(P)$ be a soft-max corresponding to the vector P , i.e., the i -th component of $\text{Exp}^\alpha(P)$ is

$$\frac{\exp(\alpha P_i)}{\sum_i \exp(\alpha P_i)}.$$

Note from Theorem 4.4 in [29] then, we have

$$\|\text{Exp}^\alpha(P_1) - \text{Exp}^\alpha(P_2)\|_1 \leq 2\alpha \|P_1 - P_2\|_\infty \tag{68}$$

for two vectors P_1 and P_2 .

Now note that in our case for a given state x , π is equivalent to $\text{Exp}^\alpha(Q_{r,h}^k(x, \cdot) + Y_k Q_{g,h}^k(x, \cdot))$, and $\tilde{\pi}$ is equivalent to $\text{Exp}^\alpha(\tilde{Q}_r(x, \cdot) + \tilde{Y}_k \tilde{Q}_g(x, \cdot))$. Then from (68) and the fact that $\text{dist}(Q_{r,h}^k + Y_k Q_{g,h}^k, \tilde{Q}_r + \tilde{Y}_k \tilde{Q}_g) \leq \epsilon'(1 + \xi + H)$ (by Corollary 1) we have

$$\|\pi(\cdot|x) - \tilde{\pi}(\cdot|x)\|_1 \leq 2\alpha\epsilon'(1 + \xi + H) \tag{69}$$

Hence, the result follows. \square

Based on the above two lemmas we show that when the Q -functions are close, the value functions in the class \mathcal{V}_j are also close.

Lemma 16. *There exists $\tilde{V}_j \in \mathcal{V}_j$ such that*

$$\text{DIST}(V_j^k, \tilde{V}_j) \leq H2\alpha\epsilon'(1 + \xi + H) + \epsilon', \quad (70)$$

where $\text{dist}(\tilde{Q}_j, Q_j) \leq \epsilon'$, $\tilde{Q}_j \in \mathcal{Q}_j$ for all j ;

$$\tilde{V}_j(\cdot) = \sum_a [\tilde{\pi}(a|\cdot)\tilde{Q}_j(\cdot)],$$

$$\tilde{\pi}(a|\cdot) = \text{SOFT-MAX}_\alpha^a((\tilde{Q}_r(\cdot, \cdot) + \tilde{Y}_k\tilde{Q}_g(\cdot, \cdot)), \quad \forall a \in \mathcal{A}$$

$$|\tilde{Y}_k - Y_k| \leq \epsilon'.$$

Proof. For any x ,

$$\begin{aligned} & V_j^k(x) - \tilde{V}_j(x) \\ &= \left| \sum_a \pi(a|x)Q_j^k(x, a) - \sum_a \tilde{\pi}(a|x)\tilde{Q}_j(x, a) \right| \\ &= \left| \sum_a \pi(a|x)\tilde{Q}_j^k(x, a) - \sum_a \pi(a|x)\tilde{Q}_j(x, a) + \sum_a \pi(a|x)\tilde{Q}_j(x, a) - \sum_a \tilde{\pi}(a|x)\tilde{Q}_j(x, a) \right| \\ &\leq \left| \sum_a \pi(a|x)Q_j^k(x, a) - \sum_a \pi(a|x)\tilde{Q}_j(x, a) \right| + \left| \sum_a \pi(a|x)\tilde{Q}_j(x, a) - \sum_a \tilde{\pi}(a|x)\tilde{Q}_j(x, a) \right| \\ &\leq \epsilon' + \|\pi(\cdot|x) - \tilde{\pi}(\cdot|x)\|_1 \|\tilde{Q}_j(x)\|_\infty \\ &\leq \epsilon' + H2\alpha\epsilon'(1 + \xi + H) \end{aligned} \quad (71)$$

where we use the fact that $\text{dist}(Q_j^k, \tilde{Q}_r) \leq \epsilon'$, and $\sum_a \pi(a|x) = 1$ for the first term and the Holder's inequality in the second term for the second last inequality. For the last inequality, we use Lemma 15, and the fact that $\tilde{Q}_j(x, a) \leq H$ for any (x, a) . Hence, we have the result. \square

Note that when $\alpha = \frac{\log(|\mathcal{A}|)K}{2(1 + \xi + H)}$ as we have in Algorithm 1, the right hand side in (70) becomes

$$\epsilon' + \log(|\mathcal{A}|)KH\epsilon' \quad (72)$$

We introduce one more notation which we use to prove Lemma 13.

Definition 8. *Let \mathcal{C}_ξ^ϵ be an ϵ cover for $Y \in [0, \xi]$. Hence, $|\mathcal{C}_\xi^\epsilon| \leq \left(1 + \frac{\xi}{\epsilon}\right)$*

Note that \mathcal{C}_ξ^ϵ consists of points which is ϵ -close to any point within the interval $[0, \xi]$. Since we have defined ϵ -cover for all the parameters, we are now ready to prove Lemma 13.

Proof. Fix an ϵ . Let $\epsilon' = \frac{\epsilon}{H2\alpha(1 + \xi + H) + 1}$, then from Lemma 16, we have $\text{DIST}(V_j^k, \tilde{V}_j) \leq \epsilon$.

Thus, we only need to find parameters in the ϵ' -covering of the Q -functions as described in Lemma 14 in order to obtain ϵ -close value function.

Recall the Definition 7. Then, there exists $\tilde{w}_r, \tilde{w}_g \in \mathcal{C}_w^{\epsilon'}$ such that $\|\tilde{w}_r - w_r\| \leq \frac{\epsilon'}{2}$, $\|\tilde{w}_g - w_g\| \leq \frac{\epsilon'}{2}$.

Further, there exists $\tilde{\mathbf{A}}_2 \in \mathcal{C}_\Lambda^{\epsilon'}$ such that $\|\tilde{\mathbf{A}} - \tilde{\mathbf{A}}\|_F \leq \frac{\epsilon'^2}{4}$, $\tilde{\mathbf{A}} = \beta^2(\Lambda^k)^{-1}$, $\tilde{\mathbf{A}} = \beta^2(\tilde{\Lambda})^{-1}$, for some $\tilde{\Lambda}$, and \tilde{Y}_k, \tilde{Y}_k such that $|Y_k - \tilde{Y}_k| \leq \epsilon'$. Then we obtain \tilde{Q}_j parameterized by $(\tilde{w}_j, \beta, \tilde{\Lambda})$ for $j = r, g$, such that $\text{dist}(Q_j, \tilde{Q}_j) \leq \epsilon'$ (by Lemma 14).

Now define $\tilde{V}_j = \sum_a \tilde{\pi}(a|\cdot)\tilde{Q}_j$, where

$$\tilde{\pi}(a|\cdot) = \text{SOFT-MAX}_\alpha^a((\tilde{Q}_r(\cdot, \cdot) + \tilde{Y}_k\tilde{Q}_g(\cdot, \cdot))).$$

Thus, from Lemma 16, we have $\text{DIST}(V_j^k, \tilde{V}_j) \leq \epsilon$. Hence, there exists \tilde{V}_j parameterized by $\tilde{w}_r, \tilde{w}_g, \tilde{Y}_k, \tilde{\mathbf{A}}$, such that $\text{Dist}(\tilde{V}_j, V_j^k) \leq \epsilon$. Hence, $N_\epsilon^V \leq |\mathcal{C}_w^{\epsilon'}| |\mathcal{C}_A^{\epsilon'}| |\mathcal{C}_\xi^{\epsilon'}|$. Thus, from Lemma 14 and Definition 8, the ϵ -covering number $N_\epsilon^{V_j}$ for the set \mathcal{V}_j satisfies the following

$$\log N_\epsilon^{V_j} \leq d \log \left(1 + 8H \frac{\sqrt{dk}}{\sqrt{\lambda \epsilon'}} \right) + d^2 \log \left[1 + 8d^{1/2} \beta^2 / (\lambda (\epsilon')^2) \right] + \log \left(1 + \frac{\xi}{\epsilon'} \right).$$

Hence, the result follows. \square

From Lemma 13, note that we need ϵ' covering for the Q -functions where $\epsilon' = \frac{\epsilon}{(H2\alpha(1 + \xi + H) + 1)}$ if we need to bound $\text{DIST}(V_j, \tilde{V}_j)$ by ϵ .

Now, we are ready to prove Lemma 8.

Proof. By Lemma 13, we know that there exists \tilde{V}_j in the ϵ -covering for \mathcal{V}_j such that for every x ,

$$V_j(x) = \tilde{V}_j(x) + \Delta V(x) \quad (73)$$

where $\sup_x \Delta V(x) \leq \epsilon$.

Hence,

$$\begin{aligned} \left\| \sum_{\tau=1}^k \phi^\tau (V_j(x_\tau) - \mathbb{E}[V_j(x_\tau) | \mathcal{F}_{\tau-1}]) \right\|_{(\Lambda^k)^{-1}}^2 &\leq 2 \left\| \sum_{\tau=1}^k \phi^\tau (\tilde{V}_j(x_\tau) - \mathbb{E}[\tilde{V}_j(x_\tau) | \mathcal{F}_{\tau-1}]) \right\|_{(\Lambda^k)^{-1}}^2 \\ &\quad + 2 \left\| \sum_{\tau=1}^k \phi^\tau (\Delta V(x_\tau) - \mathbb{E}[\Delta V(x_\tau) | \mathcal{F}_{\tau-1}]) \right\|_{(\Lambda^k)^{-1}}^2 \end{aligned} \quad (74)$$

The last expression is bounded by $\frac{8k^2\epsilon^2}{\lambda}$.

Now, we bound the first term. Note from Lemma 13 that in order to obtain \tilde{V}_j which satisfies (73), we need to obtain we need N_ϵ^V number of elements to obtain such $(\tilde{w}_r, \tilde{w}_g, \beta, \tilde{\Lambda}, \tilde{Y})$. Such \tilde{V}_j is independent of samples. Hence, we can use the Elliptical lemma for self-normalization (Theorem 2). From Theorem 2 and the union bound we obtain

$$\left\| \sum_{\tau=1}^k \phi^\tau (\tilde{V}_j(x_\tau) - \mathbb{E}[\tilde{V}_j(x_\tau) | \mathcal{F}_{\tau-1}]) \right\|_{(\Lambda^k)^{-1}}^2 \leq 2H^2 \left[d \log \left(\frac{k + \lambda}{\lambda} \right) + \log \left(\frac{N_\epsilon^V}{\delta} \right) \right] \quad (75)$$

where N_ϵ^V is upper bounded in (62). β is equal to $C_1 dH \sqrt{\iota}$ for some constant C_1 , and $\iota = \log(\log(|\mathcal{A}|)4dT/p)$. Further, $\xi = 2H/\gamma$ (by Definition 3). We obtain from (75)

$$\begin{aligned} \left\| \sum_{\tau=1}^k \phi^\tau (\tilde{V}_j(x_\tau) - \mathbb{E}[\tilde{V}_j(x_\tau) | \mathcal{F}_{\tau-1}]) \right\|_{(\Lambda^k)^{-1}}^2 &\leq \\ 4H^2 \left[\frac{d}{2} \log \left(\frac{k + \lambda}{\lambda} \right) + d \log \left(1 + \frac{8H\sqrt{dk}}{\epsilon' \sqrt{\lambda}} \right) + d^2 \log \left(1 + \frac{8d^{1/2} \beta^2}{\epsilon'^2 \lambda} \right) + \log \left(1 + \frac{2H}{\gamma \epsilon'} \right) + \log \left(\frac{4}{p} \right) \right] \end{aligned} \quad (76)$$

where $\epsilon' = \frac{\epsilon}{(H2\alpha(1 + \xi + H) + 1)}$. Set $\epsilon = \frac{dH}{k}$, $\lambda = 1$. Thus, $\epsilon' = \frac{dH}{(2H\alpha(1 + \xi + H) + 1)k}$.

Plugging in the above, and putting $\alpha = \frac{\log(|\mathcal{A}|)K}{2(1 + \xi + H)}$, we obtain from (76)

$$\left\| \sum_{\tau=1}^k \phi^\tau (\tilde{V}_j(x_\tau) - \mathbb{E}[\tilde{V}_j(x_\tau) | \mathcal{F}_{\tau-1}]) \right\|_{\Lambda_k}^2 \leq C_2 H^2 d^2 \log \left(\frac{4(C_1 + 1) \log(|\mathcal{A}|)dT}{p} \right) \quad (77)$$

for some constant C_2 where $C_1 = 256(1 + 1/\gamma)$. Hence, the result follows. \square

F Supporting Results

The following result is shown in [30] and in Lemma D.2 in [14].

Lemma 17. *Let $\{\phi_t\}_{t \geq 0}$ be a sequence in \mathbb{R}^d satisfying $\sup_{t \geq 0} \|\phi_t\| \leq 1$. For any $t \geq 0$, we define $\Lambda_t = \Lambda_0 + \sum_{j=0}^t \phi_j \phi_j^T$. Then if the smallest eigen value of Λ_0 be at least 1, we have*

$$\log \left[\frac{\det(\Lambda_h^{k+1})}{\det(\Lambda_h^1)} \right] \leq \sum_{k=1}^K (\phi_h^k)^T (\Lambda_h^k)^{-1} \phi_h^k \leq 2 \log \left[\frac{\det(\Lambda_h^{k+1})}{\det(\Lambda_h^1)} \right] \quad (78)$$

Theorem 2. [Concentration of Self-Normalized Process [30]] *Let $\{\epsilon_t\}_{t=1}^\infty$ be a real-valued stochastic process with corresponding filtration $\{\mathcal{F}_t\}_{t=0}^\infty$. Let $\epsilon_t | \mathcal{F}_{t-1}$ be a zero mean and σ sub-Gaussian, i.e., $\mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}] = 0$, and*

$$\forall \zeta \in \mathbb{R}, \quad \mathbb{E}[e^{\zeta \epsilon_t} | \mathcal{F}_{t-1}] \leq e^{\zeta^2 \sigma^2 / 2}. \quad (79)$$

Let $\{\phi_t\}_{t=1}^\infty$ be a \mathbb{R}^d -valued Stochastic process where $\phi_t \in \mathcal{F}_{t-1}$. Assume $\Lambda_0 \in \mathbb{R}^{d \times d}$ is a positive-definite matrix, let, $\Lambda_t = \Lambda_0 + \sum_{j=0}^t \phi_j \phi_j^T$. Then for any $\delta > 0$ with probability at least $1 - \delta$, we have

$$\left\| \sum_{s=1}^t \phi_s \epsilon_s \right\|_{\Lambda_t^{-1}}^2 \leq 2\sigma^2 \log \left[\frac{\det(\Lambda_t)^{1/2} \det(\Lambda_0)^{-1/2}}{\delta} \right] \quad (80)$$

The next result characterizes the covering number of an Euclidean ball (Lemma 5.2 in [31]).

Lemma 18. [Covering Number of Euclidean Ball] *For any $\epsilon > 0$, the ϵ -covering number of the Euclidean ball in \mathbb{R}^d with radius R is upper bounded by $(1 + 2R/\epsilon)^d$.*

We have used the following result from the optimization which is proved in Lemma 9 in [7].

Lemma 19. *Let Y^* be the optimal dual variable, and $C \geq 2Y^*$, then, if*

$$V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\tilde{\pi}}(x_1) + C[b - V_{g,1}^{\tilde{\pi}}(x_1)]_+ \leq \delta \quad (81)$$

then

$$[b - V_{g,1}^{\tilde{\pi}}(x_1)]_+ \leq \frac{2\delta}{C}. \quad (82)$$

G Why does the Greedy Policy Fail?

In this section, using an example we show that the greedy-policy is not Lipschitz. Further, we can not use the greedy-policy on the composite Q -function to obtain the ϵ -close covering for individual reward and utility value functions.

Consider the following toy-example: Suppose that the cardinality of the action space $|A|$ is 2. $Q_{r,h}(x, a_1) = M$, $Q_{r,h}(x, a_2) = 1$, $Q_{g,h}(x, a_1) = 1$, $Q_{g,h}(x, a_2) = M + \epsilon/2$. Consider $Y = 1$, then, greedy policy based on the composite Q -function is to choose action a_2 .

Note the ϵ -closest values can be anywhere in the ball within ϵ radius centered around $Q_{j,h}$, and Y . Assume that the closest ϵ -cover for $Q_{j,h}$ be $\tilde{Q}_{j,h}$ such that $\tilde{Q}_{r,h}(x, a_1) = M + \epsilon/2$, $\tilde{Q}_{r,h}(x, a_2) = 1 - \epsilon/2$, $\tilde{Q}_{g,h}(x, a_1) = 1 + \epsilon/2$, $\tilde{Q}_{g,h}(x, a_2) = M$, and $\tilde{Y} = 1 - \epsilon/2$. Then, we have the composite Q -functions as

$$\begin{aligned} \tilde{Q}_{r,h}(x, a_1) + \tilde{Y} \tilde{Q}_{g,h}(x, a_1) &= M + \epsilon/2 + (1 - \epsilon/2)(1 + \epsilon/2) \\ \tilde{Q}_{r,h}(x, a_2) + \tilde{Y} \tilde{Q}_{g,h}(x, a_2) &= 1 - \epsilon/2 + (1 - \epsilon/2)M \end{aligned} \quad (83)$$

Hence, it is clear that the greedy policy based on the composite Q -function is to choose action a_1 . Hence, the policy is not Lipschitz. Thus, even though the change in the Q -function is only by ϵ -amount the decision changes from taking action a_2 to taking action a_1 in a deterministic fashion.

Now, we see the changes in the value function. Since the policy is greedy and the policy is to choose a_2 for the Q -function $Q_{j,h}$, then, $V_{r,h}(x) = Q_{r,h}(x, a_2) = 1$. On the other hand, the policy is to choose a_1 for the Q -functions $\tilde{Q}_{j,h}$. Hence, $\tilde{V}_{r,h}(x) = M + \epsilon/2$, hence, $|V_{r,h}^k(x) - \tilde{V}_{r,h}(x)| > M - 1$, and can be made arbitrarily large by making M arbitrarily large. Thus, the individual value functions can not be made close even though $Q_{j,h}$ and $\tilde{Q}_{j,h}$ are close. This is the reason we can not obtain ϵ -covering number if the greedy policy is based on the composite Q -functions.

Note that in the unconstrained case (equivalent to $Y_k = 0$), the decision would be to choose a_1 for both $Q_{r,h}$ and $\tilde{Q}_{r,h}$ if the policy is set at the greedy one. Hence, the value function would only differ by at most ϵ -amount. Hence, the greedy policy works for the unconstrained case.

H Analysis for Zero Constraint Violation

The main idea behind attaining zero constraint violation is to consider the following tighter optimization problem—

$$\text{maximize } \pi \in \Delta(\mathcal{A}|\mathcal{S}, \mathcal{H}) V_{r,1}^\pi(x_1) \quad \text{subject to } V_{g,1}^\pi(x_1) \geq b + \zeta. \quad (84)$$

Since we replace b by $b + \zeta$, we are basically solving the above tighter optimization problem. By ensuring that $\zeta \leq \gamma/2$, we can ensure that Slater's condition is always satisfied, and strong duality holds. We show that by carefully choosing ζ , we can achieve zero constraint violation with the same order on regret with respect to T . First, we introduce some notations which we use throughout this section.

Let $\pi^{\zeta,*}$ be the optimal solution of the optimization problem in (84). Since the Slater's condition holds, the strong duality holds by [19]. The optimal dual variable Y^ζ of this tighter problem is

$$Y^\zeta \leq \frac{V^{\pi^{\zeta,*}}(x_1) - V_{r,1}^{\pi^{\zeta,*}}(x_1)}{b + \gamma - (b + \zeta)} \leq 4H/\gamma \quad (85)$$

where the last inequality follows from the fact that $\zeta \leq \gamma/2$.

Now, we state the main result—

Theorem 3. *In Algorithm 1, replacing $b = b + \zeta$, and $\xi = 4H/\gamma$. We obtain, with probability $1 - p$,*

$$\begin{aligned} \text{Regret}(K) &\leq C\mathcal{O}(\sqrt{d^3 H^3 T \iota^2}) + KH\zeta/\delta \\ \text{Violation}(K) &\leq \max\{C'\mathcal{O}\left(\frac{2(1+\xi)}{\xi} \sqrt{d^3 H^3 T \iota^2}\right) - K\zeta, 0\} \end{aligned} \quad (86)$$

where $\zeta = \min\{C'\mathcal{O}\left(\frac{2(1+\xi)}{\xi} \frac{\sqrt{d^3 H^3 T \iota^2}}{K}\right), \gamma/2\}$.

When $C'\mathcal{O}\left(\frac{2(1+\xi)}{\xi} \frac{\sqrt{d^3 H^3 T \iota^2}}{K}\right) \leq \gamma/2$, then the constraint is upper bounded by 0. Hence, for large enough K , we can achieve zero violation. However, by plugging the value of ζ , we obtain the upper bound on regret as

$$\text{Regret}(K) \leq C\mathcal{O}(\sqrt{d^3 H^3 T \iota^2}) + C'H\mathcal{O}\left(\frac{2(1+\xi)}{\xi} \sqrt{d^3 H^3 T \iota^2}\right)$$

where we replace the upper bound of ζ by $C'\mathcal{O}\left(\frac{2(1+\xi)}{\xi} \frac{\sqrt{d^3 H^3 T \iota^2}}{K}\right)$. Thus, the upper bound on

regret is $\tilde{\mathcal{O}}(\sqrt{d^3 H^5 T})$. Hence, the order on regret with respect to T is maintained. However, there is an additional H factor in front of the regret. A concurrent work [32] on model-based discounted horizon tabular setup using a generator model shows that extra H is unavoidable if one wants to

achieve zero violation.* Even though the setup is different, it seems that the extra H factor is unavoidable.

Proof. First, we prove the upper bound on regret. The regret can be decomposed as the following -

$$\text{Regret}(K) = \sum_{k=1}^K (V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi^{\zeta,*}}(x_1)) + \sum_{k=1}^K (V_{r,1}^{\pi^{\zeta,*}}(x_1) - V^{\pi^k}(x_1)) \quad (87)$$

The first term can be bounded with the help of the following lemma (the proof for finite state is in [13], the extension to linear MDP is provided after this proof)–

Lemma 20. *If $\pi^{\zeta,*}$ is the optimal solution of (84), then*

$$V_{r,1}^{\pi^*}(x_1) - V_{r,1}^{\pi^{\zeta,*}}(x_1) \leq H \frac{\zeta}{\gamma}. \quad (88)$$

Since the tighter optimization problem is also CMDP, we note that the second term in the right hand side of (87) is essentially the regret of the tighter CMDP.

Hence, from Theorem 1 and Lemma 20 we obtain the expression of the regret bound in Theorem 3.

Constraint Violation: Again applying Theorem 1 to the tighter optimization problem (84), we obtain

$$\begin{aligned} \sum_{k=1}^K (b + \zeta - V_{g,1}^{\pi^k}(x_1))_+ &\leq C' \mathcal{O}\left(\frac{2(1+\xi)}{\xi} \sqrt{d^3 H^3 T \iota^2}\right) \\ \sum_{k=1}^K (b + \zeta - V_{g,1}^{\pi^k}(x_1)) &\leq \sum_{k=1}^K (b + \zeta - V_{g,1}^{\pi^k}(x_1))_+ \leq C' \mathcal{O}\left(\frac{2(1+\xi)}{\xi} \sqrt{d^3 H^3 T \iota^2}\right) \end{aligned}$$

Hence,

$$\sum_{k=1}^K (b + \zeta - V_{g,1}^{\pi^k}(x_1)) \leq C' \mathcal{O}\left(\frac{2(1+\xi)}{\xi} \sqrt{d^3 H^3 T \iota^2}\right)$$

Thus,

$$\sum_{k=1}^K (b - V_{g,1}^{\pi^k}(x_1)) \leq C' \mathcal{O}\left(\frac{2(1+\xi)}{\xi} \sqrt{d^3 H^3 T \iota^2}\right) - K\zeta$$

Hence, we have

$$\sum_{k=1}^K [b - V_{g,1}^{\pi^k}(x_1)]_+ \leq \max\{C' \mathcal{O}\left(\frac{2(1+\xi)}{\xi} \sqrt{d^3 H^3 T \iota^2}\right) - K\zeta, 0\}$$

Thus, the result follows. \square

Proof of Lemma 20: We, first, introduce a few notations.

Let $\nu_h^\pi(x)$ for $h = 2, \dots, H$ be

$$\nu_h^\pi(x) = \int_{x'} \sum_a \pi_h(a|x') \phi(x', a)^T \mu_{h-1}(x) d\nu_{h-1}^\pi(x')$$

and $\nu_1(x)$ is the initial distribution of the state. $\nu_h^\pi(x)$ is the distribution of the state at step h while following the policy π . It is the state occupation measure at step h .

Also, $\nu_h(x, a) = \pi_h(a|x) \nu_h(x)$ is the state-action occupation measure at step h . Hence,

$$V_{j,1}^\pi(x_1) = \sum_h \int_{x,a} j_h(x, a) d\nu_h(x, a) \quad (89)$$

*[32] provides a sample complexity guarantee for the discounted horizon setup with discount factor γ . One can convert the result in the discounted setup to the episodic setup by equating $1/(1-\gamma) = H$.

Now, $\nu_h^*(x, a)$ corresponds to the state-action occupancy measure for the optimal policy π^* . Then, $\nu_h^\zeta(x, a) = (1 - \zeta/\gamma)\nu_h^*(x, a) + \zeta/\gamma\nu_h^{\bar{\pi}}(x, a)$. Now, we are resy to prove Lemma 20.

We have

$$\sum_h \int_{x,a} g_h(x, a) d\nu_h^\zeta(x, a) \geq (1 - \zeta/\gamma)b + \zeta/\gamma(b + \gamma) = b + \zeta \quad (90)$$

Hence, the state-action occupancy measure $\nu_1^\zeta(x, a)$ is feasible for the tightened CMDP. Now, we have

$$\begin{aligned} \sum_h \int_{x,a} r_h(x, a) d\nu_h^\zeta(x, a) &= (1 - \zeta/\delta) \sum_h \int_{x,a} r_h(x, a) d\nu_h^*(x, a) + \zeta/\delta \sum_h \int_{x,a} r_h(x, a) d\nu_h^{\bar{\pi}}(x, a) \\ &\geq (1 - \zeta/\delta)V_{r,1}^*(x_1) \end{aligned}$$

Since $\nu_1^\zeta(x, a)$ is feasible, then $V_{r,1}^{\pi^{\zeta,*}}(x_1) \geq \sum_h \int_{x,a} r_h(x, a) d\nu_h^\zeta(x, a)$. Thus,

$$V_{r,1}^*(x_1) - V_{r,1}^{\zeta,*}(x_1) \leq \zeta/\delta V_{r,1}^*(x_1) \leq \frac{\zeta}{\delta} H \quad (91)$$

Hence, the result follows. \square

Remark 3. Note that [13] and [6] use Lyapunov Drift analysis to obtain constraint violation for finite state space (tabular setting). To obtain the zero violation, their approach is also similar to ours where they also consider a tighter optimization problem, and then, carefully choosing the parameter of the tighter optimization. However, one key difference is that– they did not use any upper bound on the dual variable. Rather they rely on the Hajek’s Lemma [33] to establish a finite bound on the dual variable for the good event. The question is whether we can use similar technique in the model-free linear function approximation.

Note that we need to have an upper bound on the dual-variable (irrespective of the good and bad event) since the ϵ -covering number (Lemma 13) depends on ξ , the upper bound on the dual variable. Hence, we need to truncate the dual variable if it exceeds ξ . However, if we truncate the dual-variable, we can not use the Lyapunov-Drift analysis to bound the violation. Since in the Lyapunov-drift analysis, it relies on the fact that the magnitude of the dual-variable (or, the queue length) lower bounds the total violation (see the analysis at page 21 in [6]). Since the magnitude of the dual-variable is bounded for the good-event, one then obtain the upper bound of the violation. However, one can not extend the same analysis if we truncate the dual-variable.

Hence, our analysis is based on the results from the convex optimization (Lemma 19). Our approach to obtain zero violation for large enough K relies on different tools compared to the Lyapunov-Drift analysis and new of a kind.

I Comparison with other approaches to show uniform concentration lemma for individual value function in model-free setup

[34] considered a zero-sum. linear Markov game setup. The paper proposed an approach where they truncate w_h^k , Λ_h^k to ϵ -close value of w , and Λ respectively. Subsequently, they obtained equilibrium policies using these ϵ -close values. Then the proposed algorithm uses the above equilibrium policy attained using ϵ -close values. Since these ϵ -close values are predetermined using the ϵ -covering set of w and Λ (as we have described the ϵ -covering set in Lemma 14), hence, one can apply uniform concentration lemma for each individual value function [34] with error of at most ϵ . We can also apply the similar trick in our set-up where we truncate the obtained $w_{j,h}^k$, and Λ_h^k to one of the ϵ -close values and then we can set the policy as the greedy one with respect to the composite approximated state-action value function. The above method would also provide a log ϵ -covering number of $\mathcal{O}(\log(K))$ for each individual value function.

However, our soft-max based approach has several advantages compared to the above approach.

(i) *Computation efficiency.* The alternative method explicitly rounds the Q -function to its ϵ -close one in the algorithm, which requires an additional $\mathcal{O}(d^2)$ computation even with an efficient implementation (i.e., rounding on the fly) (See Section 3.3 in [34]). In contrast, our soft-max works directly

with the actual Q -function without any ϵ -close rounding. That is, ϵ -net argument is only used *in the analysis* in our setting.

(ii) *Stochastic policy.* A key fact about constrained MDPs is that the optimal policy is usually stochastic. Although the greedy-policy in the alternative method can approach the optimal policy in an average sense, in each episode, it could be far away from the optimal policy. Even though in our setting, α scales with K , our approach puts 'almost' similar probability among the composite Q -functions with 'almost' same values. However, in the greedy policy, it chooses the action corresponding to the highest state-action value function. Thus, such an approach can never be close to the optimal policy.

(iii) *General applicability.* The alternative method relies heavily on the fact that there exists an efficient implementation of rounding on the fly in linear MDPs so that there is no need to construct an explicit ϵ -net. However, beyond linear MDPs, it might not be possible to find an efficient rounding algorithm on the fly, and hence an even larger additional computation is required. In contrast, our soft-max based algorithm builds on the intrinsic smoothness-approximation trade-off in soft-max to establish uniform concentration and approximate optimism, which could potentially be generalized to other settings. Further, soft-max policy is popular, hence, our approach would provide the base for proving regret and violation bound for the general function approximation setup beyond the linear function approximation.