
Learning Policies with Zero or Bounded Constraint Violation for Constrained MDPs

Tao Liu*

Texas A&M University
 tliu@tamu.edu

Ruida Zhou*

Texas A&M University
 ruida@tamu.edu

Dileep Kalathil

Texas A&M University
 dileep.kalathil@tamu.edu

P. R. Kumar

Texas A&M University
 prk@tamu.edu

Chao Tian

Texas A&M University
 chao.tian@tamu.edu

Abstract

We address the issue of safety in reinforcement learning. We pose the problem in an episodic framework of a constrained Markov decision process. Existing results have shown that it is possible to achieve a reward regret of $\tilde{\mathcal{O}}(\sqrt{K})$ while allowing an $\tilde{\mathcal{O}}(\sqrt{K})$ constraint violation in K episodes. A critical question that arises is whether it is possible to keep the constraint violation even smaller. We show that when a strictly safe policy is known, then one can confine the system to zero constraint violation with arbitrarily high probability while keeping the reward regret of order $\tilde{\mathcal{O}}(\sqrt{K})$. The algorithm which does so employs the principle of optimistic pessimism in the face of uncertainty to achieve safe exploration. When no strictly safe policy is known, though one is known to exist, then it is possible to restrict the system to bounded constraint violation with arbitrarily high probability. This is shown to be realized by a primal-dual algorithm with an optimistic primal estimate and a pessimistic dual update.

1 Introduction

Reinforcement learning (RL) addresses the problem of learning an optimal control policy that maximizes the expected cumulative reward while interacting with an unknown environment [29]. Standard RL algorithms typically focus only on maximizing a single objective. However, in many real-world applications, the control policy learned by an RL algorithm has to additionally satisfy stringent safety constraints [10, 4]. For example, an autonomous vehicle may need to reach its destination in the minimum possible time without violating safety constraints such as crossing the middle of the road. The Constrained Markov Decision Process (CMDP) [2, 14] formalism, where one seeks to maximize a reward while satisfying safety constraints, is a standard approach for modeling the necessary safety criteria of a control problem via constraints on cumulative costs.

Several policy-gradient-based algorithms have been proposed to solve CMDPs. Lagrangian-based methods [30, 28, 25, 20] formulate the CMDP problem as a saddle-point problem and optimize it via primal-dual methods, while Constrained Policy Optimization [1, 33] (inspired by the trust region policy optimization [27]) computes new dual variables from scratch at each update to maintain constraints during learning. Although these algorithms provide ways to learn an optimal policy, performance guarantees about reward regret, safety violation or sample complexity are rare.

One class of RL algorithms for which performance guarantees are available follow the principle of Optimism in the Face of Uncertainty (OFU) [9, 11, 26], and provide an $\tilde{\mathcal{O}}(\sqrt{K})$ guarantee for the

*The first two authors contributed equally.

Table 1: Regret and constraint violation comparisons for algorithms on episodic CMDPs

Algorithm	Regret ²	Constraint violation ²
OPDOP [9]	$\tilde{\mathcal{O}}(H^3\sqrt{ \mathcal{S} ^2 \mathcal{A} K})$	$\tilde{\mathcal{O}}(H^3\sqrt{ \mathcal{S} ^2 \mathcal{A} K})$
OptCMDP [11] ³	$\tilde{\mathcal{O}}(H^2\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$	$\tilde{\mathcal{O}}(H^2\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$
OptCMDP-bonus [11] ³	$\tilde{\mathcal{O}}(H^2\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$	$\tilde{\mathcal{O}}(H^2\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$
OptDual-CMDP [11] ³	$\tilde{\mathcal{O}}(H^2\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$	$\tilde{\mathcal{O}}(H^2\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$
OptPrimalDual-CMDP [11] ³	$\tilde{\mathcal{O}}(H^2\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$	$\tilde{\mathcal{O}}(H^2\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$
C-UCRL [34] ⁴	$\tilde{\mathcal{O}}(T^{\frac{3}{4}})$	0
OptPess-LP	$\tilde{\mathcal{O}}(\frac{H^3}{\tau-c^0}\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$	0
OptPess-PrimalDual	$\tilde{\mathcal{O}}(\frac{H^3}{\tau-c^0}\sqrt{ \mathcal{S} ^3 \mathcal{A} K})$	$\mathcal{O}(1)$ ⁵

reward regret, where K is the number of episodes. However, these algorithms also have $\tilde{\mathcal{O}}(\sqrt{K})$ safety violations. Such significant violation of the safety constraints during learning may be unacceptable in many safety-critical real-world applications such as the control of autonomous vehicles or power systems. These applications demand a class of safe RL algorithms that can provably guarantee safety during learning. With this goal in mind, we aim to answer the following open theoretical question in this paper:

Can we design safe RL algorithms that can achieve an $\tilde{\mathcal{O}}(\sqrt{K})$ regret with respect to the performance objective, while guaranteeing zero or bounded safety constraint violation with arbitrarily high probability?

We answer the above question affirmatively by proposing two algorithms and establishing their stringent safety performance during learning. Our focus is on the tabular episodic constrained RL setting (unknown transition probabilities, rewards, and costs). The key idea behind both algorithms is a concept used earlier for safe exploration in constrained bandits [24, 21], which we call “Optimistic Pessimism in the Face of Uncertainty (OPFU)” here. The optimistic aspect incentivizes the algorithm for using exploration policies that can visit new state-action pairs, while the pessimistic aspect disincentivizes the algorithm from using exploration policies that can violate safety constraints. By carefully balancing optimism and pessimism, the proposed algorithms guarantee zero or bounded safety constraint violation during learning while achieving an $\tilde{\mathcal{O}}(\sqrt{K})$ regret with respect to the reward objective.

The two algorithms address two different classes of the safe learning problem: whether a strictly safe policy is known a priori or not. The resulting exploration strategies are very different in the two cases.

1. **OptPess-LP Algorithm:** This algorithm assumes the prior knowledge of a strictly safe policy. It ensures zero safety constraint violation during learning with high probability and utilizes the linear programming (LP) approach for solving a CMDP problem. The algorithm achieves a reward regret of $\tilde{\mathcal{O}}(\frac{H^3}{\tau-c^0}\sqrt{|\mathcal{S}|^3|\mathcal{A}|K})$ with respect to the performance objective, where H is the number of steps per episode, τ is the given constraint on safety violation, c^0 is the known safety constraint value of a strictly safe policy π^0 , and $|\mathcal{S}|$ and $|\mathcal{A}|$ are the number of states and actions respectively.
2. **OptPess-PrimalDual Algorithm:** This algorithm addresses the case where no strictly safe policy, but a feasible strictly safe cost is known. By allowing a bounded (in K) safety cost, it opens up space for exploration. The OptPess-PrimalDual algorithm avoids linear

²This table is presented for $K \geq \text{poly}(|\mathcal{S}|, |\mathcal{A}|, H)$, with polynomial terms independent of K omitted.

³Efroni et al. [11] use \mathcal{N} , the maximum number of non-zero transition probabilities across the entire state-action space, in their regret and constraint violation analysis. For consistency, we use $|\mathcal{S}|^2|\mathcal{A}|$ to bound \mathcal{N} .

⁴Zheng et al. [34] assumes a known transition kernel and analyzes regret in the long-term average setting.

⁵The detailed constraint violation is $\mathcal{O}\left(C''H + H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|C''\log(C''/\delta')}\right)$, which is independent of K . Here, $\delta' = \delta/(16|\mathcal{S}|^2|\mathcal{A}|H)$ and $C'' = \mathcal{O}(\frac{H^4|\mathcal{S}|^3|\mathcal{A}|}{(\tau-c^0)^2}\log\frac{H^4|\mathcal{S}|^3|\mathcal{A}|}{(\tau-c^0)^2\delta'})$.

programming and its attendant complexity and exploits the primal-dual approach for solving a CMDP problem. The proposed approach improves the computational tractability, while ensuring a bounded safety constraint violation during learning and a reward regret of $\tilde{\mathcal{O}}\left(\frac{H^3}{\tau-c}\sigma\sqrt{|\mathcal{S}|^3|\mathcal{A}|K}\right)$ with respect to the objective.

Compared with the other methods listed in Table 1, though the proposed algorithms have an additional $H/(\tau-c^0)$ or $\sqrt{|\mathcal{S}|}/(\tau-c^0)$ factor in the regret bounds, they are able to reduce the constraint violation to zero or constant with high probability. This improvement in safety can be extremely important for many mission-critical applications.

1.1 Related work

The problem of learning an optimal control policy that satisfies safety constraints has been studied both in the RL setting and the multi-armed bandits setting.

Constrained RL: Several policy-gradient algorithms have seen success in practice [30, 28, 25, 20, 1, 33]. Also of interest are works which utilize Gaussian processes to model the transition probabilities and value functions [5, 31, 19, 7].

Several algorithms with provable guarantees are closely related to our work. Zheng et al. [34] considered the constrained RL problem in an infinite horizon setting with unknown reward and cost functions. The approach is similar to the UCRL2 algorithm [16] and achieves a sub-linear reward regret $\tilde{\mathcal{O}}(T^{\frac{3}{4}})$ while satisfying constraints with high probability during learning. In contrast to this work, we consider the setting of unknown transition probabilities. Efroni et al [11] focused on the episodic setting of unknown non-stationary transitions over a finite horizon, attaining both a reward regret and a constraint violation of $\tilde{\mathcal{O}}(H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|K})$. Ding et al. [9] studied an episodic setting with linear function approximation (suitable for large state space cases), and proposed algorithms that can achieve $\tilde{\mathcal{O}}(dH^3\sqrt{K})$ for both the regret and the constraint violation (where d is the dimension of the feature mapping). The regret analysis in [9] can be easily extended to the tabular case, yielding $\tilde{\mathcal{O}}(H^3\sqrt{|\mathcal{S}|^2|\mathcal{A}|K})$.

Constrained Multi-Armed Bandits: Multi-armed bandit problems are special cases of MDPs, with both the number of states as well as the episode length being one. Linear bandits with constraints (satisfied with high probability) have been investigated in different settings. One setting, referred to as conservative bandits [32, 18, 13], requires the cumulative reward to remain above a fixed percentage of the cumulative reward of a given baseline policy. Another setting is where each arm is associated with two unknown distributions (similar to our setting), generating reward and cost signals respectively [3, 24, 21, 22].

2 Problem formulation

A finite-horizon constrained non-stationary MDP model is defined as a tuple $M = (\mathcal{S}, \mathcal{A}, H, P, r, c, \tau, \mu)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, H is the number of steps in each episode, $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the unknown reward function of interest, $c : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the unknown safety cost function used to model the constraint violation, $\tau \in (0, H]$ is the given constant used to define the safety constraint, and μ is the known initial distribution of the state. $P(\cdot|s, a) \in \Delta_{\mathcal{S}}^H, \forall s \in \mathcal{S}, \forall a \in \mathcal{A}$, where $\Delta_{\mathcal{S}}$ is the $|\mathcal{S}|$ -dimensional probability simplex, and $P_h(s'|s, a)$ is the unknown transition probability that the next state is s' when action a is taken for state s at step h .

Some further notation is necessary to define the problem. A (randomized Markov) policy is defined by a map $\pi : \mathcal{S} \times [H] \rightarrow \Delta_{\mathcal{A}}$, with $\pi_h(a|s)$ being the probability of taking action a in state s at time step h . With S_t and A_t representing the state and the action at time t respectively, let

$$V_h^{\pi}(s; g, P) := \mathbb{E}_{P, \pi} \left[\sum_{t=h}^H g(S_t, A_t) | S_h = s \right], \quad \forall s \in \mathcal{S}$$

denote the expected cumulative value with respect to a function $g : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}_+$ under P for a policy π over a time interval $[h, h+1, \dots, H]$. With slightly abuse of notation, we use $V_1^\pi(\mu; g, P)$ to denote $\mathbb{E}_{S_1 \sim \mu}[V_1^\pi(S_1; g, P)]$.

In the formulation below there are a total of K episodes with H steps each. Each episode $k \in [K]$ begins with an initial probability distribution μ for \mathcal{S}_1 . Then, the agent determines a randomized Markov policy π^k for that episode based on the information gathered from the previous episodes, and executes it. At time step h during the execution of the k -th episode, after taking action A_h^k at state S_h^k , the agent receives a noisy reward and cost of $R_h^k(S_h^k, A_h^k) = r_h(S_h^k, A_h^k) + \xi_h^k(S_h^k, A_h^k; r)$ and $C_h^k(S_h^k, A_h^k) = c_h(S_h^k, A_h^k) + \xi_h^k(S_h^k, A_h^k; c)$, respectively.

Assumption 2.1 (Sub-Gaussian noise). *For all $h \in [H], k \in [K]$, the reward and cost noise random variables are conditionally independent zero-mean $1/2$ -sub-Gaussian, i.e., $\mathbb{E}[\xi_h^k | \mathcal{F}_{k-1}] = 0$, $\mathbb{E}[\exp(\lambda \xi_h^k) | \mathcal{F}_{k-1}] \leq \exp(\lambda^2/4)$, $\forall \lambda \in \mathbb{R}$. Here \mathcal{F}_k is the σ -algebra generated by the random variables up to episode k .*

Let π^* denote the optimal policy of the following CMDP model:

$$\max_{\pi} V_1^\pi(\mu; r, P) \quad \text{s.t.} \quad V_1^\pi(\mu; c, P) \leq \tau. \quad (1)$$

A policy π is said to be *strictly safe* if $V_1^\pi(\mu; c, P) < \tau$.

Assumption 2.2. *There exists a strictly safe policy π^0 with $V_1^{\pi^0}(\mu; c, P) = c^0 < \tau$.*

There are two important cases of the safe learning problem.

Zero constraint violation case: The agent has prior knowledge of a strictly safe policy π^0 and its safety cost value $c^0 := V_1^{\pi^0}(\mu; c, P)$. The agent wishes to attain a sublinear (in K) *cumulative regret*,

$$Reg(K; r) := \sum_{k=1}^K \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; r, P) \right), \quad (2)$$

while incurring zero constraint violation with at least a specified high probability $(1 - \delta)$, i.e.,

$$\mathbb{P} \left(V_1^{\pi^k}(\mu; c, P) \leq \tau, \forall k \in [K] \right) \geq 1 - \delta.$$

Bounded constraint violation case: The agent knows that there exists a strictly safe policy with a known safety cost value c^0 , but does not know any strictly safe policy. It aims to achieve a cumulative regret (2) that grows sublinearly with K , while ensuring that the *regret of constraint violation*,

$$Reg(K; c) := \left(\sum_{k=1}^K \left(V_1^{\pi^k}(\mu; c, P) - \tau \right) \right)_+ \quad (\text{where } (a)_+ := \max\{a, 0\}),$$

satisfies $\sup_K Reg(K; c) < +\infty$ with at least a specified high probability $(1 - \delta)$.

Remarks. For the zero constraint violation case, the assumption of knowing c^0 can be relaxed, as shown in Appendix E.

3 Zero constraint violation case

We start by considering the zero constraint violation case. On one hand, to balance the exploration-exploitation trade-off, we employ an optimistic estimate of the reward function, as embodied in the OFU principle. On the other hand, to maintain absolute safety with high probability during the exploration, we employ a pessimistic estimate of the safety cost. Such an OPFU principle was previously discussed in constrained bandits [3] and we adapt it to the unknown CMDP setting.

At each episode k , we begin by forming empirical estimates of the transition probabilities, the reward function, and the cost function from step h of all previous episodes:

$$\hat{P}_h^k(s'|s, a) := \frac{\sum_{k'=1}^{k-1} \mathbb{1}(S_h^{k'} = s, A_h^{k'} = a, S_{h+1}^{k'} = s')}{N_h^k(s, a) \vee 1}, \quad (3)$$

$$\hat{g}_h^k(s, a) := \frac{\sum_{k'=1}^{k-1} \mathbb{1}(S_h^{k'} = s, A_h^{k'} = a)(g_h(s, a) + \xi_h^k(s, a; g))}{N_h^k(s, a) \vee 1}, \quad \text{for } g = r, c, \quad (4)$$

where $a \vee b := \max\{a, b\}$, and

$$N_h^k(s, a) := \# \text{ of visits to state-action pair } (s, a) \text{ at step } h \text{ in episodes } [1, 2, \dots, k-1]. \quad (5)$$

Next we fix some $\delta \in (0, 1)$ and form a common (for notational simplicity) confidence radius $\beta_h^k(s, a)$ for the transition probabilities, the rewards, and the costs,

$$\beta_h^k(s, a) := \sqrt{\frac{1}{N_h^k(s, a) \vee 1} Z}, \quad \text{where } Z := \log(16|\mathcal{S}|^2|\mathcal{A}|HK/\delta).$$

At each episode k , we define the *optimistically biased reward estimate* as

$$\bar{r}_h^k(s, a) := \hat{r}_h^k(s, a) + \alpha_r \beta_h^k(s, a), \quad \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H], \quad (6)$$

where the scaling factor is

$$\alpha_r := 1 + |\mathcal{S}|H + \frac{4H(1 + |\mathcal{S}|H)}{\tau - c^0}. \quad (7)$$

To guarantee safe exploration, define the *pessimistically biased safety cost estimate* at episode k as

$$\underline{c}_h^k(s, a) := \hat{c}_h^k(s, a) + (1 + H|\mathcal{S}|) \beta_h^k(s, a), \quad \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]. \quad (8)$$

The policy we execute at episode k is chosen from a “pessimistically safe” policy set Π^k , defined as

$$\Pi^k := \begin{cases} \{\pi^0\} & \text{if } V_1^{\pi^0}(\mu; \underline{c}^k, \hat{P}^k) \geq (\tau + c^0)/2, \\ \{\pi : V_1^\pi(\mu; \underline{c}^k, \hat{P}^k) \leq \tau\} & \text{otherwise.} \end{cases} \quad (9)$$

We simply use the strictly safe policy π^0 until $V_1^{\pi^0}(\mu; \underline{c}^k, \hat{P}^k) < (\tau + c^0)/2$, which is a sufficient condition to guarantee that the set $\{\pi : V_1^\pi(\mu; \underline{c}^k, \hat{P}^k) \leq \tau\}$ is non-empty. Within this set, we choose the optimistically best reward earning policy π^k , which can be solved by linear programming with $\Theta(|\mathcal{S}||\mathcal{A}|H)$ decision variables and constraints [11]. The resulting Optimistic Pessimism-based Linear Programming (OptPess-LP) algorithm is presented below:

Algorithm 1: OptPess-LP

Input: $K, \delta, \pi^0, c^0, \tau$;

Initialization: $N_h^1(s, a) = 0, \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$;

for $k = 1, 2, \dots, K$ **do**

Update empirical model (i.e., $\hat{P}^k, \hat{r}^k, \hat{c}^k$) as in Equations (3)-(5);
 Update $\bar{r}^k, \underline{c}^k$, and Π^k as in Equations (6)-(9);
 Calculate $\pi^k \in \arg \max_{\pi \in \Pi^k} V_1^\pi(\mu; \bar{r}^k, \hat{P}^k)$;
 Execute π^k and collect a trajectory $(S_h^k, A_h^k, R_h^k, C_h^k), \forall h \in [H]$;
 Update counters $N_h^{k+1}(S_h^k, A_h^k), \forall h \in [H]$;

end

Theorem 3.1 (Regret and constraint violation bounds for OptPess-LP). *Fix any $\delta \in (0, 1)$. With probability at least $(1 - \delta)$, OptPess-LP has zero constraint violation with*

$$\text{Reg}^{\text{OPLP}}(K; r) = \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + \frac{H^5 |\mathcal{S}|^3 |\mathcal{A}|}{(\tau - c^0)^2 \wedge (\tau - c^0)} \right), \quad \text{where } a \wedge b := \min\{a, b\}.$$

Theorem 3.1 shows that it is possible to achieve sublinear regret in K , while simultaneously incurring no constraint violation with arbitrarily high probability. The proof is sketched in Section 5.1, with detailed proofs presented in Appendix B.

4 Bounded constraint violation case

Without prior knowledge of a strictly safe policy π^0 , we may not be able to guarantee zero constraint violation with high probability. However, by relaxing the requirement to bounded (in K) safety constraint violation, we can incorporate more exploration and design a more efficient algorithm by a primal-dual approach. It is inspired by the design of the pessimistic term in constrained bandits [21].

Different from traditional optimistic dual analysis methods in [11], we introduce an additive pessimistic term ϵ_k at each episode k in the original optimization problem (1), i.e.,

$$\max_{\pi} V_1^{\pi}(\mu; r, P) \quad \text{s.t.} \quad V_1^{\pi}(\mu; c, P) + \epsilon_k \leq \tau. \quad (10)$$

The pessimistic term restrains the constraint violation and will be progressively decreased as learning proceeds.

The CMDP problem in (10) may however not have any feasible solution. To overcome this, we consider the Lagrangian,

$$L^k(\pi, \lambda) := V_1^{\pi}(\mu; r, P) + \lambda(\tau - \epsilon_k - V_1^{\pi}(\mu; c, P)),$$

for which, given any Lagrange multiplier λ , we can always solve for the optimizer $\max_{\pi} L^k(\pi, \lambda)$ by dynamic programming.

With the introduction of the additive term ϵ_k , the dual variable (denoted by λ^k) governed by the subgradient algorithm grows faster, which enhances safety constraints in the next episode.

In order to guarantee sufficient optimism of rewards and costs, we integrate the uncertainty of transitions into rewards and costs, i.e.,

$$\begin{aligned} \tilde{r}_h^k(s, a) &:= \hat{r}_h^k(s, a) + \beta(n_k(s, a, h)) + H|\mathcal{S}|\beta(n_k(s, a, h)), \quad \forall(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H], \\ \tilde{c}_h^k(s, a) &:= \hat{c}_h^k(s, a) - \beta(n_k(s, a, h)) - H|\mathcal{S}|\beta(n_k(s, a, h)), \quad \forall(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]. \end{aligned} \quad (11)$$

Note that in contrast to the zero constraint violation case, we *optimistically* estimate the safety cost function, with the pessimism only governed by ϵ_k .

We employ truncated value functions due to the additional uncertainties from transitions, i.e.,

$$\hat{V}_1^{\pi}(\mu; \tilde{r}^k, \hat{P}^k) := \min\{H, V_1^{\pi}(\mu; \tilde{r}^k, \hat{P}^k)\}, \quad \hat{V}_1^{\pi}(\mu; \tilde{c}^k, \hat{P}^k) := \max\{0, V_1^{\pi}(\mu; \tilde{c}^k, \hat{P}^k)\}. \quad (12)$$

For the policy update of the primal variable (denoted by π^k), we can apply standard dynamic programming by viewing $\tilde{r}_h^k(s, a) - \frac{\lambda^k}{\eta^k}(\tilde{c}_h^k(s, a) - \tau)$ as the reward function. Specifically, we apply backward induction to solve for the optimal policy:

$$Q_h^k(s, a) = \tilde{r}_h^k(s, a) - \frac{\lambda^k}{\eta^k}(\tilde{c}_h^k(s, a) - \tau) + \sum_{s' \in \mathcal{S}} \hat{P}_h^k(s'|s, a) \max_{a'} Q_{h+1}^k(s', a'), \quad \forall h \in [H]$$

with $Q_{H+1}^k(s, a) = 0, \forall(s, a) \in \mathcal{S} \times \mathcal{A}$. Then, $\pi_h^k \in \arg \max_a Q_h^k(s, a)$, which is computationally efficient (as efficient as policy-gradient-based algorithms in the tabular case).

The resulting Optimistic Pessimism-based Primal-Dual (OptPess-PrimalDual) algorithm is shown in Algorithm 2. It chooses $\epsilon_k := 5H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|}(\log \frac{k}{\delta'} + 1)/\sqrt{k \log \frac{k}{\delta'}}$, $\delta' = \delta/(16|\mathcal{S}|^2|\mathcal{A}|H)$, the scaling parameter $\eta^k := (\tau - c^0)H\sqrt{k}$ in the primal policy update, and, for convenience, a step size of 1 for the dual update.

Theorem 4.1 (Regret and constraint violation bounds for OptPess-PrimalDual). *Fix any $\delta \in (0, 1)$. Then,*

$$\begin{aligned} \text{Reg}^{\text{OPPD}}(K; r) &= \tilde{\mathcal{O}}\left(\frac{H^3}{\tau - c^0}\sqrt{|\mathcal{S}|^3|\mathcal{A}|K} + \frac{H^5|\mathcal{S}|^3|\mathcal{A}|}{(\tau - c^0)^2}\right), \\ \text{Reg}^{\text{OPPD}}(K; c) &= \mathcal{O}\left(C''(H - \tau) + H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|C''}\right) = \mathcal{O}(1), \end{aligned}$$

where $C'' = \mathcal{O}\left(\frac{H^4|\mathcal{S}|^3|\mathcal{A}|}{(\tau - c^0)^2} \log \frac{H^4|\mathcal{S}|^3|\mathcal{A}|}{(\tau - c^0)^2\delta'}\right)$ is a coefficient independent of K .

Algorithm 2: OptPess-PrimalDual

Input: K, δ, c^0, τ ;
Initialization: $N_h^1(s, a) = 0, \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H], \lambda^1 = 0$;
for $k = 1, 2, \dots, K$ **do**

Set $\epsilon_k = 5H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|}(\log \frac{k}{\delta} + 1)/\sqrt{k \log \frac{k}{\delta}}, \delta' = \delta/(16|\mathcal{S}|^2|\mathcal{A}|H)$,
 $\eta^k = (\tau - c^0)H\sqrt{k}$;
Update empirical model (i.e., $\hat{P}^k, \hat{r}^k, \hat{c}^k$) as in Equations (3)-(5);
Update \tilde{r}^k, \tilde{c}^k as in Equation (11);
(*Policy Update*) $\pi^k \in \arg \max_{\pi \in \Pi} \hat{V}_1^\pi(\mu; \tilde{r}^k, \hat{P}^k) - \frac{\lambda^k}{\eta^k} (\hat{V}_1^\pi(\mu; \tilde{c}^k, \hat{P}^k) - \tau)$;
(*Dual Update*) $\lambda^{k+1} = \left(\lambda^k + \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau \right)_+$;
Execute π^k and collect a trajectory $(S_h^k, A_h^k, R_h^k, C_h^k), \forall h \in [H]$;
Update counters $N_h^{k+1}(S_h^k, A_h^k), \forall h \in [H]$;

end

Theorem 4.1 shows that it is possible to achieve an $\tilde{\mathcal{O}}(\sqrt{K})$ reward regret, while only allowing bounded constraint violation with arbitrarily high probability. Detailed proofs are presented in Section 5.2 and Appendix C.

5 Regret and constraint violation analysis

We sketch the key steps in the proofs of Theorems 3.1 and 4.1. Full details are relegated to Appendices B and C, respectively. Note that our analysis and results are conditioned on the same high probability event (specifically defined in Appendix A), which occurs with probability at least $(1 - \delta)$.

5.1 Analysis of OptPess-LP (Algorithm 1)

Constraint violation analysis The zero constraint violation of OptPess-LP follows from the following property of the pessimistic policy set Π^k :

Lemma 5.1. *With probability at least $(1 - \delta)$, for any $k \in [K]$ and policy $\pi \in \Pi^k, V_1^\pi(\mu; c, P) \leq \tau$.*

Regret of reward analysis When the parameters c, P are not well estimated, there may not exist any policy π such that $V_1^\pi(\mu; \underline{c}^k, \hat{P}^k) \leq \tau$. Hence, as defined in (9), Π^k is a singleton set $\{\pi^0\}$, and accordingly π^0 is executed, even though π^0 is not safe for $(\underline{c}^k, \hat{P}^k)$. It subsequently takes several episodes of exploration using π^0 until it becomes strictly safe for $(\underline{c}^k, \hat{P}^k)$. At that time, policies close enough to π^0 , of which there are infinitely many, are also safe for $(\underline{c}^k, \hat{P}^k)$, and so $|\Pi^k| = +\infty$. At this point the learning algorithm can proceed to enhance reward while maintaining safety with respect to $(\underline{c}^k, \hat{P}^k)$.

To analyze the algorithm, we decompose the reward regret as follows:

$$\begin{aligned} Reg^{\text{OPLP}}(K; r) &= \sum_{k=1}^K \mathbb{1}(|\Pi^k| = 1) \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^0}(\mu; r, P) \right) \\ &\quad + \sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) \right) \\ &\quad + \sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) \left(V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P) \right). \end{aligned} \quad (13)$$

To bound the first term on the right-hand side (RHS) of (13), we have the following lemma which gives an upper bound on the number of episodes for exploration by policy π^0 :

Lemma 5.2. *With probability at least $(1 - \delta)$, $\sum_{k=1}^K \mathbb{1}(|\Pi^k| = 1) \leq C'$, where $C' = \tilde{\mathcal{O}}(H^4|\mathcal{S}|^3|\mathcal{A}|/((\tau - c^0)^2 \wedge (\tau - c^0)))$.*

Turning to the second term on the RHS of (13), we first note that π^* may not be in Π^k . To ensure that the term $V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k)$ is nevertheless non-positive even when $\pi^* \notin \Pi^k$, we set α_r to the large value shown in (7). This increases \bar{r}^k , and hence also $V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k)$. In addition, we show that there is a policy $\tilde{\pi}^k$ that attains the same reward as the (non-Markov) probabilistic mixed policy, $\tilde{\pi}^k := B_{\gamma_k} \pi^* + (1 - B_{\gamma_k}) \pi^0$, where B_{γ_k} is a Bernoulli distributed random variable with mean γ_k for $\gamma_k \in [0, 1]$. γ_k will be chosen as the largest coefficient such that $V_1^{\tilde{\pi}^k}(\mu; \underline{c}^k, \hat{P}^k) \leq \tau$. This latter policy in turn has a larger reward than π^0 since it is a mixture with π^* , yielding the following:

Lemma 5.3. *With probability at least $(1 - \delta)$,*

$$\sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) \right) \leq 0.$$

Finally, concerning the third term on the RHS of (13), akin to the closed-loop identifiability property [6, 23], while \hat{P}^k may not converge to P , the difference in the rewards $V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P)$ grows sublinearly in k since the same policy π^k is used in both values:

Lemma 5.4. *With probability at least $(1 - \delta)$,*

$$\sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) \left(V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P) \right) = \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + \frac{H^5 |\mathcal{S}|^3 |\mathcal{A}|}{\tau - c^0} \right).$$

Combining Lemmas 5.2, 5.3, and 5.4 yields Theorem 3.1.

5.2 Analysis of OptPess-PrimalDual (Algorithm 2)

In this section, we outline the steps in the proof of Theorem 4.1 by analyzing regret and constraint violation of OptPess-PrimalDual respectively.

Recall $\epsilon_k = 5H^2 \sqrt{|\mathcal{S}|^3 |\mathcal{A}|} (\log \frac{k}{\delta'} + 1) / \sqrt{k \log \frac{k}{\delta'}}$, where $\delta' = \delta / (16|\mathcal{S}|^2 |\mathcal{A}| H)$. The existence of a feasible solution to (10) can be guaranteed if $\epsilon_k \leq \tau - c^0$. Let C'' be the smallest value such that $\forall k \geq C''$, $\epsilon_k \leq (\tau - c^0)/2$. Then the perturbed optimization problem (10) has at least one feasible solution for any $k \geq C''$. By simple calculation, one can verify that $C'' = \mathcal{O}(\frac{H^4 |\mathcal{S}|^3 |\mathcal{A}|}{(\tau - c^0)^2} \log \frac{H^4 |\mathcal{S}|^3 |\mathcal{A}|}{(\tau - c^0)^2 \delta'})$. Notice that ϵ_k is a function not depending on K , and so is the coefficient C'' .

Constraint violation analysis The bounded constraint violation of OptPess-PrimalDual is established as follows. We first decompose the constraint violation as

$$\begin{aligned} \text{Reg}_{\text{OPPD}}(K; c) &= \left(\sum_{k=1}^K \left(V_1^{\pi^k}(\mu; c, P) - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right) + \sum_{k=1}^K \left(\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) - \tau \right) \right)_+ \\ &\leq \left(\sum_{k=1}^K \left(V_1^{\pi^k}(\mu; c, P) - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right) + \lambda^{K+1} - \sum_{k=1}^K \epsilon_k \right)_+. \end{aligned} \quad (14)$$

The first summation term and λ^{K+1} in (14) can be bounded as follows:

Lemma 5.5. *Recall $\delta' = \delta / (16|\mathcal{S}|^2 |\mathcal{A}| H)$, with probability at least $(1 - \delta)$,*

$$\sum_{k=1}^K \left(V_1^{\pi^k}(\mu; c, P) - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right) \leq 8H^2 \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K \log \frac{K}{\delta'}} + \mathcal{O}(\text{PolyLog}(K)).$$

Lemma 5.6. For any $k \geq C''$, with probability at least $(1 - \delta)$,

$$\lambda^k \leq \frac{1}{\zeta} \ln \frac{11\nu_{\max}^2}{3\rho^2} + C''(H - \tau) + \sum_{u=1}^{C''} \epsilon_u + H + \frac{4(H^2 + \epsilon_k^2 + \eta^k H)}{\tau - c^0},$$

where $\rho = (\tau - c^0)/4$, $\nu_{\max} = H$, $\zeta = \rho/(\nu_{\max}^2 + \nu_{\max}\rho/3)$.

To guarantee bounded violation, we ensure that $\sum_{k=1}^K \epsilon_k$ in (14) can cancel the dominant terms in the two lemmas above. According to Lemmas 5.5 and 5.6, with probability at least $(1 - \delta)$, the violation is bounded as

$$Reg^{\text{OPPD}}(K; c) = \mathcal{O} \left(C''H + H^2 \sqrt{|\mathcal{S}|^3 |\mathcal{A}| C'' \log(C''/\delta')} \right).$$

Regret of reward analysis For episode k with $k \geq C''$, let $\pi^{\epsilon_k, *}$ be the optimal policy for (10), which is well-defined by the definition of C'' . We decompose the reward regret as

$$\begin{aligned} Reg^{\text{OPPD}}(K; r) &= \sum_{k=1}^{C''} \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; r, P) \right) \\ &+ \sum_{k=C''}^K \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) \right) + \sum_{k=C''}^K \left(V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) - \hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) \right) \\ &+ \sum_{k=C''}^K \left(\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) - \hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) \right) + \sum_{k=C''}^K \left(\hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P) \right). \end{aligned} \quad (15)$$

We upper bound each term on the RHS of (20). Since $V_1^{\pi}(\mu; r, P) \in [0, H]$ for any policy π , the first term is upper bounded by HC'' . The second and third terms can be bounded by the following two lemmas:

Lemma 5.7. With probability at least $(1 - \delta)$,

$$\sum_{k=C''}^K \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) \right) \leq \sum_{k=C''}^K \frac{\epsilon_k H}{\tau - c^0} = \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} \right).$$

Lemma 5.8. With probability at least $(1 - \delta)$, $\sum_{k=C''}^K \left(V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) - \hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) \right) \leq 0$.

The pivotal step is to leverage optimism of π^k to further decompose the fourth term on the RHS of (20), and utilize the projected dual update to transfer it into the form of $\lambda^k(\lambda^k - \lambda^{k+1})$. The following lemmas provide high probability bounds for the remaining two terms on the RHS of (20):

Lemma 5.9. With probability at least $(1 - \delta)$,

$$\sum_{k=C''}^K \left(\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) - \hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) \right) = \tilde{\mathcal{O}} \left(\frac{H}{\tau - c^0} \sqrt{K} \right).$$

Lemma 5.10. With probability at least $(1 - \delta)$,

$$\sum_{k=C''}^K \left(\hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P) \right) = \tilde{\mathcal{O}} \left(H^2 \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + H^4 |\mathcal{S}|^3 |\mathcal{A}| \right).$$

Applying Lemmas 5.7, 5.8, 5.9, and 5.10 yields Theorem 4.1.

6 Concluding remarks

We present two optimistic pessimism-based algorithms that maintain stringent safety constraints (either zero or bounded safety constraint violation) for unknown CMDPs with high probability, while still attaining an $\tilde{\mathcal{O}}(\sqrt{K})$ regret of reward. The algorithms employ, respectively, a pessimistically safe policy set Π^k or an additional pessimistic term ϵ_k into the safety constraint. The first algorithm, OptPess-LP, guarantees zero violation with high probability by solving a linear programming with $\Theta(|\mathcal{S}||\mathcal{A}|H)$ decision variables, while the second algorithm, OptPess-PrimalDual, is as efficient as policy-gradient-based algorithms in the tabular case, but violates constraints during initial episodes. A possible future direction for exploration is the application of the above OPFU principle in model-free algorithms.

Acknowledgement

P. R. Kumar's work is partially supported by US National Science Foundation under CMMI-2038625, HDR Tripods CCF-1934904; US Office of Naval Research under N00014-21-1-2385; US ARO under W911NF1810331, W911NF2120064; and U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Solar Energy Technologies Office Award Number DE-EE0009031. The views expressed herein and conclusions contained in this document are those of the authors and should not be interpreted as representing the views or official policies, either expressed or implied, of the U.S. NSF, ONR, ARO, Department of Energy or the United States Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

Dileep Kalathil gratefully acknowledges funding from the U.S. National Science Foundation (NSF) grants NSF-CRII- CPS-1850206 and NSF-CAREER-EPCN-2045783.

References

- [1] J. Achiam, D. Held, A. Tamar, and P. Abbeel. Constrained policy optimization. In *International Conference on Machine Learning*, pages 22–31. PMLR, 2017.
- [2] E. Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- [3] S. Amani. Linear stochastic bandits under safety constraints. *NeurIPS 2019*, 2019.
- [4] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [5] F. Berkenkamp, M. Turchetta, A. P. Schoellig, and A. Krause. Safe model-based reinforcement learning with stability guarantees. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 908–919, 2017.
- [6] V. Borkar and P. Varaiya. Adaptive control of Markov chains, I: Finite parameter set. *IEEE Transactions on Automatic Control*, 24(6):953–957, 1979.
- [7] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3387–3395, 2019.
- [8] C. Dann, T. Lattimore, and E. Brunskill. Unifying pac and regret: uniform pac bounds for episodic reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5717–5727, 2017.
- [9] D. Ding, X. Wei, Z. Yang, Z. Wang, and M. Jovanovic. Provably efficient safe exploration via primal-dual policy optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 3304–3312. PMLR, 2021.
- [10] G. Dulac-Arnold, D. Mankowitz, and T. Hester. Challenges of real-world reinforcement learning. *arXiv preprint arXiv:1904.12901*, 2019.

- [11] Y. Efroni, S. Mannor, and M. Pirotta. Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189*, 2020.
- [12] Y. Efroni, N. Merlis, M. Ghavamzadeh, and S. Mannor. Tight regret bounds for model-based reinforcement learning with greedy policies. *Advances in Neural Information Processing Systems*, 32:12224–12234, 2019.
- [13] E. Garcelon, M. Ghavamzadeh, A. Lazaric, and M. Pirotta. Improved algorithms for conservative exploration in bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3962–3969, 2020.
- [14] J. García and F. Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- [15] W. Hoeffding. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pages 409–426. Springer, 1994.
- [16] T. Jaksch, R. Ortner, and P. Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- [17] C. Jin, T. Jin, H. Luo, S. Sra, and T. Yu. Learning adversarial markov decision processes with bandit feedback and unknown transition. In *International Conference on Machine Learning*, pages 4860–4869. PMLR, 2020.
- [18] A. Kazerouni, M. Ghavamzadeh, Y. Abbasi-Yadkori, and B. Van Roy. Conservative contextual linear bandits. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 3913–3922, 2017.
- [19] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause. Learning-based model predictive control for safe exploration. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 6059–6066. IEEE, 2018.
- [20] Q. Liang, F. Que, and E. Modiano. Accelerated primal-dual policy optimization for safe reinforcement learning. *arXiv preprint arXiv:1802.06480*, 2018.
- [21] X. Liu, B. Li, P. Shi, and L. Ying. An efficient pessimistic-optimistic algorithm for constrained linear bandits. *arXiv preprint arXiv:2102.05295*, 2021.
- [22] A. Moradipari, S. Amani, M. Alizadeh, and C. Thrampoulidis. Safe linear thompson sampling with side information. *arXiv preprint arXiv:1911.02156*, 2019.
- [23] P. R. Kumar and A. Becker. A new family of optimal adaptive controllers for Markov chains. *IEEE Transactions on Automatic Control*, 27(1):137–146, 1982.
- [24] A. Pacchiano, M. Ghavamzadeh, P. Bartlett, and H. Jiang. Stochastic bandits with linear constraints. In *International Conference on Artificial Intelligence and Statistics*, pages 2827–2835. PMLR, 2021.
- [25] S. Paternain, M. Calvo-Fullana, L. F. Chamom, and A. Ribeiro. Safe policies for reinforcement learning via primal-dual methods. *arXiv preprint arXiv:1911.09101*, 2019.
- [26] S. Qiu, X. Wei, Z. Yang, J. Ye, and Z. Wang. Upper confidence primal-dual reinforcement learning for cmdp with adversarial loss. *arXiv e-prints*, pages arXiv–2003, 2020.
- [27] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- [28] A. Stooke, J. Achiam, and P. Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pages 9133–9143. PMLR, 2020.
- [29] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [30] C. Tessler, D. J. Mankowitz, and S. Mannor. Reward constrained policy optimization. In *International Conference on Learning Representations*, 2018.

- [31] A. Wachi, Y. Sui, Y. Yue, and M. Ono. Safe exploration and optimization of constrained mdps using gaussian processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [32] Y. Wu, R. Shariff, T. Lattimore, and C. Szepesvári. Conservative bandits. In *International Conference on Machine Learning*, pages 1254–1262. PMLR, 2016.
- [33] T.-Y. Yang, J. Rosca, K. Narasimhan, and P. J. Ramadge. Projection-based constrained policy optimization. In *International Conference on Learning Representations*, 2019.
- [34] L. Zheng and L. Ratliff. Constrained upper confidence reinforcement learning. In *Learning for Dynamics and Control*, pages 620–629. PMLR, 2020.

Appendix

The theorems and lemmas presented in the paper are provided with full details in this appendix.

First, let's recall some notations. Fix some $0 < \delta < 1$ as the input of algorithms. We say $g = \tilde{\mathcal{O}}(f)$ if there exists a universal constant C such that $g \leq C(f \log \frac{1}{\delta} + f \log f)$ for any f and δ . The filtration $\{\mathcal{F}_k\}_{k \geq 0}$ is defined as follows: \mathcal{F}_0 is the trivial sigma algebra, and for each $k \in [K]$, $\mathcal{F}_k = \sigma((S_h^{k'}, A_h^{k'}, R_h^{k'}, C_h^{k'})_{h \in [H], k' \in [k]})$. The policy process, i.e., the sequence of deployed policies $\{\pi^k\}_{k \in [K]}$, is a predictable process with respect to the filtration $\{\mathcal{F}_k\}_{k \geq 0}$. According to the definition of N_h^k , we know $N_h^k(s, a) \in \mathcal{F}_{k-1}$.

Additionally, we define expectation operator $\mathbb{E}_{\mu', P', \pi'}[\cdot]$ as the expectation with respect to a stochastic trajectory $(S_h, A_h)_{h \in [H]}$ generated according to the Markov chain induced by (μ', P', π') . When μ', P', π' are random elements, $\mathbb{E}_{\mu', P', \pi'}[\cdot]$ will be a $\sigma(\mu', P', \pi')$ -measurable random variable.

A High probability good event \mathcal{E}

We aim to give performance guarantees of the algorithms with high probability. To this end, we first consider a high probability "good event" \mathcal{E} that all the subsequent analysis is conditioned on.

First, for any predictable event sequence $\mathcal{G}_{1:K}$, i.e., $\mathcal{G}_k \in \mathcal{F}_{k-1}$, $\forall k \in [K]$, define the event

$$\begin{aligned} \mathcal{E}_{\mathcal{G}}(\delta) := & \left\{ \forall K' \in [K], \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h)}{N_h^k(s, a) \vee 1} \leq 4H|\mathcal{S}||\mathcal{A}| + 2H|\mathcal{S}||\mathcal{A}| \ln K'_{\mathcal{G}} + 4 \ln \frac{2HK}{\delta}, \right. \\ & \left. \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h)}{\sqrt{N_h^k(s, a) \vee 1}} \leq 6H|\mathcal{S}||\mathcal{A}| + 2H\sqrt{|\mathcal{S}||\mathcal{A}|K'_{\mathcal{G}}} + 2H|\mathcal{S}||\mathcal{A}| \ln K'_{\mathcal{G}} + 5 \ln \frac{2HK}{\delta} \right\}, \end{aligned}$$

where $K'_{\mathcal{G}} := \sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k)$ and q^{π^k} is the occupancy measure of policy π^k , i.e., $q^{\pi^k}(s, a, h) = \mathbb{E}_{\mu, P, \pi^k}[\mathbb{1}(S_h^k = s, A_h^k = a) | \mathcal{F}_{k-1}]$.

A trivial predictable event sequence is $\mathcal{G}_{1:k}$ with $\mathcal{G}_k = \Omega, \forall k \in [K]$, where Ω is the sample space. Let $\mathcal{E}_{\Omega}(\delta)$ be the event defined by this trivial event sequence.

Consider another event sequence $\mathcal{G}'_{1:K} = \{V_1^{\pi^0}(\mu; \underline{c}^k, \hat{P}^k) \geq \frac{\tau+c^0}{2}\}_{k \in [K]}$, which is predictable with respect to $\{\mathcal{F}_k\}_{k \geq 0}$. Let $\mathcal{E}_0(\delta)$ be the event $\mathcal{E}_{\mathcal{G}'}(\delta)$ defined by this event sequence $\mathcal{G}'_{1:K}$. Notice that all notations (including \hat{r} , \hat{c} , \hat{P} , and \underline{c}^k) are defined in the same manner in the two algorithms, so this event sequence $\mathcal{G}'_{1:K}$ is also well-defined in Algorithm 2.

Good event \mathcal{E} . Define a "good event" \mathcal{E} as

$$\begin{aligned} \mathcal{E} := & \left\{ \forall k \in [K], \forall h \in [H], \forall s \in \mathcal{S}, \forall a \in \mathcal{A}, \right. \\ & |r_h(s, a) - \hat{r}_h^k(s, a)| \leq \beta_h^k(s, a), |c_h(s, a) - \hat{c}_h^k(s, a)| \leq \beta_h^k(s, a), \\ & |\hat{P}_h^k(s'|s, a) - P_h(s'|s, a)| \leq \beta_h^k(s, a), \forall s' \in \mathcal{S} \\ & \left. |\hat{P}_h^k(s'|s, a) - P_h(s'|s, a)| \leq \tilde{\beta}_h^k(s'|s, a), \forall s' \in \mathcal{S} \right\} \cap \mathcal{E}_{\Omega}(\delta/4) \cap \mathcal{E}_0(\delta/4), \quad (16) \end{aligned}$$

where $\tilde{\beta}_h^k(s'|s, a) := \sqrt{\frac{2P(s'|s, a)}{N_h^k(s, a) \vee 1}} Z + \frac{Z}{3N_h^k(s, a) \vee 1}$ and $Z := \log(16|\mathcal{S}|^2|\mathcal{A}|HK/\delta)$.

Lemma A.1. Fix any $\delta \in (0, 1)$ as the confidence parameter in the inputs of the proposed algorithms. The good event \mathcal{E} occurs with probability at least $1 - \delta$.

Proof of Lemma A.1. For each $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$, we take K mutually independent samples of the reward, cost, and next state with the probability distribution specified by the generative model M :

$$\{R^n(s, a, h), C^n(s, a, h), S^n(s, a, h)\}_{n=1}^K,$$

Let $(\hat{r}^n(s, a, h), \hat{c}^n(s, a, h), \hat{P}^n(\cdot|s, a, h))$ be the corresponding (running) empirical means, respectively, for the samples

$$\{R^i(s, a, h), C^i(s, a, h), S^i(s, a, h)\}_{i=1}^n.$$

Define the following failure events:

$$\begin{aligned} F_n^r &:= \{\exists s, a, h : |\hat{r}^n(s, a, h) - r_h(s, a)| \geq \beta(n)\}, \\ F_n^c &:= \{\exists s, a, h : |\hat{c}^n(s, a, h) - c_h(s, a)| \geq \beta(n)\}, \\ F_n^P &:= \{\exists s, a, s', h : |P_h(s'|s, a) - \hat{P}^n(s'|s, a, h)| \geq \beta(n)\}, \\ \tilde{F}_n^P &:= \{\exists s, a, s', h : |P_h(s'|s, a) - \hat{P}^n(s'|s, a, h)| \geq \tilde{\beta}(n, P_h(s'|s, a))\}, \end{aligned}$$

where $\beta(n) := \sqrt{\frac{1}{n\sqrt{1}}Z}$ and $\tilde{\beta}(n, p) := \sqrt{\frac{2P(s'|s, a)}{n\sqrt{1}}Z} + \frac{Z}{3n\sqrt{1}}$.

Define the event \mathcal{E}^{gen}

$$\mathcal{E}^{gen} := \left(\left(\bigcup_{n=1}^K F_n^r \cup F_n^c \cup F_n^P \cup \tilde{F}_n^P \right)^C \cap \mathcal{E}_\Omega(\delta/4) \cap \mathcal{E}_0(\delta/4) \right).$$

By the definition of $\mathcal{E}_\Omega(\delta/4)$ and $\mathcal{E}_0(\delta/4)$ and Lemma D.5, event $\mathcal{E}_\Omega(\delta/4) \cap \mathcal{E}_0(\delta/4)$ occurs with probability at least $1 - \delta/2$. Therefore, to show that $\mathbb{P}(\mathcal{E}^{gen}) \geq 1 - \delta$, it is sufficient to show that $\mathbb{P}(\bigcup_{n=1}^K F_n^r \cup F_n^c \cup F_n^P \cup \tilde{F}_n^P) \leq \delta/2$. Note that $\delta/16 \leq \delta|\mathcal{S}|/4(1 + |\mathcal{S}|) =: \delta'$. Now, it is straightforward to show the following:

- (i) Using Hoeffding's inequality, $\mathbb{P}(\bigcup_{n=1}^K F_n^r) \leq |\mathcal{S}||\mathcal{A}|HK \frac{\delta}{16|\mathcal{S}|^2|\mathcal{A}|HK} \leq \frac{\delta'}{|\mathcal{S}|}$.
- (ii) Using Hoeffding's inequality, $\mathbb{P}(\bigcup_{n=1}^K F_n^c) \leq |\mathcal{S}||\mathcal{A}|HK \frac{\delta}{16|\mathcal{S}|^2|\mathcal{A}|HK} \leq \frac{\delta'}{|\mathcal{S}|}$.
- (iii) Using Hoeffding's inequality, $\mathbb{P}(\bigcup_{n=1}^K F_n^P) \leq |\mathcal{S}|^2|\mathcal{A}|HK \frac{\delta}{16|\mathcal{S}|^2|\mathcal{A}|HK} \leq \delta'$.
- (iv) Using Bernstein's inequality, $\mathbb{P}(\bigcup_{n=1}^K \tilde{F}_n^P) \leq |\mathcal{S}|^2|\mathcal{A}|HK \frac{\delta}{16|\mathcal{S}|^2|\mathcal{A}|HK} \leq \delta'$.

Using union bound, $\mathbb{P}(\bigcup_{n=1}^K F_n^r \cup F_n^c \cup F_n^P \cup \tilde{F}_n^P) \leq \mathbb{P}(\bigcup_{n=1}^K F_n^r) + \mathbb{P}(\bigcup_{n=1}^K F_n^c) + \mathbb{P}(\bigcup_{n=1}^K F_n^P) + \mathbb{P}(\bigcup_{n=1}^K \tilde{F}_n^P) \leq (2 + 2/|\mathcal{S}|)\delta' = \delta/2$.

The episodic CMDP problem studied here can be simulated as follows: at episode k , taking action a at state s and step h returns

$$\left(R^{n_k(s, a, h)}(s, a, h), C^{n_k(s, a, h)}(s, a, h), S^{n_k(s, a, h)}(s, a, h) \right).$$

Then all the samples drawn in the episodic CMDP problem are contained in

$$\{R^n(s, a, h), C^n(s, a, h), S^n(s, a, h)\}_{n=1}^K.$$

The sample averages calculated by the algorithm are

$$\left(\hat{r}_h^k(s, a), \hat{c}_h^k(s, a), \hat{P}_h^k(\cdot|s, a) \right) = \left(\hat{r}^{n_k(s, a, h)}(s, a, h), \hat{c}^{n_k(s, a, h)}(s, a, h), \hat{P}^{n_k(s, a, h)}(\cdot|s, a, h) \right).$$

Since \mathcal{E}^{gen} implies \mathcal{E} , \mathcal{E} holds with probability at least $1 - \delta$. \square

Remark. The analysis in the rest of this appendix is conditioned on the good event \mathcal{E} . That is, if (Ω, \mathcal{F}, P) is the underlying probability space, then we suppose throughout the rest of the appendix that the sample point $\omega \in \mathcal{E}$.

B Details of analysis for the zero constraint violation case

B.1 Constraint violation analysis

We will provide a proof of a slightly stronger result. Specifically, we will suppose that we know a policy π^0 and a constant $c^0 < \tau$ such that the safety cost $V_1^{\pi^0}(\mu; c, P) \leq c^0$. The strengthening lies

in the relaxed requirement of knowing only an *upper bound* on the safety cost of π^0 , rather than its *exact* value.

The zero constraint violation of OptPess-LP is guaranteed by the pessimistic policy set Π^k , as stated in the following lemma. Recall

$$\Pi^k := \begin{cases} \{\pi^0\} & \text{if } V_1^{\pi^0}(\mu; \underline{c}^k, \hat{P}^k) \geq (\tau + c^0)/2, \\ \{\pi : V_1^\pi(\mu; \underline{c}^k, \hat{P}^k) \leq \tau\} & \text{otherwise.} \end{cases}$$

Lemma B.1 (Restatement of Lemma 5.1). *On the good event \mathcal{E} , for any $k \in [K]$ and policy $\pi \in \Pi^k$, $V_1^\pi(\mu; c, P) \leq \tau$.*

Proof of Lemma 5.1. Fix $\omega \in \mathcal{E}$ throughput. Then, for any k, s, a, h , we have

$$\begin{aligned} (\underline{c}_h^k - c_h)(s, a) &= \underline{c}_h^k(s, a) - \hat{c}_h^k(s, a) + \hat{c}_h^k(s, a) - c_h(s, a) \\ &= (1 + |\mathcal{S}|H)\beta_h^k(s, a) + \hat{c}_h^k(s, a) - c_h(s, a) \\ &\geq (1 + |\mathcal{S}|H)\beta_h^k(s, a) - \beta_h^k(s, a) = |\mathcal{S}|H\beta_h^k(s, a), \end{aligned}$$

and

$$\sum_{s'} (\hat{P}_h^k - P_h)(s'|s, a) V_{h+1}^\pi(s'; c, P) \geq -H \sum_{s'} \beta_h^k(s, a) = -|\mathcal{S}|H\beta_h^k(s, a).$$

Thus, by Lemma D.2, for any policy π , we have

$$\begin{aligned} &V_1^\pi(\mu; \underline{c}^k, \hat{P}^k) - V_1^\pi(\mu; c, P) \\ &= \mathbb{E}_{\mu, \hat{P}^k, \pi} \left[\sum_{h=1}^H \left((\underline{c}_h^k - c_h)(S_h, A_h) + \sum_{s'} (\hat{P}_h^k - P_h)(s'|S_h, A_h) V_{h+1}^\pi(s'; c, P) \right) \middle| \mathcal{F}_{k-1} \right] \\ &\geq \mathbb{E}_{\mu, \hat{P}^k, \pi} \left[\sum_{h=1}^H |\mathcal{S}|H\beta_h^k(S_h, A_h) - |\mathcal{S}|H\beta_h^k(S_h, A_h) \middle| \mathcal{F}_{k-1} \right] \geq 0. \end{aligned} \tag{17}$$

Consider any $\pi \in \Pi^k$. If $\pi = \pi^0$, then $V_1^{\pi^0}(\mu; c, P) = c^0 < \tau$. Otherwise,

$$V_1^\pi(\mu; c, P) \leq V_1^\pi(\mu; \underline{c}^k, \hat{P}^k) \leq \tau.$$

□

B.2 Regret of reward analysis

When the parameters c, P are not well estimated, there may not exist any policy π such that $V_1^\pi(\mu; \underline{c}^k, \hat{P}^k) \leq \tau$. Hence, as defined in (B.1), Π^k is a singleton set $\{\pi^0\}$, and accordingly π^0 is executed, even though π^0 is not safe for $(\underline{c}^k, \hat{P}^k)$. It subsequently takes several episodes of exploration using π^0 until it becomes strictly safe for $(\underline{c}^k, \hat{P}^k)$, which means that $\pi^k = \pi^0$ and $V_1^{\pi^k}(\mu; \underline{c}^k, \hat{P}^k) \geq (\tau + c^0)/2$ when $|\Pi^k| = 1$. At that time, policies close enough to π^0 , of which there are infinitely many, are also safe for $(\underline{c}^k, \hat{P}^k)$, and so $|\Pi^k| = +\infty$. At this point the learning algorithm can proceed to enhance reward while maintaining safety with respect to $(\underline{c}^k, \hat{P}^k)$.

To analyze the algorithm, we decompose the reward regret as follows:

$$\begin{aligned} Reg(K; r) &= \underbrace{\sum_{k=1}^K \mathbb{1}(|\Pi^k| = 1) \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^0}(\mu; r, P) \right)}_{(I)} \\ &\quad + \underbrace{\sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; \hat{r}^k, \hat{P}^k) \right)}_{(II)} \\ &\quad + \underbrace{\sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) \left(V_1^{\pi^k}(\mu; \hat{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P) \right)}_{(III)}. \end{aligned}$$

We analyze these three terms, and the results are summarized in the following lemmas, respectively.

The following lemma gives an upper bound on the number of episodes for exploration by policy π^0 . It thereby bounds the first term on the right-hand side (RHS) of (13).

Lemma B.2 (Restatement of Lemma 5.2). *On the good event \mathcal{E} , $\sum_{k=1}^K \mathbb{1}(|\Pi^k| = 1) \leq C'$, where $C' = \tilde{\mathcal{O}}(H^4|\mathcal{S}|^3|\mathcal{A}|/((\tau - c^0)^2 \wedge (\tau - c^0)))$.*

Proof of Lemma 5.2. With $K' := \sum_{k=1}^K \mathbb{1}(|\Pi^k| = 1)$, we then have

$$\begin{aligned} \frac{(\tau - c^0)K'}{2} &= \sum_{k=1}^K \mathbb{1}(|\Pi^k| = 1) \frac{(\tau - c^0)}{2} \\ &= \sum_{k=1}^K \mathbb{1}(|\Pi^k| = 1) \left(\frac{\tau + c^0}{2} - c^0 \right) \\ &\stackrel{(a)}{\leq} \sum_{k=1}^K \mathbb{1}(V_1^{\pi^k}(\mu; \underline{c}^k, \hat{P}^k) \geq \frac{\tau + c^0}{2}) \left(V_1^{\pi^k}(\mu; \underline{c}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; c, P) \right) \\ &\stackrel{(b)}{=} \tilde{\mathcal{O}}\left(H^4|\mathcal{S}|^3|\mathcal{A}| + H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|K'}\right). \end{aligned}$$

(a) holds since $\pi^k = \pi^0$, $V_1^{\pi^k}(\mu; \underline{c}^k, \hat{P}^k) \geq (\tau + c^0)/2$, and $V_1^{\pi^k}(\mu; c, P) \leq c^0$ when $|\Pi^k| = 1$. The equality (b) follows from Lemma D.4 with $|\underline{c}_h^k - c_h| = |\hat{c}_h^k - c_h + (1 + H|\mathcal{S}|)\beta_h^k| \leq (2 + |\mathcal{S}|H)\beta_h^k$. Applying Lemma D.6 for K' , there exists some parameter C' that

$$K' \leq C' = \tilde{\mathcal{O}}\left(\frac{H^4|\mathcal{S}|^3|\mathcal{A}|}{(\tau - c^0)(1 \wedge (\tau - c^0))}\right).$$

□

The following lemma gives the resulting high probability bound for Term (II).

Lemma B.3 (Restatement of Lemma 5.3). *Recall $\alpha_r = 1 + |\mathcal{S}|H + 4H(1 + |\mathcal{S}|H)/(\tau - c^0)$. On the good event \mathcal{E} ,*

$$\sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) \right) \leq 0.$$

Proof of Lemma 5.3. Consider any $k \in [K]$ with $V_1^{\pi^0}(\mu; \underline{c}^k, \hat{P}^k) < (\tau + c^0)/2$. Then $|\Pi^k| > 1$. When $\pi^* \in \Pi^k$, it is straightforward to verify that Term (II) is less or equal to zero by optimism. When $\pi^* \notin \Pi^k$, define a probabilistic mixed policy

$$\tilde{\pi}^k = B_{\gamma_k} \pi^* + (1 - B_{\gamma_k}) \pi^0,$$

where B_{γ_k} is an independent Bernoulli distributed random variable with mean γ_k . Let $\gamma_k \in [0, 1]$ be the largest coefficient such that

$$V_1^{\tilde{\pi}^k}(\mu; \underline{c}^k, \hat{P}^k) \leq \tau. \quad (18)$$

If $V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) < \tau$, then $\gamma_k = 1$. Otherwise, at γ_k , equality holds in (18). Then we have

$$\begin{aligned} \tau &= \gamma_k V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) + (1 - \gamma_k) V_1^{\pi^0}(\mu; \underline{c}^k, \hat{P}^k) \\ &\leq \gamma_k V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) + (1 - \gamma_k) \frac{\tau + c^0}{2} \\ &= \gamma_k (V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) - V_1^{\pi^*}(\mu; c, P)) + \gamma_k V_1^{\pi^*}(\mu; c, P) + (1 - \gamma_k) \frac{\tau + c^0}{2} \\ &\leq \gamma_k \left(V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) - V_1^{\pi^*}(\mu; c, P) \right) + \gamma_k \tau + \frac{\tau + c^0}{2} - \gamma_k \frac{\tau + c^0}{2} \\ &= \gamma_k \left(V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) - V_1^{\pi^*}(\mu; c, P) + \frac{\tau - c^0}{2} \right) + \frac{\tau + c^0}{2}. \end{aligned}$$

From (17), $V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) - V_1^{\pi^*}(\mu; c, P) + (\tau - c^0)/2 > 0$, from which it follows that

$$\gamma_k \geq \frac{\tau - c^0}{\tau - c^0 + 2 \left(V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) - V_1^{\pi^*}(\mu; c, P) \right)}.$$

By Lemma D.2, for any policy π ,

$$\begin{aligned} V_1^\pi(\mu; \underline{c}^k, \hat{P}^k) - V_1^\pi(\mu; c, P) &= \mathbb{E}_{\mu, \hat{P}^k, \pi} \left[\sum_{h=1}^H (\underline{c}_h^k - c_h)(S_h, A_h) \mid \mathcal{F}_{k-1} \right] \\ &\quad + \mathbb{E}_{\mu, \hat{P}^k, \pi} \left[\sum_{h=1}^H \langle (\hat{P}_h^k - P_h)(\cdot | S_h, A_h), V_{h+1}^\pi(\cdot; c, P) \rangle \mid \mathcal{F}_{k-1} \right] \\ &\stackrel{(a)}{\leq} \mathbb{E}_{\mu, \hat{P}^k, \pi} \left[\sum_{h=1}^H 2(1 + |\mathcal{S}|H) \beta_h^k(S_h, A_h) \mid \mathcal{F}_{k-1} \right] \\ &= 2(1 + |\mathcal{S}|H) V_1^\pi(\mu; \beta^k, \hat{P}^k), \end{aligned}$$

where (a) holds since for any k, s, a, h , we have

$$\begin{aligned} (\underline{c}_h^k - c_h)(s, a) &= \underline{c}_h^k(s, a) - \hat{c}_h^k(s, a) + \hat{c}_h^k(s, a) - c_h(s, a) \\ &\leq (1 + |\mathcal{S}|H) \beta_h^k(s, a) + \beta_h^k(s, a) \\ &= (2 + |\mathcal{S}|H) \beta_h^k(s, a), \end{aligned}$$

and

$$\sum_{s'} (\hat{P}_h^k - P_h)(s' | s, a) V_{h+1}^\pi(s'; c, P) \leq H \sum_{s'} \beta_h^k(s, a) = |\mathcal{S}|H \beta_h^k(s, a).$$

Though policy $\tilde{\pi}^k$ is not a (randomized or non-randomized) Markov policy, we can find a randomized Markov policy $\hat{\pi}^k$ such that the occupation distributions of the state-action pair at any time of $\hat{\pi}^k$ coincides with $\tilde{\pi}^k$ under the transition probabilities \hat{P}^k . Therefore, as stated in by Lemma D.3, $V_1^{\hat{\pi}^k}(\mu; g, \hat{P}^k) = V_1^{\tilde{\pi}^k}(\mu; g, \hat{P}^k)$ for any g .

Since $\hat{\pi}^k \in \Pi^k$, by the definition of π^k , we then have

$$\begin{aligned} V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) &\geq V_1^{\hat{\pi}^k}(\mu; \bar{r}^k, \hat{P}^k) = V_1^{\tilde{\pi}^k}(\mu; \bar{r}^k, \hat{P}^k) \\ &= \gamma_k V_1^{\pi^*}(\mu; \bar{r}^k, \hat{P}^k) + (1 - \gamma_k) V_1^{\pi^0}(\mu; \bar{r}^k, \hat{P}^k) \\ &\geq \gamma_k V_1^{\pi^*}(\mu; \bar{r}^k, \hat{P}^k) \\ &= \frac{\tau - c^0}{\tau - c^0 + 2 \left(V_1^{\pi^*}(\mu; \underline{c}^k, \hat{P}^k) - V_1^{\pi^*}(\mu; c, P) \right)} V_1^{\pi^*}(\mu; \bar{r}^k, \hat{P}^k) \\ &\geq \frac{\tau - c^0}{\tau - c^0 + 4(1 + |\mathcal{S}|H) V_1^{\pi^*}(\mu; \beta^k, \hat{P}^k)} V_1^{\pi^*}(\mu; \bar{r}^k, \hat{P}^k). \end{aligned}$$

To make $V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) \geq V_1^{\pi^*}(\mu; r, P)$, it suffices that

$$\frac{\tau - c^0}{\tau - c^0 + 4(1 + |\mathcal{S}|H) V_1^{\pi^*}(\mu; \beta^k, \hat{P}^k)} V_1^{\pi^*}(\mu; \bar{r}^k, \hat{P}^k) \geq V_1^{\pi^*}(\mu; r, P),$$

which follows since

$$(\tau - c^0)(V_1^{\pi^*}(\mu; \bar{r}^k, \hat{P}^k) - V_1^{\pi^*}(\mu; r, P)) \geq 4(1 + |\mathcal{S}|H) V_1^{\pi^*}(\mu; \beta^k, \hat{P}^k) V_1^{\pi^*}(\mu; r, P).$$

Now notice that from Lemma D.2,

$$\begin{aligned}
& V_1^{\pi^*}(\mu; \bar{r}^k, \hat{P}^k) - V_1^{\pi^*}(\mu; r, P) \\
&= \mathbb{E}_{\mu, \hat{P}^k, \pi^*} \left[\sum_{h=1}^H \left((\bar{r}_h^k - r_h)(S_h, A_h) + \sum_{s'} (\hat{P}_h^k - P_h)(s' | S_h, A_h) V_{h+1}^{\pi^*}(s'; r, P) \right) \mid \mathcal{F}_{k-1} \right] \\
&\geq \mathbb{E}_{\mu, \hat{P}^k, \pi^*} \left[\sum_{h=1}^H (\alpha_r - 1 - H|\mathcal{S}|) \beta_h^k(S_h, A_h) \mid \mathcal{F}_{k-1} \right] \\
&= (\alpha_r - 1 - H|\mathcal{S}|) V_1^{\pi^*}(\mu; \beta^k, \hat{P}^k).
\end{aligned}$$

Thus, the choice of $\alpha_r = 4H(1 + |\mathcal{S}|H)/(\tau - c^0) + 1 + |\mathcal{S}|H$ suffices to guarantee that $V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) \geq V_1^{\pi^*}(\mu; r, P)$ for any k with $|\Pi^k| > 1$. \square

Finally, Term (III) is bounded as follows by employing the value of α_r .

Lemma B.4 (Restatement of Lemma 5.4). *On the good event \mathcal{E} ,*

$$\sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) (V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P)) = \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + \frac{H^5 |\mathcal{S}|^3 |\mathcal{A}|}{\tau - c^0} \right).$$

Proof of Lemma 5.4. Since $|\bar{r}_h^k - r_h| = |\hat{r}_h^k - r_h + \alpha_r \beta_h^k| \leq (1 + \alpha_r) \beta_h^k$, by Lemma D.4, we have

$$\begin{aligned}
& \sum_{k=1}^K \mathbb{1}(|\Pi^k| > 1) (V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P)) \\
&\leq \sum_{k=1}^K |V_1^{\pi^k}(\mu; \bar{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P)| \\
&= \tilde{\mathcal{O}} \left((\alpha_r + H\sqrt{|\mathcal{S}|}) H \sqrt{|\mathcal{S}| |\mathcal{A}| K} + H^3 |\mathcal{S}|^2 |\mathcal{A}| (\alpha_r + |\mathcal{S}|) \right) \\
&= \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + \frac{H^5 |\mathcal{S}|^3 |\mathcal{A}|}{\tau - c^0} \right).
\end{aligned}$$

\square

B.3 Proof of Theorem 3.1

Theorem B.5 (Regret and constraint violation bounds for OptPess-LP (Restatement of Theorem 3.1)). *On the good event \mathcal{E} , OptPess-LP has zero constraint violation with*

$$\text{Reg}^{\text{OPLP}}(K; r) = \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + \frac{H^5 |\mathcal{S}|^3 |\mathcal{A}|}{(\tau - c^0)^2 \wedge (\tau - c^0)} \right), \text{ where } a \wedge b := \min\{a, b\}.$$

Proof of Theorem 3.1. The zero constraint violation of OptPess-LP is guaranteed by Lemma 5.1. Combining Lemmas 5.2, 5.3, and 5.4, which give upper bounds to Terms (I)-(III) respectively, we have

$$\sum_{k=1}^K (V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; r, P)) = \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + \frac{H^5 |\mathcal{S}|^3 |\mathcal{A}|}{(\tau - c^0) \wedge (\tau - c^0)^2} \right).$$

\square

C Details of analysis for bounded constraint violation case

Recall $\epsilon_k = 5H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|}(\log \frac{k}{\delta'} + 1)/\sqrt{k \log \frac{k}{\delta'}}$, where $\delta' = \delta/(16|\mathcal{S}|^2|\mathcal{A}|H)$. The existence of a feasible solution to (10) can be guaranteed if $\epsilon_k \leq \tau - c^0$. Let C'' be the smallest value such that $\forall k \geq C''$, $\epsilon_k \leq (\tau - c^0)/2$. Then the perturbed optimization problem (10) has at least one feasible solution for any $k \geq C''$. By simple calculation, one can verify that $C'' = \mathcal{O}(\frac{H^4|\mathcal{S}|^3|\mathcal{A}|}{(\tau - c^0)^2} \log \frac{H^4|\mathcal{S}|^3|\mathcal{A}|}{(\tau - c^0)^2\delta'})$. Notice that since ϵ_k is a function not depending on K , so is the coefficient C'' .

Lemma C.1 (Optimistic of cost value). *On the good event \mathcal{E} , for any $C'' \leq k \leq K$ and any policy π ,*

$$\hat{V}_1^\pi(\mu; \tilde{c}^k, \hat{P}^k) \leq V_1^\pi(\mu; c, P).$$

Proof of Lemma C.1. If $\hat{V}_1^\pi(\mu; \tilde{c}^k, \hat{P}^k) = V_1^\pi(\mu; \tilde{c}^k, \hat{P}^k)$, then

$$\begin{aligned} & \sum_{k=C''}^K \left(\hat{V}_1^\pi(\mu; \tilde{c}^k, \hat{P}^k) - V_1^\pi(\mu; c, P) \right) \\ &= \sum_{k=C''}^K \mathbb{E}_{\mu, \hat{P}^k, \pi} \left[\sum_{h=1}^H (\tilde{c}_h^k(S_h, A_h) - c_h(S_h, A_h)) + \langle (\hat{P}_h^k - P_h)(\cdot | S_h, A_h), V_{h+1}^\pi(\cdot; c, P) \rangle \mid \mathcal{F}_{k-1} \right] \\ &\leq \sum_{k=C''}^K \mathbb{E}_{\mu, \hat{P}^k, \pi} \left[\sum_{h=1}^H (\tilde{c}_h^k(S_h, A_h) - c_h(S_h, A_h)) - \beta_h^k - H|\mathcal{S}|\beta_h^k + H|\mathcal{S}|\beta_h^k \mid \mathcal{F}_{k-1} \right] \leq 0. \end{aligned}$$

Otherwise, $\hat{V}_1^\pi(\mu; \tilde{c}^k, \hat{P}^k) = 0 \leq V_1^\pi(\mu; c, P)$. \square

C.1 Constraint violation analysis

Lemma C.2 (Restatement of Lemma 5.5). *Recall $\delta' = \delta/(16|\mathcal{S}|^2|\mathcal{A}|H)$. On the good event \mathcal{E} ,*

$$\sum_{k=1}^K \left(V_1^{\pi^k}(\mu; c, P) - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right) \leq 8H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|K \log \frac{K}{\delta'}} + \tilde{\mathcal{O}}(H^4|\mathcal{S}|^3|\mathcal{A}|),$$

where $\tilde{\mathcal{O}}(H^4|\mathcal{S}|^3|\mathcal{A}|)$ contains factors poly-logarithmic in K .

Proof of Lemma 5.5. For any k, s, a, h , we have

$$\begin{aligned} |(\tilde{c}_h^k - c_h)(s, a)| &= |\tilde{c}_h^k(s, a) - \hat{c}_h^k(s, a) + \hat{c}_h^k(s, a) - c_h(s, a)| \\ &\leq (1 + |\mathcal{S}|H)\beta_h^k(s, a) + |\hat{c}_h^k(s, a) - c_h(s, a)| \\ &\leq (1 + |\mathcal{S}|H)\beta_h^k(s, a) + \beta_h^k(s, a) = (2 + |\mathcal{S}|H)\beta_h^k(s, a). \end{aligned}$$

By Lemma D.4, we have

$$\begin{aligned} & \sum_{k=1}^K \left(V_1^{\pi^k}(\mu; c, P) - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right) \\ &\leq (3(2 + |\mathcal{S}|H) + 3\sqrt{2H}\sqrt{|\mathcal{S}|})H\sqrt{|\mathcal{S}||\mathcal{A}|KZ} + \tilde{\mathcal{O}}(H^4|\mathcal{S}|^3|\mathcal{A}|) \\ &= (6 + 3|\mathcal{S}|H + 3H\sqrt{2|\mathcal{S}|})H\sqrt{|\mathcal{S}||\mathcal{A}|KZ} + \tilde{\mathcal{O}}(H^4|\mathcal{S}|^3|\mathcal{A}|) \\ &\stackrel{(a)}{\leq} 8H^2|\mathcal{S}|\sqrt{|\mathcal{S}||\mathcal{A}|KZ} + \tilde{\mathcal{O}}(H^4|\mathcal{S}|^3|\mathcal{A}|), \end{aligned}$$

where $\tilde{\mathcal{O}}(H^4|\mathcal{S}|^3|\mathcal{A}|)$ contains factors poly-logarithmic in K . (a) holds since we assume that $|\mathcal{S}| \geq 2$ and $H \geq 2$. \square

Lemma C.3 (Restatement of Lemma 5.6). *On the good event \mathcal{E} , for any $k \geq C''$,*

$$\lambda^k \leq \frac{1}{\zeta} \ln \frac{11\nu_{\max}^2}{3\rho^2} + C''(H - \tau) + \sum_{u=1}^{C''} \epsilon_u + H + \frac{4(H^2 + \epsilon_k^2 + \eta^k H)}{\tau - c^0},$$

where $\rho := (\tau - c^0)/4$, $\nu_{\max} := H$, $\zeta := \rho/(\nu_{\max}^2 + \nu_{\max}\rho/3)$.

Proof of Lemma 5.6. We will first verify conditions (i) and (ii) of Lemma D.7 for $k \geq C''$ by defining $\Phi(k) = \lambda^k$. Suppose $\Phi(k) = \lambda^k \geq \varphi_k := 4(H^2 + \epsilon_k^2 + \eta^k H)/(\tau - c^0)$. It is straightforward to verify that $\lambda^k + \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau$ is positive, which implies

$$\lambda^{k+1} = \lambda^k + \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau.$$

Let $L(k) = (\lambda^k)^2/2$, we have

$$\begin{aligned} L(k+1) - L(k) &= \frac{1}{2}(\lambda^{k+1})^2 - \frac{1}{2}(\lambda^k)^2 \\ &= \lambda^k(\lambda^{k+1} - \lambda^k) + \frac{1}{2}(\lambda^{k+1} - \lambda^k)^2 \\ &= \lambda^k \left(\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau \right) + \frac{1}{2}(\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau)^2 \\ &= \left(\lambda^k \left(\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau \right) - \eta^k \hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) \right) \\ &\quad + \eta^k \hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) + \frac{1}{2}(\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau)^2 \\ &\stackrel{(a)}{\leq} \left(\lambda^k \left(\hat{V}_1^{\pi^0}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau \right) - \eta^k \hat{V}_1^{\pi^0}(\mu; \tilde{r}, \hat{P}^k) \right) \\ &\quad + \eta^k \hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) + (\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) - \tau)^2 + \epsilon_k^2 \\ &\stackrel{(b)}{\leq} \lambda^k \left(V_1^{\pi^0}(\mu; c, P) + \epsilon_k - \tau \right) + \eta^k H + H^2 + \epsilon_k^2 \\ &\leq -\frac{\tau - c^0}{2} \lambda^k + H^2 + \epsilon_k^2 + \eta^k H. \end{aligned}$$

(a) is by the optimality of π^k and $(a+b)^2/2 \leq (a^2 + b^2)$. (b) holds since $\hat{V}_1^{\pi^0}(\mu; \tilde{c}^k, \hat{P}^k) \leq V_1^{\pi^0}(\mu; c, P)$ according to Lemma C.1.

It then follows that $L(k+1) < L(k)$ and

$$\begin{aligned} \lambda^{k+1} - \lambda^k &= \frac{(\lambda^{k+1})^2 - (\lambda^k)^2}{\lambda^{k+1} + \lambda^k} \\ &\leq \frac{(\lambda^{k+1})^2 - (\lambda^k)^2}{2\lambda^k} \\ &\leq -\frac{\tau - c^0}{2} + \frac{H^2 + \epsilon_k^2 + \eta^k H}{\varphi_k} = -\frac{\tau - c^0}{4} =: -\rho. \end{aligned}$$

In addition, since $\forall k \geq C''$, $\epsilon_k \leq (\tau - c^0)/2$,

$$|\lambda^{k+1} - \lambda^k| \leq |\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau| \leq H =: \nu_{\max}.$$

With the definitions of $\varphi_k, \rho, \nu_{\max}$, we apply Lemma D.7 starting from C'' , which gives

$$\begin{aligned} \lambda^k &\leq \frac{1}{\zeta} \ln \left(e^{\zeta \lambda^{C''}} + \frac{2e^{\zeta(v_{\max} + \varphi_k)}}{\zeta \rho} \right) \\ &\leq \frac{1}{\zeta} \ln \left(\frac{11v_{\max}^2 e^{\zeta(\lambda^{C''} + v_{\max} + \varphi_k)}}{3\rho^2} \right) \\ &= \frac{1}{\zeta} \ln \frac{11\nu_{\max}^2}{3\rho^2} + \lambda^{C''} + v_{\max} + \varphi_k \\ &\leq \frac{1}{\zeta} \ln \frac{11\nu_{\max}^2}{3\rho^2} + C''(H - \tau) + \sum_{u=1}^{C''} \epsilon_u + H + \frac{4(H^2 + \epsilon_k^2 + \eta^k H)}{\tau - c^0}, \end{aligned}$$

where $\zeta = \rho/(\nu_{\max}^2 + \nu_{\max}\rho/3) \geq 3(\tau - c^0)/(13H^2)$, and the last inequality is by

$$\begin{aligned} \lambda^{C''} &\leq \lambda^1 + \sum_{k=1}^{C''-1} (\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau)_+ \\ &\leq \sum_{k=1}^{C''} \epsilon_k + C''(H - \tau). \end{aligned} \tag{19}$$

□

The bounded constraint violation of OptPess-PrimalDual is established as follows. Notice that

$$\begin{aligned} \lambda^{K+1} &= \left(\lambda^K + \hat{V}_1^{\pi^K}(\mu; \tilde{c}^K, \hat{P}^K) + \epsilon_K - \tau \right)_+ \\ &\geq \epsilon_K + \hat{V}_1^{\pi^K}(\mu; \tilde{c}^K, \hat{P}^K) - \tau + \lambda^K \\ &\geq \epsilon_K + \epsilon_{K-1} + \hat{V}_1^{\pi^K}(\mu; \tilde{c}^K, \hat{P}^K) - \tau + \hat{V}_1^{\pi^{K-1}}(\mu; \tilde{c}^{K-1}, \hat{P}^{K-1}) - \tau + \lambda^{K-1} \\ &\geq \dots \\ &\geq \sum_{k=1}^K \epsilon_k + \sum_{k=1}^K \left(\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) - \tau \right) + \lambda^1. \end{aligned}$$

Since $\lambda^1 = 0$ according to the initial condition of Algorithm 2, we can decompose the constraint violation:

$$\begin{aligned} \text{Reg}^{\text{OPPD}}(K; c) &= \left(\sum_{k=1}^K \left(V_1^{\pi^k}(\mu; c, P) - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right) + \sum_{k=1}^K \left(\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) - \tau \right) \right)_+ \\ &\leq \left(\sum_{k=1}^K \left(V_1^{\pi^k}(\mu; c, P) - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right) + \lambda^{K+1} - \sum_{k=1}^K \epsilon_k \right)_+. \end{aligned}$$

The first summation term can be bounded by Lemma 5.5. The second λ^{K+1} term can be bounded by Lemma 5.6. Finally, recall $\epsilon_k = 5H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|}(\log \frac{k}{\delta'} + 1)/\sqrt{k \log \frac{k}{\delta'}}$, we have

$$\begin{aligned} \sum_{k=1}^K \epsilon_k &\geq \int_1^{K+1} \epsilon_u du \\ &\geq 10H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|K \log \frac{K}{\delta'}} - 10H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}| \log \frac{1}{\delta'}}. \end{aligned}$$

The negative $\sum_{k=1}^K \epsilon_k$ term can cancel the dominant terms. Thus, with probability at least $(1 - \delta)$, the violation is bounded as

$$\text{Reg}^{\text{OPPD}}(K; c) = \mathcal{O}\left(C''H + H^2\sqrt{|\mathcal{S}|^3|\mathcal{A}|C'' \log(C''/\delta')}\right).$$

C.2 Regret of reward analysis

For episode k with $k \geq C''$, let $\pi^{\epsilon_k, *}$ be the optimal policy for (10), which is well-defined by the definition of C'' . We decompose the reward regret as

$$\begin{aligned} \text{Reg}^{\text{OPPD}}(K; r) &= \sum_{k=1}^{C''} \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^k}(\mu; r, P) \right) \\ &\quad + \sum_{k=C''}^K \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) \right) + \sum_{k=C''}^K \left(V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) - \hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) \right) \\ &\quad + \sum_{k=C''}^K \left(\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) - \hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) \right) + \sum_{k=C''}^K \left(\hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P) \right). \end{aligned} \tag{20}$$

We upper bound each term in the RHS of (20). Since $V_1^\pi(\mu; r, P) \in [0, H]$ for any policy π , the first term is upper bounded by HC'' . The second and third terms can be bounded by the following two lemmas.

Lemma C.4 (Restatement of Lemma 5.7). *On the good event \mathcal{E} ,*

$$\sum_{k=C''}^K \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) \right) \leq \sum_{k=C''}^K \frac{\epsilon_k H}{\tau - c^0} = \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} \right).$$

Proof of Lemma 5.7. Define a probabilistic mixed policy

$$\pi^{\epsilon_k} = (1 - B_k)\pi^* + B_k\pi^0, \quad \forall k \geq C'',$$

where B_k is a Bernoulli distributed random variable with mean $\epsilon_k/(\tau - c^0)$.

We have that

$$\begin{aligned} V_1^{\pi^{\epsilon_k}}(\mu; c, P) &= (1 - \frac{\epsilon_k}{\tau - c^0})V_1^{\pi^*}(\mu; c, P) + \frac{\epsilon_k}{\tau - c^0}V_1^{\pi^0}(\mu; c, P) \\ &\leq (1 - \frac{\epsilon_k}{\tau - c^0})\tau + \frac{\epsilon_k}{\tau - c^0}c^0 = \tau - \epsilon_k. \end{aligned}$$

Though π^{ϵ_k} is not a Markov policy, by Lemma D.3, there exists a Markov policy $\hat{\pi}^{\epsilon_k}$, which has the same performance as π^{ϵ_k} under transition probabilities P . $\hat{\pi}^{\epsilon_k}$ is a feasible solution to the following problem

$$\begin{aligned} \max_{\pi} \quad & V_1^\pi(\mu; r, P) \\ \text{s.t.} \quad & V_1^\pi(\mu; c, P) + \epsilon_k \leq \tau. \end{aligned}$$

We then have

$$\begin{aligned} \sum_{k=C''}^K V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) &\leq \sum_{k=C''}^K V_1^{\pi^*}(\mu; r, P) - V_1^{\hat{\pi}^{\epsilon_k}}(\mu; r, P) \\ &= \sum_{k=C''}^K V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^{\epsilon_k}}(\mu; r, P) \\ &\leq \sum_{k=C''}^K \frac{\epsilon_k}{\tau - c^0} \left(V_1^{\pi^*}(\mu; r, P) - V_1^{\pi^0}(\mu; r, P) \right) \\ &\leq \sum_{k=C''}^K \frac{\epsilon_k H}{\tau - c^0}. \end{aligned}$$

□

Lemma C.5 (Restatement of Lemma 5.8). *On the good event \mathcal{E} ,*

$$\sum_{k=C''}^K \left(V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) - \hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) \right) \leq 0.$$

Proof of Lemma 5.8. For any $k \in [C'', K]$, $\pi^{\epsilon_k, *}$ is well-defined. If $\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) = V_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k)$, then by the value difference lemma (Lemma D.2), we know

$$\begin{aligned} & V_1^{\pi^{\epsilon_k, *}}(\mu; r, P) - \hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) \\ &= \mathbb{E}_{\mu, \hat{P}^k, \hat{\pi}^{\epsilon_k, *}} \left[\sum_{h=1}^H \left(r_h(S_h, A_h) - \tilde{r}_h^k(S_h, A_h) + \sum_{s'} (P_h - \hat{P}_h^k)(s' | S_h, A_h) V_{h+1}^\pi(s'; r, P) \right) \mid \mathcal{F}_{k-1} \right] \\ &\leq \sum_{k=C''}^K \mathbb{E}_{\mu, \hat{P}^k, \hat{\pi}^{\epsilon_k, *}} \left[\sum_{h=1}^H (r_h(S_h, A_h) - \hat{r}_h^k(S_h, A_h)) - \beta_h^k - H|\mathcal{S}|\beta_h^k + H|\mathcal{S}|\beta_h^k \mid \mathcal{F}_{k-1} \right] \leq 0. \end{aligned}$$

Otherwise, $\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) = H \geq V_1^{\pi^{\epsilon_k, *}}(\mu; r, P)$. Thus $\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) \geq V_1^{\pi^{\epsilon_k, *}}(\mu; r, P)$, for any $k \in [C'', K]$, and the lemma is proved. □

The pivotal step is to leverage optimism of π^k to further decompose the fourth term on the RHS of (20) and utilize the projected dual update to transfer it into the form of $\lambda^k(\lambda^k - \lambda^{k+1})$. The following lemmas provide high probability bounds for the remaining two terms in the RHS of (20):

Lemma C.6 (Restatement of Lemma 5.9). *On the good event \mathcal{E} ,*

$$\sum_{k=C''}^K \left(\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) - \hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) \right) = \tilde{\mathcal{O}}\left(\frac{H}{\tau - c^0} \sqrt{K}\right).$$

Proof of Lemma 5.9. Taking advantage of the optimism of π^k and the definition of the dual update, we have the following decomposition:

$$\begin{aligned} \sum_{k=C''}^K \left(\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) - \hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) \right) &= \underbrace{\sum_{k=C''}^K \left(\frac{\lambda^k}{\eta^k} (\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{c}^k, \hat{P}^k) - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k)) \right)}_{\text{(I)}} \\ &+ \underbrace{\sum_{k=C''}^K \left(\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{r}^k, \hat{P}^k) - \frac{\lambda^k}{\eta^k} \hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; \tilde{c}^k, \hat{P}^k) \right) - \left(\hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) - \frac{\lambda^k}{\eta^k} \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right)}_{\text{(II)}} \end{aligned}$$

Owing to the optimism of π^k , we know (II) ≤ 0 . Term (I) can then be upper bounded as follows:

$$\begin{aligned} \text{(I)} &\stackrel{(a)}{\leq} \sum_{k=C''}^K \frac{\lambda^k}{\eta^k} \left(\tau - \epsilon_k - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) \right) \\ &\stackrel{(b)}{\leq} \sum_{k=C''}^K \frac{1}{\eta^k} (\lambda^k(\lambda^k - \lambda^{k+1}) + \tau^2) \\ &\stackrel{(c)}{=} \sum_{k=C''}^K \frac{1}{\eta^k} \left(\frac{1}{2}(\lambda^k)^2 - \frac{1}{2}(\lambda^{k+1})^2 \right) + \sum_{k=C''}^K \frac{1}{2\eta^k} (\lambda^{k+1} - \lambda^k)^2 + \sum_{k=C''}^K \frac{1}{\eta^k} \tau^2 \\ &\stackrel{(d)}{\leq} \frac{(\lambda^{C''})^2}{2\eta^{C''}} + \sum_{k=C''}^K \frac{H^2}{2\eta^k} + \tau^2 \sum_{k=C''}^K \frac{1}{\eta^k} \\ &\stackrel{(e)}{\leq} \frac{(\sum_{k=1}^{C''} \epsilon_k + C''(H - \tau))^2}{2\eta^{C''}} + \sum_{k=C''}^K \frac{H^2}{2\eta^k} + \tau^2 \sum_{k=C''}^K \frac{1}{\eta^k}. \end{aligned}$$

(a) holds by applying Lemma C.1 for $\pi = \pi^{\epsilon_k, *}$ and $\hat{V}_1^{\pi^{\epsilon_k, *}}(\mu; c, P) \leq \tau - \epsilon_k$.

(b) holds as follows: if $\lambda^{k+1} > 0$, we know

$$\tau - \epsilon_k - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) = \lambda^k - \lambda^{k+1}.$$

If $\lambda^{k+1} = 0$, we have

$$\lambda^k \leq \tau - \epsilon_k - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) < \tau,$$

which indicates that $\lambda^k(\tau - \epsilon_k - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k)) \leq \tau^2$. Thus,

$$\lambda^k(\tau - \epsilon_k - \hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k)) \leq \lambda^k(\lambda^k - \lambda^{k+1}) + \tau^2.$$

Then, (c) holds since $-\epsilon\Delta = \frac{1}{2}\epsilon^2 - \frac{1}{2}(\epsilon + \Delta)^2 + \frac{1}{2}\Delta^2$, where $\epsilon = \lambda^k, \Delta = \lambda^{k+1} - \lambda^k$, (d) holds since η^k is monotonically increasing and $(\lambda^k - \lambda^{k+1})^2 \leq (\hat{V}_1^{\pi^k}(\mu; \tilde{c}^k, \hat{P}^k) + \epsilon_k - \tau)^2 \leq H^2$, and (e) holds according to (19). \square

Lemma C.7 (Restatement of Lemma 5.10). *On the good event \mathcal{E} ,*

$$\sum_{k=C''}^K \left(\hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P) \right) = \tilde{\mathcal{O}}(H^2 \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + H^4 |\mathcal{S}|^3 |\mathcal{A}|).$$

Proof of Lemma 5.10. For any k, s, a, h , we have

$$\begin{aligned} |(\tilde{r}_h^k - r_h)(s, a)| &= |\tilde{r}_h^k(s, a) - \hat{r}_h^k(s, a) + \hat{r}_h^k(s, a) - r_h(s, a)| \\ &\leq (1 + |\mathcal{S}|H)\beta_h^k(s, a) + |\hat{r}_h^k(s, a) - r_h(s, a)| \\ &\leq (1 + |\mathcal{S}|H)\beta_h^k(s, a) + \beta_h^k(s, a) = (2 + |\mathcal{S}|H)\beta_h^k(s, a). \end{aligned}$$

By Lemma D.4, we have

$$\sum_{k=C''}^K \left(\hat{V}_1^{\pi^k}(\mu; \tilde{r}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; r, P) \right) \leq \tilde{\mathcal{O}}(H^2 \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + H^4 |\mathcal{S}|^3 |\mathcal{A}|).$$

□

C.3 Proof of Theorem 4.1

Theorem C.8 (Regret and constraint violation bounds for OptPess-PrimalDual (Restatement of Theorem 4.1)). *On the good event \mathcal{E} ,*

$$\begin{aligned} \text{Reg}^{\text{OPPD}}(K; r) &= \tilde{\mathcal{O}} \left(\frac{H^3}{\tau - c^0} \sqrt{|\mathcal{S}|^3 |\mathcal{A}| K} + \frac{H^5 |\mathcal{S}|^3 |\mathcal{A}|}{(\tau - c^0)^2} \right), \\ \text{Reg}^{\text{OPPD}}(K; c) &= \mathcal{O} \left(C''(H - \tau) + H^2 \sqrt{|\mathcal{S}|^3 |\mathcal{A}| C''} \right) = \mathcal{O}(1), \end{aligned}$$

where $C'' = \mathcal{O} \left(\frac{H^4 |\mathcal{S}|^3 |\mathcal{A}|}{(\tau - c^0)^2} \log \frac{H^4 |\mathcal{S}|^3 |\mathcal{A}|}{(\tau - c^0)^2 \delta'} \right)$ is a coefficient independent of K .

Proof of Theorem 4.1. The bounded constraint analysis has been provided in the previous subsection. Applying Lemmas 5.7, 5.8, 5.9, and 5.10 yields the regret bound in Theorem 4.1. □

D Other supporting lemmas

Lemma D.1 (Hoeffding's inequality [15]). *For independent zero-mean 1/2-sub-Gaussian random variables X_1, X_2, \dots, X_n ,*

$$\begin{aligned} \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n X_i \geq \epsilon \right) &\leq \exp(-n\epsilon^2), \\ \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n X_i \leq -\epsilon \right) &\leq \exp(-n\epsilon^2). \end{aligned}$$

Lemma D.2 (Value difference lemma, [8, Lemma E.15]).

$$\begin{aligned} &V_1^\pi(\mu; g', P') - V_1^\pi(\mu; g, P) \\ &= \mathbb{E}_{\mu, P, \pi} \left[\sum_{h=1}^H \left(g'(S_h, A_h) - g(S_h, A_h) + \sum_{s'} (P'_h - P_h)(s' | S_h, A_h) V_{h+1}^\pi(s'; g', P') \right) \mid \mathcal{F}_{k-1} \right] \\ &= \mathbb{E}_{\mu, P', \pi} \left[\sum_{h=1}^H \left(g'(S_h, A_h) - g(S_h, A_h) + \sum_{s'} (P'_h - P_h)(s' | S_h, A_h) V_{h+1}^\pi(s'; g, P) \right) \mid \mathcal{F}_{k-1} \right], \end{aligned}$$

where $g = g' = r, c$.

Lemma D.3 ([2, Theorem 6.1(i)]). *Suppose the transition function is P . For any mixed policy $\tilde{\pi} = B_\gamma \pi^1 + (1 - B_\gamma) \pi^2$, where B_γ is a Bernoulli distributed random variable with mean γ . Then there exists a Markov policy $\hat{\pi}$ that*

$$V_h^{\hat{\pi}}(s; r, P) = V_h^{\tilde{\pi}}(s; r, P), \quad \forall r, s, h.$$

Lemma D.4. Let $\mathcal{G}_{1:K}$ be a sequence of events such that $\mathcal{G}_k \in \mathcal{F}_{k-1}$ for each $k \in [K]$. Suppose $|\tilde{g}^k - g| \leq \alpha\beta^k$, $\alpha \geq 1$. On the good event \mathcal{E} , for any $K' \leq K$,

$$\begin{aligned} \sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k) |V_1^{\pi^k}(\mu; \tilde{g}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; g, P)| &\leq (3\alpha + 3\sqrt{2}H\sqrt{|\mathcal{S}|})H\sqrt{|\mathcal{S}||\mathcal{A}|K'_G Z} \\ &\quad + \tilde{\mathcal{O}}(\alpha H^3 |\mathcal{S}|^2 |\mathcal{A}|), \end{aligned}$$

where $\mathcal{G}_{1:K} = (\Omega)_{1:K}$ or $\mathcal{G}_{1:k} = \{V_1^{\pi^0}(\mu; \underline{g}^k, \hat{P}^k) \geq \frac{\tau+c^0}{2}\}_{k \in [K]}$, and $K'_G = \sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k)$.

Proof of Lemma D.4. It can be proved following steps similar to those in the proof of Lemma 32 in [11], but utilizing Lemma D.5 for taking account of the predictable event sequence $\mathcal{G}_{1:K}$, and Lemma D.6 for bounding the leading order explicitly. For completeness, we include the detailed proof here.

$$\begin{aligned} &\sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k) \left| V_1^{\pi^k}(\mu; \tilde{g}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; g, P) \right| \\ &= \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \left| \tilde{g}_h^k(s, a) - g_h(s, a) + \sum_{s'} (\hat{P}_h^k - P_h)(s'|s, a) V_{h+1}^{\pi^k}(s'; \tilde{g}^k, \hat{P}^k) \right| \\ &\stackrel{(a)}{\leq} \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \left[\alpha\beta_h^k(s, a) + \sqrt{\sum_{s'} 2P(s'|s, a)(\beta_h^k(s, a))^2} \sqrt{|\mathcal{S}|H^2} + \frac{ZH|\mathcal{S}|}{N_h^k(s, a) \vee 1} \right] \\ &\quad + \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \left| \sum_{s'} (\hat{P}_h^k - P_h)(s'|s, a) \left(V_{h+1}^{\pi^k}(s'; \tilde{g}^k, \hat{P}^k) - V_{h+1}^{\pi^k}(s'; g, P) \right) \right| \\ &= \sum_{k=1}^K \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \left[(\alpha + \sqrt{2|\mathcal{S}|}H) \sqrt{\frac{Z}{N_h^k(s, a) \vee 1}} + \frac{ZH|\mathcal{S}|}{N_h^k(s, a) \vee 1} \right] \\ &\quad + \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \left| \sum_{s'} (\hat{P}_h^k - P_h)(s'|s, a) \left(V_{h+1}^{\pi^k}(s'; \tilde{g}^k, \hat{P}^k) - V_{h+1}^{\pi^k}(s'; g, P) \right) \right|. \end{aligned}$$

Since

$$\begin{aligned} &\sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \left| \sum_{s'} (\hat{P}_h^k - P_h)(s'|s, a) \left(V_{h+1}^{\pi^k}(s'; \tilde{g}^k, \hat{P}^k) - V_{h+1}^{\pi^k}(s'; g, P) \right) \right| \\ &\stackrel{(b)}{\leq} \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \sum_{s'} \frac{\sqrt{2P_h(s'|s, a)Z}}{\sqrt{N_h^k(s, a) \vee 1}} \left| V_{h+1}^{\pi^k}(s'; \tilde{g}^k, \hat{P}^k) - V_{h+1}^{\pi^k}(s'; g, P) \right| \\ &\quad + \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \frac{(1+\alpha)|\mathcal{S}|HZ}{N_h^k(s, a) \vee 1} \\ &\stackrel{(c)}{\leq} \tilde{\mathcal{O}}(\alpha H^2 |\mathcal{S}|^2 |\mathcal{A}|) + |\mathcal{S}| \sqrt{\alpha |\mathcal{A}| H^3} \sqrt{\sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k) \left| V_1^{\pi^k}(\mu; \tilde{g}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; g, P) \right|} + |\mathcal{S}| \sqrt{\alpha |\mathcal{A}| H^3} \\ &\quad \times \sqrt{\sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \mathbb{1}(\mathcal{G}_k) q^{\pi^k}(s, a, h) \left| \sum_{s'} (\hat{P}_h^k - P_h)(s'|s, a) \left(V_{h+1}^{\pi^k}(s'; \tilde{g}^k, \hat{P}^k) - V_{h+1}^{\pi^k}(s'; g, P) \right) \right|} \\ &\stackrel{(d)}{\leq} \tilde{\mathcal{O}}(\alpha H^3 |\mathcal{S}|^2 |\mathcal{A}|) + \frac{3}{2} |\mathcal{S}| H \sqrt{\alpha |\mathcal{A}| H} \sqrt{\sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k) \left| V_1^{\pi^k}(\mu; \tilde{g}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; g, P) \right|}, \end{aligned}$$

we have

$$\begin{aligned}
& \sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k) \left| V_1^{\pi^k}(\mu; \tilde{g}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; g, P) \right| \\
& \stackrel{(e)}{\leq} (2\alpha + 2\sqrt{2}H\sqrt{|\mathcal{S}|})H\sqrt{|\mathcal{S}||\mathcal{A}|K'_G Z} + \tilde{\mathcal{O}}(H|\mathcal{S}||\mathcal{A}|(\alpha + |\mathcal{S}|H)) + \tilde{\mathcal{O}}(\alpha H^3|\mathcal{S}|^2|\mathcal{A}|) \\
& \quad + \frac{3}{2}|\mathcal{S}|H\sqrt{\alpha|\mathcal{A}|H} \sqrt{\sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k) \left| V_1^{\pi^k}(\mu; \tilde{g}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; g, P) \right|} \\
& \stackrel{(f)}{\leq} (3\alpha + 3\sqrt{2}H\sqrt{|\mathcal{S}|})H\sqrt{|\mathcal{S}||\mathcal{A}|K'_G Z} + \tilde{\mathcal{O}}(\alpha H^3|\mathcal{S}|^2|\mathcal{A}|).
\end{aligned}$$

(a) and (b) hold due to the triangle inequality, the Cauchy-Schwarz inequality, and since $|V_1^{\pi^k}(\mu; \tilde{g}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; g, P)| \leq |V_1^{\pi^k}(\mu; \tilde{g}^k, \hat{P}^k) - V_1^{\pi^k}(\mu; g, \hat{P}^k)| + |V_1^{\pi^k}(\mu; g, \hat{P}^k) - V_1^{\pi^k}(\mu; g, P)| \leq (\alpha + 1)H$. (c) follows by steps similar to those in Lemma 32 in [12]. (d) and (f) hold due to Lemma D.6. (e) holds due to Lemma D.5. \square

Lemma D.5 ([17, Lemma 10] with a predictable events sequence). *Given a sequence of events $\mathcal{G}_{1:K}$ that $\mathcal{G}_k \in \mathcal{F}_{k-1}$ for each $k \in [K]$. With probability at least $1 - \delta$, for any $K' \leq K$*

$$\begin{aligned}
& \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)q^{\pi^k}(s,a,h)}{N_h^k(s,a) \vee 1} \leq 4H|\mathcal{S}||\mathcal{A}| + 2H|\mathcal{S}||\mathcal{A}| \ln K'_G + 4 \ln \frac{2HK}{\delta}, \\
& \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)q^{\pi^k}(s,a,h)}{\sqrt{N_h^k(s,a) \vee 1}} \leq 6H|\mathcal{S}||\mathcal{A}| + 2H\sqrt{|\mathcal{S}||\mathcal{A}|K'_G} + 2H|\mathcal{S}||\mathcal{A}| \ln K'_G + 5 \ln \frac{2HK}{\delta},
\end{aligned}$$

where $K'_G = \sum_{k=1}^{K'} \mathbb{1}(\mathcal{G}_k)$ and q^{π^k} is the occupancy measure of policy π^k , i.e., $q^{\pi^k}(s,a,h) = \mathbb{E}_{\mu, P, \pi^k}[\mathbb{1}(S_h^k = s, A_h^k = a) | \mathcal{F}_{k-1}]$.

Proof of Lemma D.5. This lemma can be proved following steps similar to those in the proof of Lemma 10 in [17]. For completeness, we include the detailed proof here.

Let $\mathbb{1}_k(s, a, h)$ be the indicator of whether the pair (s, a) is visited at step h in episode k , so that $\mathbb{E}_k[\mathbb{1}_k(s, a, h)] = q^{\pi^k}(s, a, h)$. We decompose the first quantity as

$$\begin{aligned}
& \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)q^{\pi^k}(s,a,h)}{N_h^k(s,a) \vee 1} = \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)\mathbb{1}_k(s,a,h)}{N_h^k(s,a) \vee 1} \\
& \quad + \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)(q^{\pi^k}(s,a,h) - \mathbb{1}_k(s,a,h))}{N_h^k(s,a) \vee 1}.
\end{aligned}$$

The first term can be bounded as

$$\sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)\mathbb{1}_k(s,a,h)}{N_h^k(s,a) \vee 1} \leq 2H|\mathcal{S}||\mathcal{A}| + H|\mathcal{S}||\mathcal{A}| \ln K'_G.$$

To bound the second term, we apply Lemma 9 in [17] with $Y_k = \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)(q^{\pi^k}(s,a,h) - \mathbb{1}_k(s,a,h))}{N_h^k(s,a) \vee 1}$, $\lambda = 1/2$, and the fact

$$\begin{aligned}
\mathbb{E}_k[Y_k^2] & \leq \mathbb{E}_k \left[\left(\sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)\mathbb{1}_k(s,a,h)}{N_h^k(s,a) \vee 1} \right)^2 \right] \\
& = \mathbb{E}_k \left[\sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)\mathbb{1}_k(s,a,h)}{(N_h^k(s,a))^2 \vee 1} \right] \\
& \leq \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)q^{\pi^k}(s,a,h)}{N_h^k(s,a) \vee 1},
\end{aligned}$$

which show that with probability at least $1 - \delta/(2HK)$,

$$\sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)(q^{\pi^k}(s,a,h) - \mathbb{1}_k(s,a,h))}{N_h^k(s,a) \vee 1} \leq \frac{1}{2} \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)q^{\pi^k}(s,a,h)}{N_h^k(s,a) \vee 1} + 2 \ln \frac{2HK}{\delta}.$$

Combining these two bounds, rearranging, and applying a union bound over k proves the first part of the lemma.

Similarly, we decompose the second quantity as

$$\begin{aligned} \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)q^{\pi^k}(s,a,h)}{\sqrt{N_h^k(s,a) \vee 1}} &= \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)\mathbb{1}_k(s,a,h)}{\sqrt{N_h^k(s,a) \vee 1}} \\ &\quad + \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)(q^{\pi^k}(s,a,h) - \mathbb{1}_k(s,a,h))}{\sqrt{N_h^k(s,a) \vee 1}}. \end{aligned}$$

The first term is bounded as

$$\sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)\mathbb{1}_k(s,a,h)}{\sqrt{N_h^k(s,a) \vee 1}} \leq 2H|\mathcal{S}||\mathcal{A}| + 2H\sqrt{|\mathcal{S}||\mathcal{A}|K'_G}.$$

To bound the second term, we again apply Lemma 9 in [17] with $Y_k = \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)(q^{\pi^k}(s,a,h) - \mathbb{1}_k(s,a,h))}{\sqrt{N_h^k(s,a) \vee 1}} \leq 1, \lambda = 1$, and the fact

$$\mathbb{E}_k[Y_k^2] \leq \mathbb{E}_k \left[\left(\sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)\mathbb{1}_k(s,a,h)}{\sqrt{N_h^k(s,a) \vee 1}} \right)^2 \right] = \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)q^{\pi^k}(s,a,h)}{N_h^k(s,a) \vee 1},$$

which shows that with probability at least $1 - \delta/(2HK)$,

$$\sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)(q^{\pi^k}(s,a,h) - \mathbb{1}_k(s,a,h))}{\sqrt{N_h^k(s,a) \vee 1}} \leq \sum_{k=1}^{K'} \sum_{h=1}^H \sum_{s,a} \frac{\mathbb{1}(\mathcal{G}_k)q^{\pi^k}(s,a,h)}{N_h^k(s,a) \vee 1} + \ln \frac{2HK}{\delta}.$$

Combining the first part of the lemma and employing a union bound proves the second part of the lemma. \square

Lemma D.6. Suppose $0 \leq x \leq a + b\sqrt{x}$, for some $a, b > 0$,

$$x \leq \frac{3}{2}a + \frac{3}{2}b^2.$$

Proof of Lemma D.6. By solving the equation of $x_0 = a + b\sqrt{x_0}$, we know

$$\begin{aligned} x \leq x_0 &= \frac{1}{2}(\sqrt{4ab^2 + b^4} + 2a + b^2) \\ &\leq \sqrt{ab^2} + \frac{b^2}{2} + a + \frac{b^2}{2} \\ &\leq \frac{3}{2}a + \frac{3}{2}b^2. \end{aligned}$$

\square

Lemma D.7 ([21, Lemma 11]). *Let $S(k)$ be a random process, $\Phi(k)$ be its Lyapunov function with $\Phi(0) = \Phi_0$ and $\Delta(k) = \Phi(k+1) - \Phi(k)$ be the Lyapunov drift. Given an increasing sequence $\{\varphi_t\}$, ρ and ν_{\max} with $0 < \rho \leq \nu_{\max}$, if the expected drift $\mathbb{E}[\Delta(k)|S(k) = s]$ satisfies the following conditions:*

- (i) *There exists constants $\rho > 0$ and $\varphi_t > 0$ s.t. $\Delta(k) \leq -\rho$ when $\Phi(k) \geq \varphi_k$, and*
- (ii) *$|\Phi(k+1) - \Phi(k)| \leq \nu_{\max}$ holds with probability one;*

then we have

$$e^{\zeta\Phi(k)} \leq e^{\zeta\Phi_0} + \frac{2e^{\zeta(\nu_{\max} + \varphi_k)}}{\zeta\rho},$$

where $\zeta = \rho/(\nu_{\max}^2 + \nu_{\max}\rho/3)$.

E Extensions (unknown c^0 with prior knowledge of π^0)

It may be noticed that the OptPess-LP algorithm can guarantee zero violation with high probability if a coefficient $c^{0'}$ with $V_1^{\pi^0}(\mu; c, P) \leq c^{0'} < \tau$ is used to replace c^0 in the inputs. Thus when the value $V_1^{\pi^0}(\mu; c, P)$ is not known, one may first estimate an appropriate upper bound of $V_1^{\pi^0}(\mu; c, P)$, and call OptPess-LP using this upper bound. The details of choosing the upper bound are as follows.

We can execute policy π^0 sequentially until an estimate of c^0 with enough precision is obtained. Denote $\hat{c}^0(k)$ as the mean estimate after executing policy π^0 k times. We stop the estimation process after K'' executions, where

$$K'' := \min \left\{ k \geq 1 : \tau - \hat{c}^0(k) \geq 3\sqrt{\frac{1}{kH} \log \frac{2K}{\delta''}} \right\} \wedge K.$$

By Hoeffding's inequality, with probability at least $1 - \delta''$,

$$|\hat{c}^0(k) - c^0| \leq \sqrt{\frac{1}{kH} \log \frac{2K}{\delta''}}.$$

Thus with probability at least $1 - \delta''$, for $k = K'' - 1$,

$$\tau - 4\sqrt{\frac{1}{kH} \log \frac{2K}{\delta''}} < \hat{c}^0(k) - \sqrt{\frac{1}{kH} \log \frac{2K}{\delta''}} \leq c^0,$$

which implies

$$K'' = k + 1 \leq 1 + \frac{16 \log(2K/\delta'')}{H(\tau - c^0)^2}.$$

Then $K'' < K$ can then be guaranteed if $\frac{16 \log(2K/\delta'')}{H(\tau - c^0)^2} < K - 1$. After K'' executions, with probability at least $1 - \delta''$,

$$\begin{aligned} c^{0'} &= \hat{c}^0(K'') + \sqrt{\frac{1}{K''H} \log \frac{2K}{\delta''}} \\ &= \frac{1}{2} \left(\hat{c}^0(K'') - \sqrt{\frac{1}{K''H} \log \frac{2K}{\delta''}} \right) + \frac{1}{2} \left(\hat{c}^0(K'') + 3\sqrt{\frac{1}{K''H} \log \frac{2K}{\delta''}} \right) \\ &\leq \frac{1}{2}c^0 + \frac{1}{2}\tau = \frac{1}{2}(c^0 + \tau). \end{aligned}$$

We can apply the OptPess-LP algorithm with a pessimistic estimate $\left(\hat{c}^0(K'') + \sqrt{\frac{1}{K''H} \log \frac{2K}{\delta''}} \right)$ of c^0 . It also maintains zero constraint violation with the same order of regret since $1/(\tau - c^{0'}) \leq 1/(\tau - \frac{1}{2}(c^0 + \tau)) = 2/(\tau - c^0)$.