

Policy-based Primal-Dual Methods for Concave CMDP with Variance Reduction

DONGHAO YING*, University of California, Berkeley, USA

MENGZI AMY GUO, University of California, Berkeley, USA

HYUNIN LEE, University of California, Berkeley, USA

YUHAO DING, Cubist Systematic Strategies, USA

JAVAD LAVAEI, University of California, Berkeley, USA

ZUO-JUN MAX SHEN, University of California, Berkeley, USA

We study Concave Constrained Markov Decision Processes (Concave CMDPs) where both the objective and constraints are defined as concave functions of the state-action occupancy measure. We propose the Variance-Reduced Primal-Dual Policy Gradient Algorithm (VR-PDPG), which updates the primal variable via policy gradient ascent and the dual variable via projected sub-gradient descent. Despite the challenges posed by the loss of additivity structure and the nonconcave nature of the problem, we establish the global convergence of VR-PDPG by exploiting a form of hidden concavity. In the exact setting, we prove an $O(T^{-1/3})$ convergence rate for both the average optimality gap and constraint violation, which further improves to $O(T^{-1/2})$ under strong concavity of the objective in the occupancy measure. In the sample-based setting, we demonstrate that VR-PDPG achieves an $\tilde{O}(\epsilon^{-4})$ sample complexity for ϵ -global optimality. Moreover, by incorporating a diminishing pessimistic term into the constraint, we show that VR-PDPG can attain a zero constraint violation without compromising the convergence rate of the optimality gap. Finally, we validate our methods through numerical experiments.

JAIR Associate Editor: Marek Petrik

JAIR Reference Format:

Donghao Ying, Mengzi Amy Guo, Hyunin Lee, Yuhao Ding, Javad Lavaei, and Zuo-Jun Max Shen. 2025. Policy-based Primal-Dual Methods for Concave CMDP with Variance Reduction. *Journal of Artificial Intelligence Research* 83, Article 31 (August 2025), 53 pages. DOI: [10.1613/jair.1.18129](https://doi.org/10.1613/jair.1.18129)

1 Introduction

Reinforcement Learning (RL), typically formulated by Markov Decision Processes (MDPs), focuses on learning a policy that maps situations to actions to maximize the expected cumulative sum of rewards. Mathematically, this objective can be reformulated as maximizing the inner product between the policy-induced state-action occupancy measure and a policy-independent reward for each state-action pair. Nonetheless, many decision-making problems of practical interest take a form beyond the cumulative reward. For instance, in apprenticeship learning, the objective shifts towards minimizing the distance between the occupancy measures of an expert and

*Corresponding Author.

Authors' Contact Information: Donghao Ying, ORCID: [0009-0001-7329-5917](https://orcid.org/0009-0001-7329-5917), donghaoy@berkeley.edu, University of California, Berkeley, Berkeley, CA, USA; Mengzi Amy Guo, ORCID: [0000-0001-8869-4673](https://orcid.org/0000-0001-8869-4673), mengzi_guo@berkeley.edu, University of California, Berkeley, Berkeley, CA, USA; Hyunin Lee, ORCID: [0000-0003-1718-9594](https://orcid.org/0000-0003-1718-9594), hyunin@berkeley.edu, University of California, Berkeley, Berkeley, CA, USA; Yuhao Ding, yuhao.ding3@gmail.com, Cubist Systematic Strategies, New York City, NY, USA; Javad Lavaei, ORCID: [0000-0003-4294-1338](https://orcid.org/0000-0003-4294-1338), lavaei@berkeley.edu, University of California, Berkeley, Berkeley, CA, USA; Zuo-Jun Max Shen, ORCID: [0000-0003-4538-8312](https://orcid.org/0000-0003-4538-8312), maxshen@berkeley.edu, University of California, Berkeley, Berkeley, CA, USA.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2025 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.18129](https://doi.org/10.1613/jair.1.18129)

Table 1. We summarize our results for concave CMDPs in the following table, where the concavity is in terms of the occupancy measure and T denotes the total number of iterations. The number of samples required by VR-PDPG is $\mathcal{O}(\log T)$ in each iteration, thereby implying the overall sample complexity of $\tilde{\mathcal{O}}(\epsilon^{-4})$ to achieve an $\mathcal{O}(\epsilon)$ optimality gap and constraint violation. For all cases in the table, a zero constraint violation can be achieved while maintaining the same order of convergence rate for the optimality gap.

Setting	Algorithm	Objective	Optimality Gap	Constraint Violation
Exact	PDPG	Concave	$\mathcal{O}(T^{-1/3})$	$\mathcal{O}(T^{-1/3})$
		Strongly Concave	$\mathcal{O}(T^{-1/2})$	$\mathcal{O}(T^{-1/2})$
Sample-based	VR-PDPG	Concave	$\mathcal{O}(T^{-1/4})$	$\mathcal{O}(T^{-1/4})$

the learning agent [1]. In the pursuit of diverse skill discovery, the agent seeks to maximize the Kullback–Leibler (KL) divergence between the state occupancy measures that correspond to different skills [19]. Additionally, in pure exploration problems, the goal becomes maximizing the entropy/or a KL-divergence of the state occupancy measure [23, 26, 33]. Recently, [50] and [48] abstract such problems as Concave (Convex) Markov Decision Processes, which focus on finding a policy to maximize (minimize) a concave (convex) function of the induced state-action occupancy measure. We refer the reader to [50, 48, 32] and the references therein for more examples of concave MDPs.

However, the application of concave MDP frameworks in safety-critical domains—such as autonomous driving [21], cyber-security [52], and financial management [2]—introduces additional complexity. To this end, we consider the Concave Constrained Markov Decision Process (Concave CMDP) problem where both the objective and constraints are defined as concave functions of the state-action occupancy measure. This formulation strictly generalizes the classical safe RL and CMDPs [5], which assume that the objective and constraints are expressed by cumulative rewards, i.e., being linear in the state-action occupancy measure.

In this paper, we study concave CMDPs from the optimization perspective and aim to develop a theoretical foundation for direct policy search methods. When moving beyond linear/additive structures in the objective and the constraints, we quickly face several technical challenges. Firstly, the problem has a nonconcave objective and nonconvex constraints with respect to the policy even under the simplest direct policy parameterization. Thus, the existing tools from the convex constrained optimization literature are not applicable. Secondly, the gradient of the objective/constraint with respect to the occupancy measure becomes policy-dependent, thereby making the evaluation for the single-step improvement of the algorithm harder without knowing the occupancy measure. Yet, accurately computing the occupancy measure for a given policy can be inefficient [38]. Thirdly, the performance difference lemma introduced by [25], which is key to the analysis of policy-based primal-dual methods for standard CMDPs [13, 29], is no longer helpful for concave CMDPs.

In view of the aforementioned challenges, our main contributions to the policy search of concave CMDPs are summarized in Table 1 and provided below:

- Despite being nonconcave with respect to the policy and nonlinear with respect to the state-action occupancy measure, we prove that the strong duality still holds for concave CMDP problems under some mild conditions.
- In the exact setting, we propose a simple but effective algorithm—Primal-Dual Policy Gradient method (PDPG) that updates the primal variable via policy gradient ascent and the dual variable via projected sub-gradient descent. In the sample-based setting, we propose the Variance-Reduced PDPG method (VR-PDPG) by incorporating momentum-based recursive estimators for the occupancy measure and the policy gradient. Strong bounds

on the optimality gap and the constraint violation are established for both algorithms, and the effectiveness of variance reduction is illustrated by numerical experiments.

- Inspired by the idea of “optimistic pessimism in the face of uncertainty”, we further demonstrate that by including a diminishing pessimistic term to the constraint, a zero constraint violation can be achieved in both exact and sample-based settings while maintaining the same order of convergence rate for the optimality gap.

1.1 Related Work

RL with General Utilities. Motivated by emerging applications in RL that extend beyond cumulative rewards [36, 1, 24, 23, 35, 26], recent research has focused on developing general approaches for concave (or convex) MDPs. Specifically, [50] introduces a new policy gradient approach called variational policy gradient, and by exploiting the hidden concavity of the problem, it establishes the global convergence of the gradient ascent method in the exact setting. Following this, [51] studies a REINFORCE-based policy gradient method and its variance-reduced version, where the sample complexity of the algorithm is analyzed. More recently, [7] achieves a similar sample complexity result with a new single-loop variance-reduced algorithm. Additionally, [48] transforms the concave MDP to a saddle-point problem using Fenchel duality and proposes a meta-algorithm to solve the problem with standard RL techniques. [22] proves the equivalence between concave MDPs and mean-field games (MFGs) and shows that algorithms for MFGs can be used to solve concave MDPs. However, these studies only address unconstrained problems, which may lead to undesirable policies in safety-critical applications. Thus, additional efforts are required to deal with the emerging safety constraints, and our work addresses this challenge.

CMDP. Our work is also pertinent to policy-based CMDP algorithms [5, 8, 3, 16, 10, 18]. Notably, [13] develops a natural policy gradient-based primal-dual algorithm, demonstrating an $\mathcal{O}(T^{-1/2})$ global convergence rate for both the optimality gap and the constraint violation under standard soft-max parameterization. [42] considers a primal-based approach, achieving a similar convergence rate. More recently, the papers [44, 29, 27] incorporate entropy regularization and obtain improved convergence rates with dual methods. Nonetheless, these papers focus on cumulative rewards/utilities and do not directly generalize to a broader class of safe RL problems, such as safe imitation learning [53] and safe exploration [23]. Beyond CMDPs with cumulative rewards/utilities, [46, 6] also study the concave CMDP problem. Specifically, [46] considers a multi-agent setting and designs a scalable algorithm that converges to a stationary point of the problem. The algorithm by [6] is based on the randomized linear programming proposed by [39] and it achieves the sample complexity of $\tilde{\mathcal{O}}(1/\epsilon^2)$. However, as their approach works directly in the space of occupancy measures, it is thus not applicable to more general problems where the state-action spaces are large and a function approximation is needed. In comparison, our work addresses this issue by focusing on the policy-based primal-dual method and adopting a general soft-max policy parameterization.

1.2 Notations

For a finite set \mathcal{S} , let $\Delta(\mathcal{S})$ denote the probability simplex over \mathcal{S} , and let $|\mathcal{S}|$ denote its cardinality. For a number x , let $[x]_+ := \max\{x, 0\}$. For a random variable x following distribution ρ , we write it as $x \sim \rho$. Let $\mathbb{E}[\cdot]$ (resp. $\mathbb{P}[\cdot]$) and $\mathbb{E}[\cdot | \cdot]$ (resp. $\mathbb{P}[\cdot | \cdot]$) denote the expectation (resp. probability) and conditional expectation (resp. probability) of a random variable. For brevity, we also use the shorthand notation $\mathbb{E}[x]_+ := \mathbb{E}[[x]_+]$. Let \mathbb{R} denote the set of real numbers. For a vector x , we use x^\top to denote the transpose of x and use $\langle x, y \rangle$ to denote the inner product $x^\top y$. We use the convention that $\|x\|_1 = \sum_i |x_i|$, $\|x\| = \sqrt{\sum_i x_i^2}$, and $\|x\|_\infty = \max_i |x_i|$. For a set X in the Euclidean space, let $\text{cl}(X)$ denote the closure of X . Let \mathcal{P}_X denote the projection onto X , defined as $\mathcal{P}_X(y) := \arg \min_{x \in X} \|x - y\|$. For a matrix A , let $\|A\|$ stand for the spectral norm, i.e., $\|A\| = \max_{\|x\| \neq 0} \{\|Ax\|/\|x\|\}$. For a function $f(x)$, let

$\arg \min f(x)$ (resp. $\arg \max f(x)$) denote any global minimum (resp. global maximum) of $f(x)$ and let $\nabla_x f(x)$ denote its gradient with respect to x .

1.3 Organizations

The rest of the paper is organized as follows. In Section 2, we formulate the problem of concave CMDP and introduce the Lagrangian duality used in the paper. Then, in Section 3, we introduce the PDPG and VR-PDPG algorithms and discuss our main analysis technique. Next, in Section 4, we provide the convergence proofs for our algorithms, both in the exact and sample-based settings. In Section 5, we further show that a zero constraint violation can be achieved without sacrificing the convergence rate of the optimality gap. Finally, we validate our algorithms through numerical experiments in Section 6 and close the main paper with concluding remarks in Section 7.

2 Problem Formulation

Standard CMDP. Consider an infinite-horizon CMDP over a finite state space \mathcal{S} and a finite action space \mathcal{A} with a discount factor $\gamma \in [0, 1)$. Let ρ be the initial distribution. The transition dynamics is given by $\mathbb{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$, where $\mathbb{P}(s'|s, a)$ is the probability of transition from state s to state s' when action a is taken. A policy is a function $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ that represents the decision rule the agent uses, i.e., the agent takes action a with probability $\pi(a|s)$ in state s . We denote the set of all stochastic policies as Π . The goal of the agent is to find a policy that maximizes some long-term objective. In standard CMDPs, the agent aims at maximizing the expected (discounted) cumulative reward for a given initial distribution ρ while satisfying constraints on the expected (discounted) cumulative cost, i.e.,

$$\begin{aligned} \max_{\pi \in \Pi} V^\pi(r) &:= \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r(s_k, a_k) \mid \pi, s_0 \sim \rho \right], \\ \text{s.t. } V^\pi(c) &:= \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k c(s_k, a_k) \mid \pi, s_0 \sim \rho \right] \leq 0, \end{aligned} \quad (1)$$

where the expectation is taken over all possible trajectories under policy π with initial distribution ρ , and $r(\cdot, \cdot)$ and $c(\cdot, \cdot)$ denote the reward and cost functions, respectively. We note that the cost can be a vector-valued function, not limited to a single constraint. For given $r(\cdot, \cdot)$, we define the state-action value function (Q-function) under policy π as

$$Q^\pi(r; s, a) := \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r(s_k, a_k) \mid \pi, s_0 = s, a_0 = a \right], \quad (2)$$

which can be interpreted as the expected total reward with an initial state $s_0 = s$ and an initial action $a_0 = a$. For each policy $\pi \in \Pi$ and state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, the unnormalized state-action occupancy measure is defined as

$$\lambda^\pi(s, a) = \sum_{k=0}^{\infty} \gamma^k \mathbb{P}(s_k = s, a_k = a \mid \pi, s_0 \sim \rho). \quad (3)$$

We use Λ to denote the set of all possible state-action occupancy measures, which is a convex polytope (see its expression in (46)). By using λ as decision variables, we observe that the CMDP problem in (1) can be re-parameterized as:

$$\max_{\lambda \in \Lambda} \langle r, \lambda \rangle, \quad \text{s.t. } \langle c, \lambda \rangle \leq 0, \quad (4)$$

where we note that the inner product is taken while viewing λ , r , and c are vectors of dimension $|\mathcal{S}||\mathcal{A}|$. Problem 4 is known as the linear programming formulation of CMDP [5]. Once a solution λ^* is computed, the corresponding policy can be recovered through the one-to-one relation $\pi(a|s) = \lambda^\pi(s, a) / \sum_{a' \in \mathcal{A}} \lambda^\pi(s, a')$.

Concave CMDP. In this work, we consider a more general problem where the agent's goal is to find a policy that maximizes a concave function of the state-action occupancy measure λ subject to constraints defined by another concave function in λ , namely

$$\max_{\lambda \in \Lambda} f(\lambda) \quad \text{s.t. } g(\lambda) \geq 0, \quad (5)$$

where $f(\cdot)$ and $g(\cdot)$ are general concave functions, named as *general utilities*. We note that problem (5) encompasses (4) as a special case with $f(\lambda) = \langle r, \lambda \rangle$ and $g(\lambda) = -\langle c, \lambda \rangle$. As problem (5) is a concave program in λ , we refer to the problem as *Concave CMDP*. We emphasize that the method proposed in this paper directly generalizes to multiple constraints, and we present the single constraint setting only for brevity. Please refer to Remark 1 for the detailed discussion on this generalization. Below, we provide several examples to illustrate the effectiveness of the general utility and the concave CMDP formulation.

EXAMPLE 1 (SAFE LEARNING). *Motivated by the applications of CMDPs in safety-critical systems, [34] develops an approach that guarantees constraint satisfaction during the learning process, in contrast to ordinary CMDP approaches that only ensure policy feasibility upon convergence. They reformulate (1) using the log barrier function, and the objective can be described as the maximization of a general utility, defined as follows*

$$\max_{\lambda \in \Lambda} f(\lambda) = \langle r, \lambda \rangle + \eta \log(-\langle c, \lambda \rangle), \quad (6)$$

where η is a positive penalty weight.

EXAMPLE 2 (SAFE EXPLORATION). *Many learning problems originate from unsupervised settings, where no explicit tasks are assigned. In these situations, agents are typically inclined to efficiently explore the environment [23]. Yet, such explorations can impose significant safety concerns. The work by [43] introduces a model where the learning agent aims to maximize the entropy under the premise of safety, as presented below*

$$\max_{\lambda \in \Lambda} f(\lambda) = \text{Entropy}(\lambda) \quad \text{s.t. } g(\lambda) = -\langle c, \lambda \rangle \geq 0, \quad (7)$$

where the entropy function under policy π is defined as $\text{Entropy}(\lambda^\pi) := -\sum_{s \in \mathcal{S}} d^\pi(s) \cdot \log(d^\pi(s))$ with $d^\pi(s) := (1 - \gamma) \sum_{a' \in \mathcal{A}} \lambda^\pi(s, a')$ denoting the state occupancy measure. This formulation decouples the safety constraints from the reward to alleviate the challenge of designing a single reward signal that must carefully balance task performance with safety.

EXAMPLE 3 (SAFETY-AWARE APPRENTICESHIP LEARNING (AL)). *In AL, instead of maximizing the long-term reward, the agent learns to mimic an expert's demonstrations. When there are critical safety requirements, the learner will also strive to satisfy given constraints on the expected total cost [53]. This problem can be formulated as*

$$\max_{\lambda \in \Lambda} f(\lambda) = -\text{dist}(\lambda, \lambda_e) \quad \text{s.t. } g(\lambda) = -\langle c, \lambda \rangle \geq 0, \quad (8)$$

where λ_e corresponds to the expert demonstration, c denotes the cost function, and $\text{dist}(\cdot, \cdot)$ can be any distance function on Λ , e.g., ℓ^2 -distance or Kullback-Liebler (KL) divergence.

EXAMPLE 4 (FEASIBILITY CONSTRAINED MDPs). *As an extension to standard CMDPs, the designer may desire to control the MDP through more general constraints described by a convex feasibility region C such that the learned policy is not too far away from C [31]. In this case, the problem can be cast as*

$$\max_{\lambda \in \Lambda} f(\lambda) = \langle r, \lambda \rangle \quad \text{s.t. } g(\lambda) = -\text{dist}(\lambda, C) \geq -d_0, \quad (9)$$

where $d_0 \geq 0$ denotes the threshold of the allowable deviation.

Policy Parameterization. Since recovering a policy from the associated state-action occupancy measure is straightforward, a natural approach to solving the concave CMDP problem is to optimize (5) directly (or equivalently (4) for standard CMDPs). However, since the decision variable λ has the size $|\mathcal{S}||\mathcal{A}|$, such approaches lack scalability and can converge extremely slowly for large state and action spaces. In this work, we consider the direct policy search method, which can handle the curse of dimensionality via policy parameterization. We assume that the policy $\pi = \pi_\theta$ is parameterized by a general soft-max function, i.e.,

$$\pi_\theta(a|s) = \frac{\exp(\psi(\theta; s, a))}{\sum_{a' \in \mathcal{A}} \exp(\psi(\theta; s, a'))}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad (10)$$

where $\psi(\cdot; s, a)$ is some smooth function and $\theta \in \mathbb{R}^K$ is the unconstrained parameter vector. We assume that θ over-parameterizes the set of all stochastic policies, in the sense that $\text{cl}(\{\pi_\theta | \theta \in \mathbb{R}^K\}) = \Pi$. Equivalently, this implies $\text{cl}(\{\lambda^{\pi_\theta} | \theta \in \mathbb{R}^K\}) = \Lambda$. This condition essentially ensures that any stochastic policy can be sufficiently approximated by the parameterized policy class. Without this assumption, some policies would be unrepresentable, thereby inherently shrinking the feasible region of the original problem. Further assumptions on the parameterization will be formally stated in Section 4. In practice, the function ψ can be chosen as a deep neural network, where θ is the parameter and the state-action pair (s, a) is the input. Under the parameterization (10), problem (5) can be rewritten as

$$\max_{\theta \in \mathbb{R}^K} f(\lambda(\theta)) \quad \text{s.t.} \quad g(\lambda(\theta)) \geq 0, \quad (11)$$

where we introduce the shorthand notations $\lambda(\theta) := \lambda^{\pi_\theta}$ and $\lambda(\theta; s, a) := \lambda^{\pi_\theta}(s, a)$. It is worth mentioning that, differently from (5), the problem (11) is nonconcave due to its nonconcave objective function and nonconvex constraint with respect to θ .

2.1 Lagrangian Duality

Consider the Lagrangian function associated with (11), defined as $L(\lambda(\theta), \mu) := f(\lambda(\theta)) + \mu g(\lambda(\theta))$. The dual function is $D(\mu) := \max_{\theta \in \mathbb{R}^K} L(\lambda(\theta), \mu)$. Let π_{θ^*} be an optimal policy such that θ^* is an optimal solution to (11), and μ^* be an optimal dual variable. In constrained optimization, strict feasibility can induce many desirable properties. We assume that the following Slater's condition holds.

ASSUMPTION 1 (SLATER'S CONDITION). *There exist $\tilde{\theta}$ and $\xi > 0$ such that $g(\lambda(\tilde{\theta})) \geq \xi$.*

Slater's condition is a standard assumption in convex optimization, and it holds when the feasible region has a non-empty interior. In practice, identifying an interior point $\tilde{\theta}$ is often straightforward with prior knowledge of the problem. In the context of CMDPs, an interior point corresponds to a safe policy that remains feasible under small perturbations, ensuring the safety constraints are still satisfied. This interior point ensures that the optimization problem satisfies the regularity conditions for the duality theory to apply effectively. The following result is a direct consequence of Slater's condition [5].

LEMMA 1 (STRONG DUALITY AND BOUNDEDNESS OF μ^*). *Let Assumption 1 hold, and suppose that $\text{cl}(\{\lambda(\theta) | \theta \in \mathbb{R}^K\}) = \Lambda$, where $\text{cl}(\cdot)$ denotes the closure operation. Then, we have that*

$$f(\lambda(\theta^*)) = D(\mu^*) = L(\lambda(\theta^*), \mu^*); \quad 0 \leq \mu^* \leq [f(\lambda(\theta^*)) - f(\lambda(\tilde{\theta}))] / \xi. \quad (12)$$

Hence, it suffices to search the optimal dual variable within the closed interval $U := [0, C_0]$, where $C_0 := 1 + (M - f(\lambda(\tilde{\theta})))/\xi$ and M is some upper bound on $f(\cdot)$.

For completeness, we provide a proof for Lemma 1 in Appendix B. We remark that in the definition of C_0 , the constant M will be formalized in Lemma 2, and the constant number 1 plays as a slackness to facilitate the

convergence analysis. The strong duality implies that (11) is equivalent to the following saddle point problem:

$$\max_{\theta \in \mathbb{R}^K} \min_{\mu \in U} L(\lambda(\theta), \mu) = \min_{\mu \in U} \max_{\theta \in \mathbb{R}^K} L(\lambda(\theta), \mu). \quad (13)$$

Motivated by this equivalence, we seek to develop a primal-dual algorithm.

3 Safe Policy Search Beyond Cumulative Rewards/Utilities

To solve (13), we propose the following Primal-Dual Policy Gradient Algorithm (PDPG):

$$\begin{cases} \theta_{t+1} = \theta_t + \eta_{\theta,t} \nabla_{\theta} L(\lambda(\theta_t), \mu_t) \\ \mu_{t+1} = \mathcal{P}_U (\mu_t - \eta_{\mu} \nabla_{\mu} L(\lambda(\theta_t), \mu_t)) \end{cases}, \quad \forall t = 0, 1, \dots, \quad (14)$$

where $\eta_{\theta,t} > 0$, $\eta_{\mu} > 0$ are step-sizes, and the dual feasible region $U = [0, C_0]$ is introduced in Lemma 1. The method (14) adopts an alternating update scheme: the primal step performs the gradient ascent in the policy space, whereas the dual step updates the multiplier with projected sub-gradient descent such that μ_{t+1} is obtained by adding a multiple of the constraint violation to μ_t . The values of $\eta_{\theta,t}$ and η_{μ} will be specified later in the paper.

In an MDP with the reward r , the *Policy gradient* theorem [37] suggests that $\nabla_{\theta} V^{\pi_{\theta}}(r)$ can be estimated via the following stochastic estimator (see Lemma 15):

$$\widehat{d}(\tau, \theta, r) = \sum_{h=0}^{H-1} \nabla \log \pi_{\theta}(a_h | s_h) \left[\sum_{k=h}^{H-1} \gamma^k r(s_k, a_k) \right], \quad (15)$$

where $\tau = \{(s_k, a_k)\}_{k=0}^{H-1}$ is a trajectory under policy π_{θ} and initial distribution ρ . In the standard CMDP where the value function is linear in the occupancy measure (see (4)), the Lagrangian can be viewed as a value function with the reward $r - \mu c$, i.e., $L(\lambda(\theta), \mu) = V^{\pi_{\theta}}(r - \mu c)$. Hence, it is natural to estimate $\nabla_{\theta} L(\lambda(\theta), \mu)$ via $\widehat{d}(\tau, \theta, r - \mu c)$. However, this favorable result is not applicable when $f(\cdot)$ and $g(\cdot)$ are general concave functions. To address this challenge, one approach is to utilize the REINFORCE-based Policy Gradient introduced by [51] to express $\nabla_{\theta} L(\lambda(\theta), \mu)$ as the policy gradient of a standard value function with reward $\nabla_{\lambda} L(\lambda(\theta), \mu)$, i.e.,

$$\nabla_{\theta} L(\lambda(\theta), \mu) = [\nabla_{\theta} \lambda(\theta)]^{\top} \nabla_{\lambda} L(\lambda(\theta), \mu) = \nabla_{\theta} V^{\pi_{\theta}}(\nabla_{\lambda} L(\lambda(\theta), \mu)), \quad (16)$$

where the first equality follows from the chain rule, and the second equality uses the fact that $\nabla_{\theta} V^{\pi_{\theta}}(r) = [\nabla_{\theta} \lambda(\theta)]^{\top} r$. To distinguish the policy-dependent reward in (16) from the standard reward function, we introduce the concept of *shadow reward*, defined as $\nabla_{\lambda} L(\lambda(\theta), \mu) = \nabla_{\lambda} f(\lambda(\theta), \mu) + \mu \nabla_{\lambda} g(\lambda(\theta), \mu)$. Equation (16) suggests the following scheme for computing $\nabla_{\theta} L(\lambda(\theta), \mu)$. Firstly, we collect a trajectory $\tau = \{(s_k, a_k)\}_{k=0}^{H-1}$ under π_{θ} and approximate the occupancy measure $\lambda(\theta)$ through the Monte Carlo estimator:

$$\widehat{\lambda}(\tau) = \sum_{k=0}^{H-1} \gamma^k \cdot \mathbb{1}(s_k, a_k), \quad (17)$$

where $\mathbb{1}(s, a) \in \mathbb{R}^{|S| \times |A|}$ is the indicator vector that takes one in the (s, a) -th entry and zero elsewhere. It is clear that $\mathbb{E}_{\tau \sim \pi_{\theta}, \rho} [\widehat{\lambda}(\tau)] = \sum_{k=0}^{H-1} \gamma^k \mathbb{P}(s_k = s, a_k = a | \pi_{\theta}, s_0 \sim \rho) =: \lambda_H(\theta)$, i.e., $\widehat{\lambda}(\tau)$ is an unbiased estimator for the truncated occupancy measure $\lambda_H(\theta)$. Secondly, we estimate the true shadow reward $\nabla_{\lambda} L(\lambda(\theta), \mu)$ by $\nabla_{\lambda} L(\widehat{\lambda}(\tau), \mu)$. Finally, we substitute the shadow reward into the stochastic estimator (15), and approximate the policy gradient $\nabla_{\theta} L(\lambda(\theta), \mu)$ by $\widehat{d}(\tau, \theta, \nabla_{\lambda} L(\widehat{\lambda}(\tau), \mu))$. Nevertheless, without additional techniques, the convergence for this vanilla REINFORCE-based approach can be slow, especially with large state and action spaces [38, 50].

3.1 Variance-Reduced Primal-Dual Policy Gradient Algorithm

In this section, we present the sample-based version of algorithm (14), where we incorporate variance reduction in the gradient evaluation of the Lagrangian function. Specifically, we adopt a recursive scheme similar to the STORM algorithm in stochastic optimization [11, 7], which only requires a single policy rollout in each iteration. Compared to the vanilla Monte Carlo estimators in (15) and (17), the variance-reduced estimators can achieve a more accurate and sample-efficient approximation.

We begin by introducing the importance sampling (IS), which plays a key role in variance-reduced policy gradient methods [47, 51, 17, 7]. Given any $\theta \in \mathbb{R}^K$, let $\tau = \{(s_k, a_k)\}_{k=0}^{H-1}$ be a trajectory generated under π_θ . Then, for any target policy θ' , the IS weight is defined as

$$w(\tau|\theta', \theta) := \prod_{k=0}^{H-1} \frac{\pi_{\theta'}(a_k|s_k)}{\pi_\theta(a_k|s_k)}. \quad (18)$$

One major effect of $w(\tau|\theta', \theta)$ is to mitigate the distribution shift brought by the policy update [11, 51, 7]. For example, it holds that

$$\mathbb{E}_{\tau \sim \pi_{\theta, \rho}} \left[w(\tau|\theta', \theta) \widehat{\lambda}(\tau) \right] = \lambda_H(\theta'), \quad \mathbb{E}_{\tau \sim \pi_{\theta, \rho}} \left[w(\tau|\theta', \theta) \widehat{d}(\tau, \theta', r) \right] = [\nabla_{\theta} \lambda_H(\theta')]^\top r, \quad (19)$$

where we recall that $\lambda_H(\theta) = \sum_{k=0}^{H-1} \gamma^k \mathbb{P}(s_k = s, a_k = a | \pi_\theta, s_0 \sim \rho)$ is the truncated occupancy measure. Consequently, even when τ is generated under policy π_θ in period t , we can still construct an unbiased estimator for $\lambda_H(\theta_{t-1})$ by scaling $\widehat{\lambda}(\tau)$ with the IS weight $w(\tau|\theta_{t-1}, \theta_t)$. This idea forms the basis for using momentum-based variance reduction methods, which typically require simulating stochastic estimators for preceding variables using the data from the current period. The same approach also applies to the estimation of the policy gradient. In Appendix B.2, we provide a proof for (19).

Now, we present the Variance-Reduced Primal-Dual Policy Gradient Algorithm (VR-PDPG), as outlined in Algorithm 1. In each iteration, VR-PDPG consists of four stages.

- **Stage 1** (lines 9-11): We collect a single trajectory τ_t and construct a vanilla occupancy measure estimator through (17), which is then used to refine the variance-reduced estimator λ_t in (20). Specifically, λ_t is updated in a recursive way akin to the STORM algorithm [11, 7]. Next, we substitute λ_t into $\nabla_{\lambda} f(\cdot)$ and $\nabla_{\lambda} g(\cdot)$ to obtain empirical shadow rewards.
- **Stage 2** (lines 12): We use a similar momentum-based scheme to estimate the policy gradients $\nabla_{\theta} f(\lambda(\theta_t))$ and $\nabla_{\theta} g(\lambda(\theta_t))$. Note that the shadow rewards from the previous two periods, i.e., $r_{\diamond, t-1}$ and $r_{\diamond, t-2}$, are intentionally used to avoid the independence issues.
- **Stage 3** (lines 13): We update the policy parameter through a gradient ascent. Motivated by [7], we use an normalized step-size such that $\eta_{\theta, t} = \eta_{\theta} / \|d_{L, t}\|$ for all $t \geq 0$, where $d_{L, t}$ is the approximate gradient direction obtained in stage 2. This normalization ensures the boundedness of the IS weight (see Appendix B.2 for further discussions).
- **Stage 4** (lines 14): Finally, we adjust the dual variable by a projected sub-gradient descent, where the dual gradient $\nabla_{\mu} L(\lambda(\theta_t), \mu_t)$ is approximated by $g(\lambda_t)$. Recall that Lemma 1 ensures that we can limit the dual variable search within the closed interval $U = [0, C_0]$.

3.2 Exploiting the Hidden Concavity

By itself, (13) is a nonconcave-linear maximin problem. The existing results [28] for the analysis of the gradient ascent descent method for such problems can only guarantee the convergence to a stationary point in the rate of $O(T^{-1/6})$. Thus, new analyses are required to obtain an improved rate and achieve global optimality. However, we comment that those analyses based on the performance difference lemma (see Lemma 18) do not apply to

Algorithm 1 Variance-Reduced Primal-Dual Policy Gradient Algorithm (VR-PDPG)

-
- 1: **Input:** Iteration number T ; initial policy θ_0 ; initial dual variable μ_0 ; step-sizes η_θ, η_μ , and $\{\alpha_t\}_{t \geq 1}$; trajectory length H ; dual feasible region $U = [0, C_0]$.
 - 2: Sample trajectory τ_0 with length H under policy π_{θ_0} .
 - 3: Compute occupancy measure: $\lambda_0 = \widehat{\lambda}(\tau_0)$.
 - 4: Compute shadow rewards $r_{f,0} = \nabla_\lambda f(\lambda_0)$ and $r_{g,0} = \nabla_\lambda g(\lambda_0)$. Let $r_{f,-1} = r_{f,0}$ and $r_{g,-1} = r_{g,0}$.
 - 5: Compute gradients: $d_{\diamond,0} = \widehat{d}(\tau_0, \theta_0, r_{\diamond,0})$, for $\diamond \in \{f, g\}$.
 - 6: Update policy: $\theta_1 = \theta_0 + \eta_{\theta,0} d_{L,0}$, where $d_{L,0} = d_{f,0} + \mu_0 d_{g,0}$ and $\eta_{\theta,0} = \eta_\theta / \|d_{L,0}\|$.
 - 7: Update dual variable: $\mu_1 = \mathcal{P}_U(\mu_0 - \eta_\mu \cdot g(\lambda_0))$.
 - 8: **for** iteration $t = 1, 2, \dots, T - 1$ **do**
 - 9: Sample trajectory τ_t with length H under policy π_{θ_t} .
 - 10: Compute occupancy measure:

$$\lambda_t = \widehat{\lambda}(\tau_t) + (1 - \alpha_t) [\lambda_{t-1} - w(\tau_t | \theta_{t-1}, \theta_t) \widehat{\lambda}(\tau_t)]. \quad (20)$$

- 11: Compute shadow rewards $r_{f,t} = \nabla_\lambda f(\lambda_t)$ and $r_{g,t} = \nabla_\lambda g(\lambda_t)$.
- 12: Compute gradients:

$$d_{\diamond,t} = \widehat{d}(\tau_t, \theta_t, r_{\diamond,t-1}) + (1 - \alpha_t) [d_{\diamond,t-1} - w(\tau_t | \theta_{t-1}, \theta_t) \widehat{d}(\tau_t, \theta_{t-1}, r_{\diamond,t-2})], \quad (21)$$

for $\diamond \in \{f, g\}$.

- 13: Update policy: $\theta_{t+1} = \theta_t + \eta_{\theta,t} d_{L,t}$, where $d_{L,t} = d_{f,t} + \mu_t d_{g,t}$ and $\eta_{\theta,t} = \eta_\theta / \|d_{L,t}\|$.
 - 14: Update dual variable: $\mu_{t+1} = \mathcal{P}_U(\mu_t - \eta_\mu \cdot g(\lambda_t))$.
 - 15: **end for**
-

concave CMDPs [4, 13]. A key insight is that, due to the loss of linearity, the performance difference lemma can only provide an upper bound for the single-step improvement with the gradient information at the current step, i.e.,

$$L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_t), \mu_t) \leq \frac{1}{1 - \gamma} \mathbb{E}_{s \sim d^{\pi_{\theta_{t+1}}}} \mathbb{E}_{a \sim \pi_{\theta_{t+1}}} [A^{\pi_{\theta_t}}(\nabla_\lambda L(\lambda(\theta_t), \mu_t); s, a)], \quad (22)$$

where $d^\pi(\cdot)$ is the state occupancy measure and $A^\pi(r; s, a) := Q^\pi(r; s, a) - \mathbb{E}_{a \sim \pi(\cdot|s)} [Q^\pi(r; s, a)]$ is the advantage function with reward r . Below, we elaborate on how the ‘‘hidden concavity’’ of (11) with respect to λ can help us relate the policy update to the optimality gap.

Assume, for illustration purposes, that the inverse mapping $\lambda^{-1} : \Lambda \rightarrow \mathbb{R}^K$ is globally well-defined with $\lambda^{-1}(\lambda(\theta)) = \theta$ for all $\theta \in \mathbb{R}^K$. Then, for every period t , we can associate θ_t with an occupancy measure constructed as $(1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*) \in \Lambda$, which admits the inverse $\theta_{\varepsilon,t} := \lambda^{-1}((1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*)) \in \mathbb{R}^K$. Since $L(\cdot, \mu)$ is concave, it follows that

$$\begin{aligned} L(\lambda(\theta_{\varepsilon,t}), \mu_t) &= L((1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*), \mu_t) \geq (1 - \varepsilon)L(\lambda(\theta_t), \mu_t) + \varepsilon L(\lambda(\theta^*), \mu_t) \\ \iff L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{\varepsilon,t}), \mu_t) &\leq (1 - \varepsilon)[L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)]. \end{aligned} \quad (23)$$

Hence, if we can demonstrate that $L(\lambda(\theta_{t+1}), \mu_t) \approx L(\lambda(\theta_{\varepsilon,t}), \mu_t)$, inequality (23) implies a nearly $(1 - \varepsilon)$ -contraction of the Lagrangian from (θ_t, μ_t) to (θ_{t+1}, μ_t) . When the dual step-size is properly chosen, we can further obtain a recursive relation between the optimality gaps $L(\lambda(\theta^*), \mu_{t+1}) - L(\lambda(\theta_{t+1}), \mu_{t+1})$ and $L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)$. This relation paves the way for controlling the average performance in terms of Lagrangian, i.e., $\frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)]$, which is the key quantity in our convergence analysis.

The argument presented above offers an initial insight for our analysis. However, it is crucial to comment that the global inverse $\lambda^{-1}(\cdot)$ typically does not exist under policy parameterization. This limitation is particularly evident for soft-max type parameterizations such as (10), where the parameter space is unconstrained (see Appendix A for further discussions). Therefore, motivated by [51], we employ a more flexible condition that only requires the local invertibility of the mapping, which will be formalized in Assumption 2. Finally, it is worth noting that although the hidden concavity/convexity has been studied in unconstrained (RL) problems [50, 51, 7, 20], the analysis diverges remarkably in the presence of constraints.

4 Convergence Analysis

By exploiting the hidden concavity of (11), we establish the global convergence for both the PDPG and VR-PDPG algorithms in the time-average sense. Particularly, we are interested in quantifying the average optimality gap $\frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_t))]$, and constraint violation $\frac{1}{T} [\sum_{t=0}^{T-1} -g(\lambda(\theta_t))]_+$, where we recall that $[x]_+ := \max\{x, 0\}$. By definition, when the average performance bound holds, an iteration selected uniformly at random from the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ would satisfy the same performance bound in expectation.

First, we formally state our assumption about parameterization (10). To avoid introducing additional biases, it is natural to assume that π_θ has enough expressibility to represent any policy, i.e., $\forall \pi \in \Pi, \exists \theta \in \mathbb{R}^K$ such that $\pi = \pi_\theta$. However, assuming a one-to-one correspondence between π and θ is too restrictive, and it is not satisfied even by the soft-max policy. In practice, using a deep neural network to represent the policy can often arrive at an over-parameterization. Therefore, following [51], we assume that π_θ is defined such that it can represent any policy and that $\lambda(\cdot) : \mathbb{R}^K \rightarrow \Lambda$ is locally continuously invertible. As briefly introduced in Section 3.2, this (local) invertibility is important in exploiting the hidden concavity. A more detailed discussion can be found in Appendix A.

ASSUMPTION 2 (PARAMETERIZATION). *Under general soft-max parameterization (10), $\psi(\cdot; s, a)$ is twice differentiable for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ and there exist $M_{\pi,1}, M_{\pi,2} > 0$ such that*

$$\max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sup_{\theta \in \mathbb{R}^K} \|\nabla_\theta \psi(\theta; s, a)\| \leq M_{\pi,1}, \quad \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sup_{\theta \in \mathbb{R}^K} \|\nabla_\theta^2 \psi(\theta; s, a)\| \leq M_{\pi,2}. \quad (24)$$

Furthermore, π_θ over-parameterizes the set of all stochastic policies in the following senses:

- (I) For every $\theta \in \mathbb{R}^K$, there exists a neighborhood $\mathcal{U}_\theta \ni \theta$ such that the restriction of $\lambda(\cdot)$ to \mathcal{U}_θ is a bijection between \mathcal{U}_θ and $\mathcal{V}_{\lambda(\theta)} := \lambda(\mathcal{U}_\theta)$, where we denote $\lambda_{\mathcal{V}_{\lambda(\theta)}}^{-1} : \mathcal{V}_{\lambda(\theta)} \rightarrow \mathcal{U}_\theta$ as the local inverse of $\lambda(\cdot)$, i.e., $\lambda_{\mathcal{V}_{\lambda(\theta)}}^{-1}(\lambda(\theta_0)) = \theta_0, \forall \theta_0 \in \mathcal{U}_\theta$.
- (II) There exists $L_\lambda > 0$ such that $\lambda_{\mathcal{V}_{\lambda(\theta)}}^{-1}(\cdot)$ is L_λ -Lipschitz continuous for all $\theta \in \mathbb{R}^K$.
- (III) There exists $\bar{\varepsilon} > 0$ such that $(1 - \varepsilon)\lambda(\theta) + \varepsilon\lambda(\theta^*) \in \mathcal{V}_{\lambda(\theta)}, \forall \varepsilon \leq \bar{\varepsilon}, \forall \theta \in \mathbb{R}^K$.

Under (24) of Assumption 2, the score function $\nabla_\theta \log \pi_\theta(a|s)$ is always upper-bounded (see Lemma 19), thereby ensuring the boundedness of the policy gradient estimator in (15). Next, we make assumptions on the Lipschitz continuity and smoothness of $f(\cdot)$ and $g(\cdot)$.

ASSUMPTION 3 (UTILITY FUNCTION). *There exist $M_\lambda, \ell_\lambda > 0$ such that*

- (I) $f(\cdot)$ and $g(\cdot)$ are M_λ -Lipschitz continuous, i.e., $\|\nabla_\lambda f(\lambda)\| \leq M_\lambda, \|\nabla_\lambda g(\lambda)\| \leq M_\lambda$ for all λ .
- (II) $f(\cdot)$ and $g(\cdot)$ are ℓ_λ -smooth in the sense that for all λ_1, λ_2 , it holds that

$$\|\nabla_\lambda f(\lambda_1) - \nabla_\lambda f(\lambda_2)\|_\infty \leq \ell_\lambda \|\lambda_1 - \lambda_2\|, \quad \|\nabla_\lambda g(\lambda_1) - \nabla_\lambda g(\lambda_2)\|_\infty \leq \ell_\lambda \|\lambda_1 - \lambda_2\|.$$

We note that Assumptions 2 and 3 are standard in the literature of RL with general utilities [51, 49, 7, 46]. Apart from the local invertibility condition of the policy parameterization, these assumptions impose only mild smoothness requirements on the policy and utility functions. For further justification of these assumptions, we

refer readers to [51, Lemma 5.3] and [46, Appendix F.1.2]. The following function properties are directly derived from Assumptions 2 and 3.

LEMMA 2. *Let Assumptions 2 and 3 hold. There exist $M, \ell_\theta, L_{\theta, \lambda} > 0$ such that*

- (I) $f(\cdot)$ and $g(\cdot)$ are bounded by M on Λ , i.e., $|f(\lambda)| \leq M, |g(\lambda)| \leq M$, for all $\lambda \in \Lambda$.
 (II) $f(\lambda(\theta))$ and $g(\lambda(\theta))$ are ℓ_θ -smooth w.r.t. θ , i.e., for all $\theta_1, \theta_2 \in \mathbb{R}^K$, it holds that

$$\|\nabla_\theta f(\lambda(\theta_1)) - \nabla_\theta f(\lambda(\theta_2))\| \leq \ell_\theta \|\theta_1 - \theta_2\|, \quad \|\nabla_\theta g(\lambda(\theta_1)) - \nabla_\theta g(\lambda(\theta_2))\| \leq \ell_\theta \|\theta_1 - \theta_2\|.$$

- (III) For every $H > 0$ and $\theta \in \mathbb{R}^K$, it holds that

$$\|\nabla_\theta f(\lambda(\theta)) - \nabla_\theta f(\lambda_H(\theta))\| \leq L_{\theta, \lambda} \gamma^H, \quad \|\nabla_\theta g(\lambda(\theta)) - \nabla_\theta g(\lambda_H(\theta))\| \leq L_{\theta, \lambda} \gamma^H,$$

where $\lambda_H(\theta) = \sum_{k=0}^{H-1} \gamma^k \mathbb{P}(s_k = s, a_k = a | \pi_\theta, s_0 \sim \rho)$.

- (IV) For any $\mu \in U = [0, C_0]$, the Lagrangian $L(\cdot, \mu)$ is bounded by $(1 + C_0)M$ on Λ ; $L(\lambda(\theta), \mu)$ is $(1 + C_0)\ell_\theta$ -smooth w.r.t. θ ; for every $H > 0$ and $\theta \in \mathbb{R}^K$, it holds that

$$\|\nabla_\theta L(\lambda(\theta), \mu) - \nabla_\theta L(\lambda_H(\theta), \mu)\| \leq (1 + C_0)L_{\theta, \lambda} \gamma^H. \quad (25)$$

It is worth mentioning that the property (III) is similar to the Lipschitz continuity of $\nabla_\theta f(\lambda)$ with respect to λ , where we use the fact that $\|\lambda(\theta) - \lambda_H(\theta)\|$ is bounded by $\gamma^H/(1 - \gamma)$. This boundedness plays a crucial role in the convergence proof for the sample-based scenario since λ_t in Algorithm 1 is indeed an unbiased estimator for the truncated occupancy measure $\lambda_H(\theta)$, instead of $\lambda(\theta)$ directly. In addition, we note that the properties in (IV) are a direct consequence of (I) to (III) since the Lagrangian is the simple sum of $f(\lambda(\theta))$ and $\mu g(\lambda(\theta))$.

The rest of this section is divided into two cases: the exact setting (PDPG in (14)) and the sample-based setting (VR-PDPG in Algorithm 1). We unfold the convergence results by first presenting a key finding regarding the average performance of the Lagrangian, which lays the foundation for deriving both the optimality gap and constraint violation.

4.1 Exact Setting

This section is dedicated to the convergence analysis of the PDPG algorithm in (14). We first characterize the average performance bounds of the Lagrangian for both general concave and σ -strongly concave frameworks in the following lemma.

LEMMA 3. *Let Assumptions 2 and 3 hold. Then, for every $T > 0$ and $\varepsilon \leq \bar{\varepsilon}$, the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ generated by PDPG in (14) with $\eta_{\theta, t} \equiv \frac{1}{(1+C_0)\ell_\theta}$ satisfies that*

$$\frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] \leq \frac{L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)}{\varepsilon T} + \frac{2\varepsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon}. \quad (26)$$

If $f(\cdot)$ is σ -strongly concave with respect to λ , $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ further satisfies that

$$\frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] \leq \frac{L(\lambda(\theta_T), \mu_{T-1}) - L(\lambda(\theta_0), \mu_0)}{\tilde{\varepsilon} T} + \frac{\eta_\mu M^2}{\tilde{\varepsilon}}, \quad (27)$$

where $\tilde{\varepsilon} = \min \left\{ \bar{\varepsilon}, \frac{\sigma}{\sigma + 2(1+C_0)\ell_\theta L_\lambda^2} \right\}$.

The core idea in proving Lemma 3 is that one can relate the primal update to the sub-optimality gap $L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)$ by leveraging the hidden concavity of (11) with respect to λ . Then, as $|\mu_{t+1} - \mu_t| = \mathcal{O}(\eta_\mu)$, we are able to draw a recursion between the sub-optimality gaps for two consecutive periods (see (68)). The

average performance in terms of the Lagrangian can be decomposed into the summation of the average optimality gap and the weighted average of the constraint violation, i.e.,

$$\frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] = \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_t))] + \frac{1}{T} \sum_{t=0}^{T-1} \mu_t [g(\lambda(\theta^*)) - g(\lambda(\theta_t))].$$

Since θ^* must be a feasible solution, the difference $g(\lambda(\theta^*)) - g(\lambda(\theta_t))$ can be interpreted as an approximate of the constraint violation. Next, in Lemma 4, we decouple the Lagrangian bound to obtain the optimality gap and the true constraint violation.

LEMMA 4. *Let Assumptions 1 and 3 hold, and the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ generated by PDPG in (14) satisfy that $\frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] \leq C$. Then, the optimality gap and constraint violation are upper-bounded as follows*

$$\frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_t))] \leq C + \frac{\mu_0^2 - \mu_T^2}{2T\eta_\mu} + \frac{\eta_\mu M^2}{2}, \quad (28a)$$

$$\frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ \leq C + \frac{\max_{\mu \in U} \{(\mu_0 - \mu)^2 - (\mu_T - \mu)^2\}}{2T\eta_\mu} + \frac{\eta_\mu M^2}{2}. \quad (28b)$$

Recall that the dual feasible region $U = [0, C_0]$ is set by taking $C_0 = 1 + (M_F - f(\lambda(\tilde{\theta}))) / \xi$ in Lemma 1. Since $\mu^* \leq [f(\lambda(\theta^*)) - f(\lambda(\tilde{\theta}))] / \xi \leq C_0 - 1$, it implies that $\hat{\mu} := \mu^* + 1 \in U$. This “slackness” is important when bounding the constraint violation, as we can write $\left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ = [(\mu^* - \hat{\mu}) \sum_{t=0}^{T-1} \nabla_\mu L(\lambda(\theta_t), \mu_t)]_+$, where the latter term can be related to the first-order expansion of $L(\lambda(\theta), \cdot)$ and bounded via the use of telescoping sums.

Combining Lemmas 3 and 4, we optimize the step-size choice and establish the global convergence rate for algorithm (14) in the following theorem.

THEOREM 1 (EXACT SETTING). *Let Assumptions 1, 2, and 3 hold. For every $T \geq (\bar{\epsilon})^{-3}$, we choose $\mu_0 = 0$, $\eta_{\theta,t} = \frac{1}{(1+C_0)t_\theta}$, and $\eta_\mu = T^{-2/3}$. Then, the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ generated by PDPG in (14) satisfies that*

$$\frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_t))] \leq \frac{2M + M^2/2}{T^{2/3}} + \frac{2(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2 T^{1/3}} + \frac{2M^2}{T^{1/3}}, \quad (29a)$$

$$\frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ \leq \frac{2M + M^2/2}{T^{2/3}} + \frac{2(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2 T^{1/3}} + \frac{2M^2 + C_0^2/2}{T^{1/3}}. \quad (29b)$$

Suppose further that $f(\cdot)$ is σ -strongly concave with respect to λ . For every $T > 0$, we choose $\mu_0 = 0$, $\eta_{\theta,t} \equiv \frac{1}{(1+C_0)t_\theta}$, and $\eta_\mu = T^{-1/2}$. Then, the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ satisfies that

$$\frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_t))] \leq \frac{(2+C_0)M}{\tilde{\epsilon}T} + \left(\frac{M^2}{\tilde{\epsilon}} + \frac{M^2}{2} \right) \frac{1}{\sqrt{T}}, \quad (30a)$$

$$\frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ \leq \frac{(2+C_0)M}{\tilde{\epsilon}T} + \left(\frac{M^2}{\tilde{\epsilon}} + \frac{M^2 + C_0^2}{2} \right) \frac{1}{\sqrt{T}}, \quad (30b)$$

where $\tilde{\epsilon} = \min \left\{ \bar{\epsilon}, \frac{\sigma}{\sigma+2(1+C_0)\ell_\theta L_\lambda^2} \right\}$.

For general concave CMDPs, Theorem 1 shows that PDPG (14) achieves a global convergence in the time-average sense such that the optimality gap and constraint violation decay to zero with the rate of $\mathcal{O}(T^{-1/3})$. In other words, to obtain an $\mathcal{O}(\epsilon)$ -approximate solution, the iteration complexity is $\mathcal{O}(\epsilon^{-3})$. When $f(\cdot)$ and $g(\cdot)$ are linear functions as in standard CMDPs, Theorem 1 shows a better convergence rate than the existing $\mathcal{O}(T^{-1/4})$ rate for the policy gradient primal-dual methods (see [14] or [15, Theorem 6]). When $f(\lambda(\theta))$ is strongly concave with respect to λ , we can further improve the convergence rate to $\mathcal{O}(T^{-1/2})$ by taking the dual step-size $\eta_\mu = \mathcal{O}(T^{-1/2})$. This rate improvement mainly results from the better average performance of the Lagrangian in the strongly concave scenario (see (27)), which does not contain an $\mathcal{O}(\bar{\epsilon})$ error term. Hence, unlike the general concave setting where we need to choose $\epsilon = \mathcal{O}(T^{-1/3})$ in (26), we can choose $\epsilon = \mathcal{O}(1)$ to obtain a sharper convergence rate in the strongly concave scenario.

4.2 Sample-based Setting

We now analyze the convergence of VR-PDPG, as presented in Algorithm 1, with a focus on the general concave setting. Before delving into the technical details, we articulate the rationale for employing variance reduction over a simpler averaged stochastic policy gradient estimator. The primary reason is sample efficiency. To achieve an ϵ -accurate policy gradient estimator, an averaged approach would typically require $\mathcal{O}(1/\epsilon^2)$ samples per iteration. Combined with the $\mathcal{O}(1/\epsilon^3)$ iteration complexity of PDPG from Theorem 1, the total sample complexity would be at least $\mathcal{O}(1/\epsilon^5)$. This is significantly less efficient than the VR-PDPG complexity shown in Theorem 2 below. Secondly, our method's exploitation of hidden concavity is incompatible with estimators that have a fixed bounded error, a typical trait of averaged stochastic gradients. Such an error would accumulate and result in a non-vanishing error term in the final performance bound.

To overcome these shortcomings, we incorporate a variance reduction technique. Furthermore, unlike in the exact setting, our sample-based algorithm adopts a normalized policy update rule, *i.e.*, $\eta_{\theta,t} = \eta_\theta / \|\nabla_\theta \mathcal{L}_t\|$, which necessitates the use of sufficiently small primal step-sizes. Consequently, in this sample-based setting, the benefits of strong concavity are insufficient to offset the approximation errors, resulting in no improvement in the convergence rate over the general concave case.

The remainder of this section is structured as follows. We begin by establishing an average performance bound for the Lagrangian function in the sample-based scenario (see Lemma 5). This bound is subsequently decoupled into an optimality gap and a constraint violation term (see Lemma 6). These intermediate results pave the way for our main convergence guarantee in Theorem 2.

LEMMA 5. *Let Assumptions 2 and 3 hold. Then, for every $T > 0$ and $\epsilon \leq \min\left\{\bar{\epsilon}, \frac{(1-\gamma)\eta_\theta}{\sqrt{2L_\lambda}}\right\}$, the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ generated by VR-PDPG in Algorithm 1 satisfies that*

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] &\leq \frac{L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)}{\epsilon T} + \frac{2\eta_\theta}{\epsilon T} \sum_{t=0}^{T-1} \|e_{L,t}\| + \frac{\epsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} \\ &\quad + \frac{1}{\epsilon} \left[(1+C_0) \left(\frac{\ell_\theta \eta_\theta^2}{2} + 2L_{\theta,\lambda} \eta_\theta \gamma^H \right) + 2\eta_\mu M^2 \right], \end{aligned} \quad (31)$$

where $e_{L,t} := d_{L,t} - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)$.

Compared to its counterpart in the exact setting (see Lemma 3), Lemma 5 yields a different bound, which is due to the choice of the normalized update rule and the sample-based estimations. In particular, the extra term $e_{L,t}$ quantifies the approximation error of the policy gradient estimator $d_{L,t}$ with respect to $\nabla_\theta L(\lambda_H(\theta_t), \mu_t)$, and

the term that contains γ^H measures the truncation error $\nabla_{\theta}L(\lambda_H(\theta_t), \mu_t) - \nabla_{\theta}L(\lambda(\theta_t), \mu_t)$. Next, we decompose the boundedness on the Lagrangian into the optimality gap and constraint violation.

LEMMA 6. *Let Assumptions 1 and 3 hold, and the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ generated by VR-PDPG in Algorithm 1 satisfies that*

$$\frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] \leq C.$$

Then, the optimality gap and constraint violation can be respectively upper-bounded as follows

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_t))] &\leq C + \frac{\mu_0^2 - \mu_T^2}{2T\eta_{\mu}} + \eta_{\mu}M^2 + \frac{C_0M_{\lambda}\gamma^H}{1-\gamma} + \frac{2\eta_{\mu}M_{\lambda}^2\gamma^{2H}}{(1-\gamma)^2} + \frac{C_0M_{\lambda}}{T} \sum_{t=0}^{T-1} \|e_{\lambda,t}\| \\ &\quad + \frac{2\eta_{\mu}M_{\lambda}^2}{T} \sum_{t=0}^{T-1} \|e_{\lambda,t}\|^2, \end{aligned} \quad (32a)$$

$$\begin{aligned} \frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_{+} &\leq C + \frac{\max_{\mu \in U} \{(\mu_0 - \mu)^2 - (\mu_T - \mu)^2\}}{2T\eta_{\mu}} + \eta_{\mu}M^2 + \frac{C_0M_{\lambda}\gamma^H}{1-\gamma} + \frac{2\eta_{\mu}M_{\lambda}^2\gamma^{2H}}{(1-\gamma)^2} \\ &\quad + \frac{C_0M_{\lambda}}{T} \sum_{t=0}^{T-1} \|e_{\lambda,t}\| + \frac{2\eta_{\mu}M_{\lambda}^2}{T} \sum_{t=0}^{T-1} \|e_{\lambda,t}\|^2, \end{aligned} \quad (32b)$$

where $e_{\lambda,t} := \lambda_H(\theta_t) - \lambda_t$, for all $t \geq 0$.

The upper bounds in (32a) and (32b) contain the error term $e_{\lambda,t}$, which stems from the difference between the estimated occupancy measure λ_t and the true truncated occupancy measure $\lambda_H(\theta_t)$. We control the size of this error (see Appendix C.8) and then fine-tune the step-sizes to arrive at our main convergence result. For the brevity of the presentation, we use the symbol “ \lesssim ” to hide numerical constants and higher-order terms.

THEOREM 2 (SAMPLE-BASED SETTING). *Let Assumptions 1, 2, and 3 hold. For every $T \geq \frac{1}{2} \left(\frac{1-\gamma}{L_{\lambda}\epsilon} \right)^2$, we choose $\mu_0 = 0$, $\eta_{\theta} = T^{-1/2}$, $\eta_{\mu} = T^{-3/4}$, $\alpha_t = (t+1)^{-1/2}$, and $H = \frac{1}{2} \log_{1/\gamma} T$. Then, the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ generated by VR-PDPG in Algorithm 1 satisfies that*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(\lambda(\theta^*)) - f(\lambda(\theta_t)) \right] \lesssim \left[C_0M_{\lambda}\sqrt{C_{\lambda}} + \frac{C_0L_{\lambda}C_e + L_{\lambda}M^2}{1-\gamma} \right] \frac{1}{T^{1/4}}, \quad (33a)$$

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_{+} \lesssim \left[C_0^2 + C_0M_{\lambda}\sqrt{C_{\lambda}} + \frac{C_0L_{\lambda}C_e + L_{\lambda}M^2}{1-\gamma} \right] \frac{1}{T^{1/4}}, \quad (33b)$$

where $\mathbb{E}[x]_{+} := \mathbb{E}[\max\{x, 0\}]$, and the constants C_{λ}, C_e are defined as follows

$$C_{\lambda} := \frac{1 + C_w}{(1-\gamma)^2}, \quad (34a)$$

$$C_e := \frac{\sqrt{M_{\lambda}^2 [(1 + C_w)M_{\pi,1}^2 + (M_{\pi,1}^2 + M_{\pi,2})^2] + (M_{\pi,1}\ell_{\lambda})^2 C_{\lambda} + M_{\pi,1}\ell_{\lambda}\sqrt{C_{\lambda}}}}{(1-\gamma)^2}, \quad (34b)$$

$$C_w := 2H [(8H + 2)M_{\pi,1}^2 + 2M_{\pi,2}]. \quad (34c)$$

Since each θ_t is a random variable in the sample-based setting, the convergence established in Theorem 2 is in the expectation sense. In each iteration of Algorithm 1, $H = \mathcal{O}(\log T)$ new samples are collected. Thus, the $\mathcal{O}(T^{-1/4})$ convergence rate implies that a total of $\tilde{\mathcal{O}}(\epsilon^{-4})$ samples are required to achieve an $\mathcal{O}(\epsilon)$ optimality gap and constraint violation. In the special case of standard CMDPs, Theorem 2 matches the sample complexity of natural policy gradient-based primal-dual methods by [15, 6].

REMARK 1 (EXTENSION TO MULTIPLE CONSTRAINTS). *Although we present the single constraint setting only, it is straightforward to extend all of our results, including the zero constraint violation result in Section 5, to the general n -constraint setting, i.e.,*

$$\max_{\theta \in \mathbb{R}^K} f(\lambda(\theta)) \quad \text{s.t. } g_i(\lambda(\theta)) \geq 0, \quad \forall i = 1, \dots, n. \quad (35)$$

Firstly, under the Slater's condition, we can generalize the strong duality in Lemma 1 and show that the optimal dual variable $\mu^* \in \mathbb{R}^n$ satisfies that $0 \leq \mu_i^* \leq [f(\lambda(\theta^*)) - f(\lambda(\tilde{\theta}))]/\xi_i$, where ξ_i is the slackness for the i -th constraint. Hence, we can replace the dual feasible region $U = [0, C_0] = [0, 1 + (M_F - f(\lambda(\tilde{\theta})))/\xi]$ with the hypercube $\{\mu \in \mathbb{R}^n \mid 0 \leq \mu_i \leq (M_F - f(\lambda(\tilde{\theta})))/\xi_i\}$. Secondly, in Algorithm 1, we shall create separate variance-reduced estimators for each constraint (similar to $d_{g,t}$ and $g(\lambda_t)$). Finally, in the convergence proof, we replace the simple square of difference (e.g., $(\mu_{t+1} - \mu)^2$ in (90)) with the squared norm and replace the scalar multiplication (e.g., $\mu_t g(\lambda_t)$ in (91)) with the inner product. Then, it is easy to verify that the same convergence holds for the optimality gap and constraint violation. Note that the constraint violation in the n -constraint setting is defined as $\frac{1}{T} \|\sum_{t=0}^{T-1} -g(\lambda(\theta_t))\|_+$, where $g(\cdot)$ represents the vector of n constraints, and the operator $[\cdot]_+$ applies entrywisely.

5 Zero Constraint Violation

In safety-critical systems where violating the constraint may induce an unexpected cost, having a zero constraint violation is of great importance. Motivated by the recent works [29, 30], we will show that a zero constraint violation can be achieved while maintaining the same order of convergence rate for the optimality gap. We remark that the proof ideas are the same for both the exact setting and the sample-based setting, with the constant of the latter one being more involved. For the ease of presentation, we use the exacting setting as the showcase. Consider the pessimistic counterpart of (11):

$$\max_{\theta \in \mathbb{R}^K} f(\lambda(\theta)) \quad \text{s.t. } g_\delta(\lambda(\theta)) := g(\lambda(\theta)) - \delta \geq 0, \quad (36)$$

where $\delta > 0$ is the pessimistic term to be determined. Suppose that $\delta < \xi$, i.e., the pessimistic term is smaller than the strict feasibility of the Slater point, so that the Slater's condition still holds for problem (36). Consider applying algorithm (14) to problem (36), i.e.,

$$\begin{cases} \theta_{t+1} = \theta_t + \eta_{\theta,t} \nabla_{\theta} L_\delta(\lambda(\theta_t), \mu_t) \\ \mu_{t+1} = \mathcal{P}_U(\mu_t - \eta_{\mu} \nabla_{\mu} L_\delta(\lambda(\theta_t), \mu_t)) \end{cases}, \quad t = 0, 1, \dots, \quad (37)$$

where $L_\delta(\lambda(\theta), \mu) := f(\lambda(\theta)) + \mu g_\delta(\lambda(\theta))$ is the Lagrangian function for problem (36). The following theorem states that, with a carefully chosen δ , the constraint violation of the iterates generated by algorithm (37) will be zero for the original problem (11) when T is reasonably large, and the optimality gap still remains the same order as before. We direct the reader to Appendix E for the proof of Theorem 3.

THEOREM 3. *Let Assumptions 1, 2, and 3 hold. For $T > 0$, let $\delta = \mathcal{O}(T^{-1/3})$ satisfies that*

$$\frac{2M + (M + \xi)^2/2}{T^{2/3}} + \frac{2(1 + C_{0,\delta})\ell_{\theta}L_{\lambda}^2}{(1 - \gamma)^2 T^{1/3}} + \frac{2(M + \xi)^2 + C_{0,\delta}^2/2}{T^{1/3}} - \delta = 0,$$

where $C_{0,\delta} = 1 + \frac{M-f(\lambda(\tilde{\theta}))}{\xi-\delta}$. For $T > 0$ with $\delta < \xi$, let $\mu_0 = 0$, $\eta_{\theta,t} = \frac{1}{(1+C_0)\ell_\theta}$, and $\eta_\mu = T^{-2/3}$. Then, the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ generated by algorithm (37) with $U = [0, C_{0,\delta}]$ satisfies that

$$\frac{1}{T} \sum_{t=0}^{T-1} \left[f(\lambda(\theta^*)) - f(\lambda(\theta_t)) \right] \leq \frac{2\delta M}{\xi} + \frac{2M + (M + \xi)^2/2}{T^{2/3}} + \frac{2(1 + C_{0,\delta})\ell_\theta L_\lambda^2}{(1-\gamma)^2 T^{1/3}} + \frac{2(M + \xi)^2}{T^{1/3}}, \quad (38)$$

$$\frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ = 0. \quad (39)$$

Suppose that $f(\cdot)$ is σ -strongly concave w.r.t. λ . For $T > 0$, let $\delta = O(T^{-1/2})$ satisfies that

$$\frac{2M + C_{0,\delta}(M + \xi)}{\tilde{\epsilon}T} + \left(\frac{(M + \xi)^2}{\tilde{\epsilon}} + \frac{(M + \xi)^2 + C_{0,\delta}^2}{2} \right) \frac{1}{\sqrt{T}} - \delta = 0,$$

For $T > 0$ with $\delta < \xi$, let $\mu_0 = 0$, $\eta_{\theta,t} = \frac{1}{(1+C_0)\ell_\theta}$, and $\eta_\mu = T^{-1/2}$. Then, the sequence $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ generated by algorithm (37) with $U = [0, C_{0,\delta}]$ satisfies that

$$\frac{1}{T} \sum_{t=0}^{T-1} \left[f(\lambda(\theta^*)) - f(\lambda(\theta_t)) \right] \leq \frac{2\delta M}{\xi} + \frac{2M + C_{0,\delta}(M + \xi)}{\tilde{\epsilon}T} + \left(\frac{(M + \xi)^2}{\tilde{\epsilon}} + \frac{(M + \xi)^2}{2} \right) \frac{1}{\sqrt{T}}, \quad (40)$$

$$\frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ = 0. \quad (41)$$

The zero constraint violation presented in Theorem 3 applies to the exact setting, which achieves the same order of convergence for the optimality gap as in Theorem 1. We briefly introduce the ideas behind Theorem 3 here. Adding the pessimistic term δ shifts the optimal solution from θ^* to a new point θ_δ^* , where the difference $|f(\lambda(\theta^*)) - f(\lambda(\theta_\delta^*))|$ is bounded by $O(\delta)$. Since the orders of convergence rates are the same for optimality gap and constraint violation in Theorem 1, we can choose δ to have the same order and then offset the constraint violation for the pessimistic problem (36). As a result, the constraint violation becomes zero for the original problem (11), and the optimality gap preserves its previous order.

6 Numerical Experiment

We evaluate Algorithm 1 on a feasibility-constrained MDP problem, as introduced in Example 4. The experiment involves maximizing the objective while following a reference trajectory, i.e., expert demonstration, which is a reformulation of Example 4, as follows:

$$\max_{\lambda \in \Lambda} f(\lambda) = \langle r, \lambda \rangle \quad \text{s.t.} \quad g(\lambda) = \|\lambda - \lambda_e\| \leq d_0, \quad (42)$$

where λ_e is the occupancy measure of the reference trajectory, computed prior to conducting the experiments¹. The experiments are conducted in two grid world environments: 8×8 and 20×20 , where the agent starts at the upper left corner (marked as 'S' in Figure 1a) and aims to reach a goal located at the right bottom corner (marked as 'G' in Figure 1a). The policy of the agent is parameterized a neural network with a single hidden layer.

For each grid size, we conduct two different experiments with varying reference trajectories, denoted as traj1 and traj2 in Figures 1b and 1c, and present the episodic returns (Figures 1b and 2a) and constraint violations (Figures 1c and 2b). In Figures 1 and 2, the solid lines represent the mean, and the shaded areas indicate the 95% confidence level among 10 different seeds. For 8×8 grid world experiment (Figure 1), we use 12 different combinations of the hyperparameters, including the constraint violation $d_0 \in \{0.001, 0.005, 0.01, 0.05\}$ and the

¹Our experiment codes are available at <https://github.com/hyunin-lee/VR-PDPG>

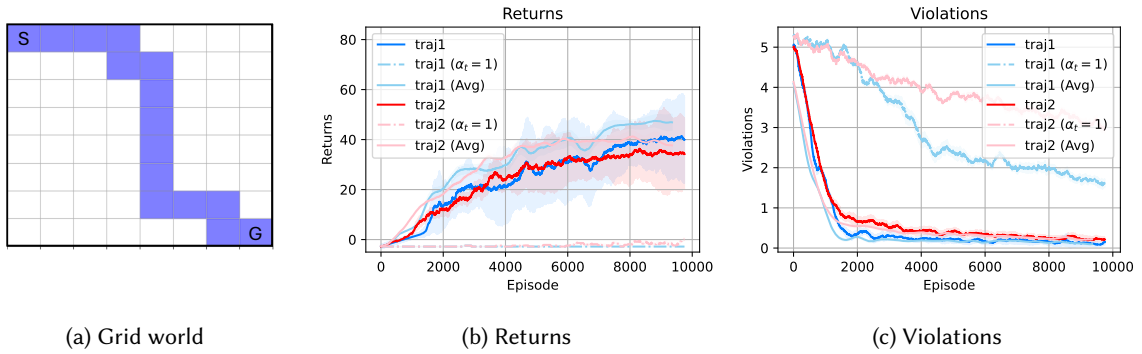


Fig. 1. Two instances of 8×8 grid world experiments under different reference trajectories. For $k \in \{1, 2\}$, traj_k , $\text{traj}_k(\alpha_t = 1)$, and $\text{traj}_k(\text{Avg})$ correspond to VR-PDPG, VR-PDPG with $\alpha_t = 1$, and VR-PDPG with $\alpha_t = 1$ and multi-trajectory estimation, respectively.

initial step size for primal parameter $\eta_\theta \in \{0.23, 0.24, 0.25\}$ to find a best return and take the average performance among different seeds, respectively. For 20×20 grid world experiment (Figure 2), we use 60 different combinations of the hyperparameters, including the initial step size for the dual parameter $\eta_\mu = \{0.01, 0.02, 0.03, 0.04, 0.05\}$, the momentum update parameter $\alpha \in \{0.03, 0.05, 0.07, 0.09\}$, and the constant violation $d_0 \in \{0.001, 0.005, 0.01\}$. We then select the best result out of 60 and compute the mean and variance among different seeds. We set the step-sizes for the primal parameter $\eta_{\theta,t}$ and the momentum parameter α_t to be time-invariant, i.e., $\eta_\theta = \eta_{\theta,t}$ and $\alpha = \alpha_t$ for all $t \geq 0$.

To demonstrate the superiority of VR-PDPG, which incorporates momentum terms, we compare it with the following two baselines: (i) VR-PDPG with $\alpha_t = 1, \forall t \geq 0$ (referred to as $\text{traj}_1(\alpha_t = 1)$ and $\text{traj}_2(\alpha_t = 1)$), which do not utilize the momentum acceleration for convergence. (ii) VR-PDPG with $\alpha_t = 1, \forall t \geq 0$ while employing multi-trajectory estimation (referred to as $\text{traj}_1(\text{Avg})$ and $\text{traj}_2(\text{Avg})$), which uses the mean of 100 trajectories with the same policy to estimate the gradient in (15) and occupancy measure in (17). We comment that our experiments focus the comparison on variants of VR-PDPG because, to the best of our knowledge, VR-PDPG is the first policy-based algorithm developed for solving concave CMDPs, and no existing methods are directly comparable. The most relevant prior work is [6]; however, as discussed in Section 1.1, their approach is based on occupancy measures and is fundamentally different from ours.

The results shown in Figures 1 and 2 demonstrate that the significant effectiveness of VR-PDPG. Specifically, comparing that trajectories traj_1 (or traj_2) and $\text{traj}_1(\alpha_t = 1)$ (or $\text{traj}_2(\alpha_t = 1)$) reveals that VR-PDPG guarantees a faster convergence and higher returns within a fixed number of episodes. Additionally, the high standard deviation observed when estimating from a single trajectory suggests that improvements may be achieved through multi-trajectory estimation since it helps reduce the variance. When comparing traj_1 (or traj_2) and $\text{traj}_1(\text{Avg})$ (or $\text{traj}_2(\text{Avg})$), we observe that averaging multiple trajectories does help improve the performance; however, the extent of this improvement is marginal. It is also worth noting that although the multi-trajectory estimation demonstrates some performance, it consumes 100 times more samples than our VR-PDPG algorithm, making it less attractive for practical purposes.

7 Conclusion

In this work, we propose primal-dual policy gradient (PDPG) algorithms to solve concave CMDP problems under the general soft-max policy with an over-parameterization assumption. We first consider the exact setting and

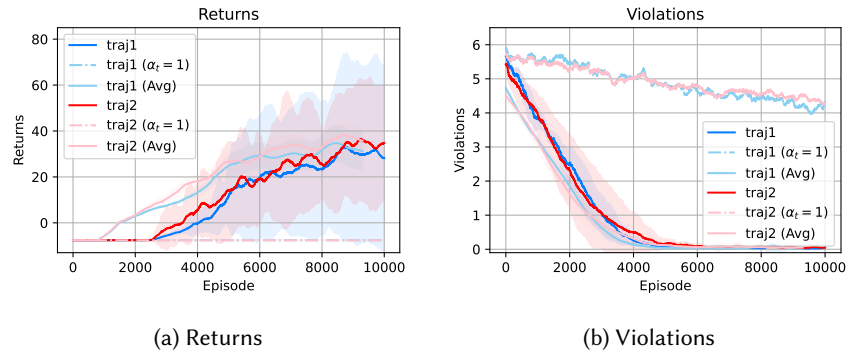


Fig. 2. Two instances of 20×20 grid world experiments under different reference trajectories.

show that the PDPG algorithm achieves an $\tilde{O}(T^{-1/3})$ rate of convergence in terms of the average optimality gap and constraint violation. This rate can be further improved to $\tilde{O}(T^{-1/2})$ when the objective is strongly concave in the state-action visitation distribution. Subsequently, in the sample-based setting, we propose the VR-PDPG algorithm and demonstrate that it can achieve an ϵ -global optimality using a total of $\tilde{O}(\epsilon^{-4})$ samples. In addition, by incorporating a diminishing pessimistic term into the original constraint, we also prove that PDPG and VR-PDPG can attain a zero constraint violation while maintaining the same convergence rate for the optimality gap.

One important direction of future work lies in establishing a lower bound for concave CMDP problems under a general soft-max parameterization to verify the optimality of our upper bounds. Furthermore, it is interesting to study whether geometric structures, such as entropy regularization [44, 29, 27] or policy mirror descent [40], can be exploited to accelerate the convergence. Finally, as opposed to the average-iterate convergence results in this paper, an alternative direction is to find last-iterate convergent algorithms for concave CMDPs [12], as such algorithms may exhibit a higher degree of stability and could directly output the final solution.

Acknowledgments

This work is an extension of [45], which only derived the convergence rate of Primal-Dual Policy Gradient algorithm (PDPG) in the exact setting. In this work, however, we study a more practical setting where the gradients of the Lagrangian need to be estimated from data. Under this sample-based setting, we incorporate variance reduction techniques and propose the Variance-Reduced PDPG algorithm (VR-PDPG), whose sample complexity is rigorously analyzed. We also conduct new experiments to validate the effectiveness of VR-PDPG.

This work was supported by the U. S. Army Research Laboratory and the U. S. Army Research Office under Grant W911NF2010219, Office of Naval Research under Grant N000142412673, NSF, and the UC Noyce Initiative.

References

- [1] P. Abbeel and A. Y. Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, 1.
- [2] N. Abe et al. 2010. Optimizing debt collections using constrained reinforcement learning. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 75–84.
- [3] J. Achiam, D. Held, A. Tamar, and P. Abbeel. 2017. Constrained policy optimization. In *International conference on machine learning*. PMLR, 22–31.
- [4] A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan. 2021. On the theory of policy gradient methods: optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22, 98, 1–76.

- [5] E. Altman. 1999. *Constrained Markov decision processes*. Vol. 7. CRC Press.
- [6] Q. Bai, A. S. Bedi, M. Agarwal, A. Koppel, and V. Aggarwal. 2023. Achieving zero constraint violation for concave utility constrained reinforcement learning via primal-dual approach. *Journal of Artificial Intelligence Research*, 78, 975–1016.
- [7] A. Barakat, I. Fatkhullin, and N. He. 2023. Reinforcement learning with general utilities: simpler variance reduction and large state-action space. In *International Conference on Machine Learning*. PMLR, 1753–1800.
- [8] V. S. Borkar. 2005. An actor-critic algorithm for constrained markov decision processes. *Systems & control letters*, 54, 3, 207–213.
- [9] S. Boyd, S. P. Boyd, and L. Vandenberghe. 2004. *Convex optimization*. Cambridge university press.
- [10] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone. 2017. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research*, 18, 1, 6070–6120.
- [11] A. Cutkosky and F. Orabona. 2019. Momentum-based variance reduction in non-convex sgd. *Advances in neural information processing systems*, 32.
- [12] D. Ding, C.-Y. Wei, K. Zhang, and A. Ribeiro. 2024. Last-iterate convergent policy gradient primal-dual methods for constrained mdps. *Advances in Neural Information Processing Systems*, 36.
- [13] D. Ding, K. Zhang, T. Basar, and M. Jovanovic. 2020. Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33, 8378–8390.
- [14] D. Ding, K. Zhang, T. Başar, and M. R. Jovanović. 2022. Convergence and optimality of policy gradient primal-dual method for constrained markov decision processes. In *2022 American Control Conference (ACC)*. IEEE, 2851–2856.
- [15] D. Ding, K. Zhang, J. Duan, T. Başar, and M. R. Jovanović. 2022. Convergence and sample complexity of natural policy gradient primal-dual methods for constrained mdps. arXiv preprint arXiv:2206.02346. (2022).
- [16] Y. Ding and J. Lavaei. 2023. Provably efficient primal-dual reinforcement learning for cmdps with non-stationary objectives and constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37, 7396–7404.
- [17] Y. Ding, J. Zhang, and J. Lavaei. 2022. On the global optimum convergence of momentum-based policy gradient. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1910–1934.
- [18] Y. Efroni, S. Mannor, and M. Pirotta. 2020. Exploration-exploitation in constrained mdps. arXiv preprint arXiv:2003.02189. (2020).
- [19] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine. 2019. Diversity is all you need: learning skills without a reward function. In *International Conference on Learning Representations*.
- [20] I. Fatkhullin, N. He, and Y. Hu. 2023. Stochastic optimization under hidden convexity. In *NeurIPS Workshop on Optimization for Machine Learning, 2023*.
- [21] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin. 2018. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64, 7, 2737–2752.
- [22] M. Geist, J. Perolat, M. Lauriere, R. Elie, S. Perrin, O. Bachem, R. Munos, and O. Pietquin. 2022. Concave utility reinforcement learning: the mean-field game viewpoint. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 489–497.
- [23] E. Hazan, S. Kakade, K. Singh, and A. Van Soest. 2019. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*. PMLR, 2681–2691.
- [24] J. Ho and S. Ermon. 2016. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29.
- [25] S. Kakade and J. Langford. 2002. Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer.
- [26] L. Lee, B. Eysenbach, E. Parisotto, E. Xing, S. Levine, and R. Salakhutdinov. 2019. Efficient exploration via state marginal matching. arXiv preprint arXiv:1906.05274. (2019).
- [27] T. Li, Z. Guan, S. Zou, T. Xu, Y. Liang, and G. Lan. 2024. Faster algorithm and sharper analysis for constrained markov decision process. *Operations Research Letters*, 54, 107107.
- [28] T. Lin, C. Jin, and M. Jordan. 2020. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*. PMLR, 6083–6093.
- [29] T. Liu, R. Zhou, D. Kalathil, P. R. Kumar, and C. Tian. 2021. Policy optimization for constrained mdps with provable fast global convergence. arXiv preprint arXiv:2111.00552. (2021).
- [30] T. Liu, R. Zhou, D. Kalathil, P. Kumar, and C. Tian. 2021. Learning policies with zero or bounded constraint violation for constrained mdps. *Advances in Neural Information Processing Systems*, 34.
- [31] S. Miryoosefi, K. Brantley, H. Daume III, M. Dudik, and R. E. Schapire. 2019. Reinforcement learning with convex constraints. *Advances in Neural Information Processing Systems*, 32.
- [32] M. Mutti, R. De Santi, P. De Bartolomeis, and M. Restelli. 2022. Challenging common assumptions in convex reinforcement learning. *Advances in Neural Information Processing Systems*, 35, 4489–4502.
- [33] M. Mutti, R. De Santi, and M. Restelli. 2022. The importance of non-markovianity in maximum state entropy exploration. In *International Conference on Machine Learning*. PMLR, 16223–16239.

- [34] T. Ni and M. Kamgarpour. 2024. A safe exploration approach to constrained markov decision processes. In *ICML 2024 Workshop: Foundations of Reinforcement Learning and Control—Connections and Perspectives*.
- [35] A. Rosenberg and Y. Mansour. 2019. Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*. PMLR, 5478–5486.
- [36] S. Schaal. 1996. Learning from demonstration. *Advances in neural information processing systems*, 9.
- [37] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. 1999. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12.
- [38] A. B. Tsybakov. 2008. *Introduction to Nonparametric Estimation*. (1st ed.). Springer Publishing Company, Incorporated. ISBN: 0387790519.
- [39] M. Wang. 2020. Randomized linear programming solves the markov decision problem in nearly linear (sometimes sublinear) time. *Mathematics of Operations Research*, 45, 2, 517–546.
- [40] L. Xiao. 2022. On the convergence rates of policy gradient methods. *Journal of Machine Learning Research*, 23, 282, 1–36.
- [41] P. Xu, F. Gao, and Q. Gu. 2019. Sample efficient policy gradient methods with recursive variance reduction. In *International Conference on Learning Representations*.
- [42] T. Xu, Y. Liang, and G. Lan. 2021. Crpo: a new approach for safe reinforcement learning with convergence guarantee. In *International Conference on Machine Learning*. PMLR, 11480–11491.
- [43] Q. Yang and M. T. Spaan. 2023. Cem: constrained entropy maximization for task-agnostic safe exploration. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37, 10798–10806.
- [44] D. Ying, Y. Ding, and J. Lavaei. 2022. A dual approach to constrained markov decision processes with entropy regularization. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1887–1909.
- [45] D. Ying, M. A. Guo, Y. Ding, J. Lavaei, and Z.-J. Shen. 2023. Policy-based primal-dual methods for convex constrained markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37, 10963–10971.
- [46] D. Ying, Y. Zhang, Y. Ding, A. Koppel, and J. Lavaei. 2024. Scalable primal-dual actor-critic method for safe multi-agent rl with general utilities. *Advances in Neural Information Processing Systems*, 36.
- [47] H. Yuan, X. Lian, J. Liu, and Y. Zhou. 2020. Stochastic recursive momentum for policy gradient methods. arXiv preprint arXiv:2003.04302. (2020).
- [48] T. Zahavy, B. O’Donoghue, G. Desjardins, and S. Singh. 2021. Reward is enough for convex mdps. *Advances in Neural Information Processing Systems*, 34.
- [49] J. Zhang, A. S. Bedi, M. Wang, and A. Koppel. 2022. Multi-agent reinforcement learning with general utilities via decentralized shadow reward actor-critic. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36, 9031–9039.
- [50] J. Zhang, A. Koppel, A. S. Bedi, C. Szepesvari, and M. Wang. 2020. Variational policy gradient method for reinforcement learning with general utilities. *Advances in Neural Information Processing Systems*, 33, 4572–4583.
- [51] J. Zhang, C. Ni, C. Szepesvari, M. Wang, et al. 2021. On the convergence and sample efficiency of variance-reduced policy gradient method. *Advances in Neural Information Processing Systems*, 34.
- [52] X. Zhang, K. Zhang, E. Miehling, and T. Basar. 2019. Non-cooperative inverse reinforcement learning. *Advances in Neural Information Processing Systems*, 32.
- [53] W. Zhou and W. Li. 2018. Safety-aware apprenticeship learning. In *International Conference on Computer Aided Verification*. Springer, 662–680.

A Discussions About Assumption 2

Below, we present a discussion on Assumption 2 regarding the parameterization. To leverage the hidden concavity of problem (11) with respect to λ , it is natural to assume that there exists some desirable correspondence between $\lambda(\theta)$ and θ . However, as mentioned in Section 4, requiring such correspondence to be one-to-one or invertible is too restrictive. Although we can show that a one-to-one correspondence indeed exists under the direct parameterization and that the inverse map is Lipschitz continuous as long as there is a universal positive lower bound for the state occupancy measure $d^\pi(\cdot)$ (see Lemma 7), this is not the case for many other parameterizations. The soft-max policy, defined as

$$\pi_\theta(a|s) = \frac{\exp(\theta_{sa})}{\sum_{a' \in \mathcal{A}} \exp(\theta_{sa'})}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad (43)$$

serves as a counterexample. For a fixed vector $\theta_0 \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$, consider the set of parameters

$$\{\theta \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|} \mid \theta_{sa} = (\theta_0)_{sa} + k, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad \forall k \in \mathbb{R}\}.$$

Then, it is clear that all parameters in the set correspond to the same policy π_{θ_0} . Thus, a one-to-one correspondence does not exist. This motivates Assumption 2, which only requires the local existence of a continuous inverse λ^{-1} . Assumption 2 is able to accommodate the soft-max policy defined in (43), as suggested by [51].

LEMMA 7 (LIPSCHITZ CONTINUITY OF λ^{-1} UNDER DIRECT PARAMETERIZATION). *Suppose that $d_0 := \min_{s \in \mathcal{S}, \pi \in \Pi} d^\pi(s) > 0^2$. For every two occupancy measures $\lambda_1, \lambda_2 \in \Lambda$, it holds that*

$$\|\lambda^{-1}(\lambda_1) - \lambda^{-1}(\lambda_2)\| \leq \frac{\sqrt{2(1 + |\mathcal{A}|)}}{d_0} \|\lambda_1 - \lambda_2\|,$$

where $\lambda^{-1}(\cdot)$ maps an occupancy measure to the corresponding policy, defined as

$$\pi(a|s) = [\lambda^{-1}(\lambda^\pi)]_{s,a} := \frac{\lambda^\pi(s, a)}{\sum_{a' \in \mathcal{A}} \lambda^\pi(s, a')}.$$

PROOF. Let $d_1(s) = \sum_a \lambda_1(s, a)$ and $d_2(s) = \sum_a \lambda_2(s, a)$. Then, it holds that

$$[\lambda^{-1}(\lambda_1)]_{s,a} - [\lambda^{-1}(\lambda_2)]_{s,a} = \frac{\lambda_1(s, a)}{d_1(s)} - \frac{\lambda_2(s, a)}{d_2(s)} = \frac{1}{d_1(s)} \left(\lambda_1(s, a) - \lambda_2(s, a) + \frac{d_2(s) - d_1(s)}{d_2(s)} \cdot \lambda_2(s, a) \right).$$

Therefore, one can compute that

$$\begin{aligned} \|\lambda^{-1}(\lambda_1) - \lambda^{-1}(\lambda_2)\|^2 &= \sum_{s,a} \left([\lambda^{-1}(\lambda_1)]_{s,a} - [\lambda^{-1}(\lambda_2)]_{s,a} \right)^2 \\ &= \sum_s \frac{1}{[d_1(s)]^2} \sum_a \left(\lambda_1(s, a) - \lambda_2(s, a) + [d_2(s) - d_1(s)] \cdot \frac{\lambda_2(s, a)}{d_2(s)} \right)^2 \\ &\leq \sum_s \frac{2}{[d_1(s)]^2} \left(\sum_a [\lambda_1(s, a) - \lambda_2(s, a)]^2 + \left([d_2(s) - d_1(s)] \cdot \frac{\lambda_2(s, a)}{d_2(s)} \right)^2 \right), \end{aligned} \quad (44)$$

where the last line follows from the inequality $(x + y)^2 \leq 2x^2 + 2y^2$. For the second term inside the summation, we have that

$$\begin{aligned} \sum_a \left([d_2(s) - d_1(s)] \cdot \frac{\lambda_2(s, a)}{d_2(s)} \right)^2 &= [d_2(s) - d_1(s)]^2 \cdot \left\| \frac{\lambda_2(s, \cdot)}{d_2(s)} \right\|^2 \\ &\stackrel{(i)}{\leq} [d_2(s) - d_1(s)]^2 \cdot \left\| \frac{\lambda_2(s, \cdot)}{d_2(s)} \right\|_1^2 \\ &= [d_2(s) - d_1(s)]^2 \\ &= \left[\sum_a \lambda_2(s, a) - \sum_a \lambda_1(s, a) \right]^2 \\ &\leq |\mathcal{A}| \cdot \sum_a [\lambda_2(s, a) - \lambda_1(s, a)]^2, \end{aligned} \quad (45)$$

²Since $d^\pi(s) \geq (1 - \gamma)\rho(s)$, this assumption is satisfied when there is an exploratory initial distribution, i.e., $\rho_0 := \min_{s \in \mathcal{S}} \rho(s) > 0$.

where (i) is due to $\|\cdot\| \leq \|\cdot\|_1$ and the last step follows from the Cauchy-Schwarz inequality. By substituting (45) into (44) and noting that $d^\pi(s) \geq d_0$, $\forall s \in \mathcal{S}, \pi \in \Pi$, it holds that

$$\begin{aligned} \|\lambda^{-1}(\lambda_1) - \lambda^{-1}(\lambda_2)\|^2 &\leq \sum_s \frac{2(1+|\mathcal{A}|)}{[d_1(s)]^2} \left(\sum_a [\lambda_1(s, a) - \lambda_2(s, a)]^2 \right) \\ &\leq \frac{2(1+|\mathcal{A}|)}{d_0^2} \sum_{s,a} [\lambda_1(s, a) - \lambda_2(s, a)]^2 \\ &\leq \frac{2(1+|\mathcal{A}|)}{d_0^2} \|\lambda_1 - \lambda_2\|^2. \end{aligned}$$

The proof is completed by taking the square root of both sides of the inequality. \square

B Supplementary Materials for Sections 2 and 3

This section documents the supplementary materials corresponding to Sections 2 and 3, including the proofs of Lemma 1 and auxiliary lemmas regarding the importance sampling weight.

B.1 Proof of Lemma 1

PROOF. We note that Λ , the set of all unnormalized state-action occupancy measures, is a convex polytope having the expression

$$\Lambda = \left\{ \lambda \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|} \mid \lambda \geq \mathbf{0}, \sum_a \lambda(s, a) = \rho(s) + \frac{\gamma}{1-\gamma} \sum_{s', a'} \mathbb{P}(s|s', a') \cdot \lambda(s', a'), \forall s \in \mathcal{S} \right\}. \quad (46)$$

Since $\text{cl}(\{\lambda(\theta) \mid \theta \in \mathbb{R}^K\}) = \Lambda$, the nonconcave problem (11) is equivalent to the convex problem (5). Therefore, the strong duality holds under Assumption 1 [9].

To prove (II), let $C \in \mathbb{R}$. For every $\mu \geq 0$ such that $\mathcal{D}(\mu) \leq C$, it holds that

$$C \geq D(\mu) \stackrel{(i)}{\geq} f(\lambda(\tilde{\theta})) + \mu g(\lambda(\tilde{\theta})) \stackrel{(ii)}{\geq} f(\lambda(\tilde{\theta})) + \mu \xi, \quad (47)$$

where (i) follows from the definition of $\mathcal{D}(\mu)$ and (ii) is due to Assumption 1.

Since $\xi > 0$, (47) gives rise to the bound $\mu \leq (C - f(\lambda(\tilde{\theta}))) / \xi$. Now, by letting $C = f(\lambda(\theta^*))$, it results from the strong duality that $\{\mu \geq 0 \mid D(\mu) \leq C = f(\lambda(\theta^*))\}$ becomes the set of optimal dual variables. This completes the proof. \square

B.2 Importance Sampling Weight

LEMMA 8. Given any $\theta \in \mathbb{R}^K$, let $\tau = \{(s_k, a_k)\}_{k=0}^{H-1}$ be a trajectory generated under π_θ . Then, for any target policy θ' , the IS weight $w(\tau \mid \theta', \theta)$ satisfies that

$$\mathbb{E}_{\tau \sim \pi_{\theta, \rho}} \left[w(\tau \mid \theta', \theta) \widehat{\lambda}(\tau) \right] = \lambda_H(\theta'), \quad \mathbb{E}_{\tau \sim \pi_{\theta, \rho}} \left[w(\tau \mid \theta', \theta) \widehat{d}(\tau, \theta', r) \right] = [\nabla_\theta \lambda_H(\theta')]^\top r, \quad (48)$$

where $\lambda_H(\theta) = \sum_{k=0}^{H-1} \gamma^k \mathbb{P}(s_k = s, a_k = a \mid \pi_\theta, s_0 \sim \rho)$.

PROOF. We first define the probability of having a trajectory τ with the length H under the policy π_θ as follows

$$\mathbb{P}(\tau \mid \pi_\theta) = \rho(s_0) \cdot \pi_\theta(a_0 \mid s_0) \cdot \prod_{k=1}^{H-1} \mathbb{P}(s_k \mid s_{k-1}, a_{k-1}) \pi_\theta(a_k \mid s_k). \quad (49)$$

By the definitions of $w(\tau|\theta', \theta)$ in (18) and $\mathbb{P}(\tau|\pi_\theta)$ in (48), we deduce that

$$w(\tau|\theta', \theta) = \prod_{k=0}^{H-1} \frac{\pi_{\theta'}(a_k|s_k)}{\pi_\theta(a_k|s_k)} = \frac{\mathbb{P}(\tau|\pi_{\theta'})}{\mathbb{P}(\tau|\pi_\theta)}. \quad (50)$$

Together with the definition of $\widehat{\lambda}(\tau)$ in (17), we have that

$$\begin{aligned} \mathbb{E}_{\tau \sim \pi_{\theta', \rho}} \left[w(\tau|\theta', \theta) \widehat{\lambda}(\tau) \right] &= \sum_{\tau} \mathbb{P}(\tau|\pi_{\theta'}) \cdot \left(\frac{\mathbb{P}(\tau|\pi_{\theta'})}{\mathbb{P}(\tau|\pi_\theta)} \cdot \sum_{k=0}^{H-1} \gamma^k \cdot \mathbb{1}(s_k, a_k) \right) \\ &= \sum_{\tau} \mathbb{P}(\tau|\pi_{\theta'}) \cdot \left(\sum_{k=0}^{H-1} \gamma^k \cdot \mathbb{1}(s_k, a_k) \right) = \lambda_H(\theta'). \end{aligned} \quad (51)$$

With similar steps, we can also derive that $\mathbb{E}_{\tau \sim \pi_{\theta', \rho}} \left[w(\tau|\theta', \theta) \widehat{d}(\tau, \theta', r) \right] = [\nabla_{\theta} \lambda_H(\theta')]^{\top} r$. \square

LEMMA 9. Let Assumption 2 hold, and suppose that $\{\theta_t\}_{t=0}^{T-1}$ is the primal sequence generated by Algorithm 1, i.e., $\theta_{t+1} = \theta_t + \eta_{\theta, t} d_{L, t}$ with $\eta_{\theta, t} = \eta_{\theta} / \|d_{L, t}\|$. Then, it holds that

(I) For every $t \geq 0$ and any trajectory τ of length H , the IS weight satisfies that $w(\tau|\theta_t, \theta_{t+1}) \leq \exp(2HM_{\pi, 1} \cdot \eta_{\theta})$.

(II) When $\eta_{\theta} = T^{-1/2}$ and $H = \frac{1}{2} \log_{1/\gamma} T$, there exists $W = o(1)$ such that $w(\tau|\theta_t, \theta_{t+1}) \leq W$. In addition, for any trajectory τ generated under $\pi_{\theta_{t+1}}$, it holds that

$$\mathbb{E}_{\tau \sim \pi_{\theta_{t+1}, \rho}} \left[w(\tau|\theta_t, \theta_{t+1}) \right] = 1, \quad (52)$$

$$\text{Var}_{\tau \sim \pi_{\theta_{t+1}, \rho}} \left(w(\tau|\theta_t, \theta_{t+1}) \right) \leq \frac{C_w}{T}, \quad (53)$$

where $C_w := 2H[(8H+2)M_{\pi, 1}^2 + 2M_{\pi, 2}]$.

PROOF. (I). Based on the policy parameterization defined in (10), for any policy $\pi_{\theta_t}, \pi_{\theta_{t+1}}$ and any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, it follows that

$$\begin{aligned} \frac{\pi_{\theta_t}(a|s)}{\pi_{\theta_{t+1}}(a|s)} &= \frac{\exp(\psi(\theta_t; s, a))}{\sum_{a' \in \mathcal{A}} \exp(\psi(\theta_t; s, a'))} \cdot \frac{\sum_{a' \in \mathcal{A}} \exp(\psi(\theta_{t+1}; s, a'))}{\exp(\psi(\theta_{t+1}; s, a))} \\ &\leq \exp(\psi(\theta_t; s, a) - \psi(\theta_{t+1}; s, a)) \cdot \max_{a' \in \mathcal{A}} \{ \exp(\psi(\theta_{t+1}; s, a') - \psi(\theta_t; s, a')) \} \\ &\leq \exp(2M_{\pi, 1} \|\theta_t - \theta_{t+1}\|) \\ &\leq \exp(2M_{\pi, 1} \cdot \eta_{\theta}), \end{aligned} \quad (54)$$

where the last two inequalities result from (24) in Assumption 2 and the update rule of θ_t , respectively. Lastly, from the definition of $w(\tau|\theta_t, \theta_{t+1})$ in (18) we have that

$$w(\tau|\theta_t, \theta_{t+1}) = \prod_{k=0}^{H-1} \frac{\pi_{\theta_t}(a_k|s_k)}{\pi_{\theta_{t+1}}(a_k|s_k)} \leq \exp(2HM_{\pi, 1} \cdot \eta_{\theta}). \quad (55)$$

(II). By (I), when $\eta_{\theta} = T^{-1/2}$ and $H = \frac{1}{2} \log_{1/\gamma} T$, it holds that

$$w(\tau|\theta_t, \theta_{t+1}) \leq \exp(2HM_{\pi, 1} \cdot \eta_{\theta}) = \exp(\log_{1/\gamma}(T) \cdot M_{\pi, 1} \cdot T^{-1/2}), \quad (56)$$

where the right hand side of (56) decays exponentially in T . Hence, there exists $W = o(1)$ such that $w(\tau|\theta_t, \theta_{t+1}) \leq W$. Then, using (50), we expand the expectation in (52) as follows

$$\mathbb{E}_{\tau \sim \pi_{\theta_{t+1}, \rho}} [w(\tau|\theta_t, \theta_{t+1})] = \sum_{\tau} \mathbb{P}(\tau|\pi_{\theta_{t+1}}) \cdot \frac{\mathbb{P}(\tau|\pi_{\theta_t})}{\mathbb{P}(\tau|\pi_{\theta_{t+1}})} = \sum_{\tau} \mathbb{P}(\tau|\pi_{\theta_t}) = 1. \quad (57)$$

Furthermore, it follows from the proof of Lemma B.1 in [41] that

$$\text{Var}_{\tau \sim \pi_{\theta_{t+1}, \rho}} (w(\tau|\theta_t, \theta_{t+1})) \leq \frac{H(W+1)[(8H+2)M_{\pi,1}^2 + 2M_{\pi,2}]}{T} \leq \frac{C_w}{T}, \quad (58)$$

where $C_w = 2H[(8H+2)M_{\pi,1}^2 + 2M_{\pi,2}]$, and the last inequality holds because $W = o(1)$. \square

C Supplementary Materials for Section 4

In this section, we present all proofs for the theoretical findings in Section 4, including our main results on convergence rate in terms of optimality gap and constraint violation.

C.1 Proof of Lemma 2

PROOF. (I). Being a polytope implies that Λ is closed and compact (see (46)). Since $f(\cdot)$ and $g(\cdot)$ are concave in $\Lambda \subset \mathbb{R}^{|\mathcal{S}| \cdot |\mathcal{A}|}$, they are also continuous. Thus, we have that $f(\cdot)$ and $g(\cdot)$ are bounded on Λ , where we denote the upper bound by M . Note that we use the same bound M for both $f(\cdot)$ and $g(\cdot)$ for simplicity. This completes the proof of (I).

(II). This part follows directly from the proof of Lemma 5.3 (iii) in [51], where the constant ℓ_θ is defined as follows

$$\ell_\theta = \frac{4\ell_\lambda M_{\pi,1}^2}{(1-\gamma)^4} + \frac{8M_{\pi,1}^2 M_\lambda}{(1-\gamma)^3} + \frac{2M_\lambda (M_{\pi,2} + M_{\pi,1}^2)}{(1-\gamma)^2}.$$

(III). See the proof of Lemma E.3 in [51], where $L_{\theta,\lambda}$ is defined as

$$L_{\theta,\lambda} = \sqrt{\frac{8\ell_\lambda^2 M_{\pi,1}^2}{(1-\gamma)^6} + 16M_{\pi,1}^2 \cdot M_\lambda^2 \left(\frac{(H+1)^2}{(1-\gamma)^2} + \frac{1}{(1-\gamma)^4} \right)}.$$

(IV). From the boundedness of $f(\cdot)$ and $g(\cdot)$ on Λ in part (I) and the fact that $\mu \in U = [0, C_0]$, we have that

$$|L(\lambda(\theta), \mu)| = |f(\lambda(\theta)) + \mu g(\lambda(\theta))| \leq |f(\lambda(\theta))| + \mu |g(\lambda(\theta))| \leq M + C_0 M, \quad \forall \mu \in U. \quad (59)$$

Similarly, as both $f(\lambda(\theta))$ and $g(\lambda(\theta))$ are ℓ_θ -smooth from part (II), we have that $L(\lambda(\theta), \mu)$ is $(1+C_0)\ell_\theta$ -smooth w.r.t. θ . In addition, by the properties in part (III), we further deduce that $\|\nabla_\theta L(\lambda(\theta), \mu) - \nabla_\theta L(\lambda_H(\theta), \mu)\| \leq (1+C_0)L_{\theta,\lambda}\gamma^H$. \square

C.2 Proof of Lemma 3

PROOF. We first prove (26) in the general concave setting. Note that computing the primal update in algorithm (14) is equivalent to solving the following sub-problem:

$$\begin{aligned} \theta_{t+1} &= \theta_t + \eta_{\theta,t} \nabla_\theta L(\lambda(\theta_t), \mu_t) \\ &= \arg\max_{\theta \in \mathbb{R}^K} \left\{ L(\lambda(\theta_t), \mu_t) + (\theta - \theta_t)^\top \nabla_\theta L(\lambda(\theta_t), \mu_t) - \frac{1}{2\eta_{\theta,t}} \|\theta - \theta_t\|^2 \right\} \\ &= \arg\max_{\theta \in \mathbb{R}^K} \left\{ L(\lambda(\theta_t), \mu_t) + (\theta - \theta_t)^\top \nabla_\theta L(\lambda(\theta_t), \mu_t) - \frac{(1+C_0)\ell_\theta}{2} \|\theta - \theta_t\|^2 \right\}. \end{aligned} \quad (60)$$

Since $L(\lambda(\theta), \mu)$ is $(1 + C_0)\ell_\theta$ -smooth w.r.t. θ by Lemma 2, we obtain for every $\theta \in \mathbb{R}^K$ that

$$\left| L(\lambda(\theta), \mu_t) - L(\lambda(\theta_t), \mu_t) - (\theta - \theta_t)^\top \nabla_\theta L(\lambda(\theta_t), \mu_t) \right| \leq \frac{(1 + C_0)\ell_\theta}{2} \|\theta - \theta_t\|^2. \quad (61)$$

Thus, the following ascent property holds

$$\begin{aligned} L(\lambda(\theta), \mu_t) &\geq L(\lambda(\theta_t), \mu_t) + (\theta - \theta_t)^\top \nabla_\theta L(\lambda(\theta_t), \mu_t) - \frac{(1 + C_0)\ell_\theta}{2} \|\theta - \theta_t\|^2 \\ &\geq L(\lambda(\theta), \mu_t) - (1 + C_0)\ell_\theta \|\theta - \theta_t\|^2, \end{aligned} \quad (62)$$

where we apply (61) two times respectively in the two inequalities above. On the basis of (60) and (62), it holds that

$$\begin{aligned} L(\lambda(\theta_{t+1}), \mu_t) &\geq L(\lambda(\theta_t), \mu_t) + (\theta_{t+1} - \theta_t)^\top \nabla_\theta L(\lambda(\theta_t), \mu_t) - \frac{(1 + C_0)\ell_\theta}{2} \|\theta_{t+1} - \theta_t\|^2 \\ &= \max_{\theta \in \mathbb{R}^K} \left\{ L(\lambda(\theta_t), \mu_t) + (\theta - \theta_t)^\top \nabla_\theta L(\lambda(\theta_t), \mu_t) - \frac{(1 + C_0)\ell_\theta}{2} \|\theta - \theta_t\|^2 \right\} \\ &\geq \max_{\theta \in \mathbb{R}^K} \left\{ L(\lambda(\theta), \mu_t) - (1 + C_0)\ell_\theta \|\theta - \theta_t\|^2 \right\}. \end{aligned} \quad (63)$$

Now, we leverage the local invertibility of $\lambda(\cdot)$ to lower-bound the right-hand side of (63). We first define that

$$\theta_{\varepsilon,t} := \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1} \left((1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*) \right). \quad (64)$$

According to Assumption 2, since $\varepsilon \leq \bar{\varepsilon}$, we have $(1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*) \in \mathcal{V}_{\lambda(\theta_t)}$. Thus, $\theta_{\varepsilon,t}$ is well-defined and $\theta_{\varepsilon,t} \in \mathcal{U}_{\theta_t}$. By definition, the composition of $\lambda : \Theta \rightarrow \Lambda$ and $\lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1} : \mathcal{V}_{\lambda(\theta_t)} \rightarrow \mathcal{U}_{\theta_t}$ is the identity map on $\mathcal{V}_{\lambda(\theta_t)}$. Together with the facts that $L(\cdot, \mu)$ is concave in λ , we have

$$\begin{aligned} L(\lambda(\theta_{\varepsilon,t}), \mu_t) &= L \left(\lambda \circ \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1} \left((1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*) \right), \mu_t \right) \\ &= L \left((1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*), \mu_t \right) \\ &\geq (1 - \varepsilon)L(\lambda(\theta_t), \mu_t) + \varepsilon L(\lambda(\theta^*), \mu_t). \end{aligned} \quad (65)$$

Additionally, the Lipschitz continuity of $\lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}(\cdot)$ implies that

$$\begin{aligned} \|\theta_{\varepsilon,t} - \theta_t\|^2 &= \left\| \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1} \left((1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*) \right) - \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}(\lambda(\theta_t)) \right\|^2 \\ &\leq L_\lambda^2 \left\| (1 - \varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*) - \lambda(\theta_t) \right\|^2 \\ &\leq \varepsilon^2 L_\lambda^2 \left\| \lambda(\theta^*) - \lambda(\theta_t) \right\|^2 \\ &\leq \frac{2\varepsilon^2 L_\lambda^2}{(1 - \gamma)^2}, \end{aligned} \quad (66)$$

where the last inequality uses the diameter of Λ , i.e., $\max_{\lambda_1, \lambda_2 \in \Lambda} \|\lambda_1 - \lambda_2\| \leq \frac{\sqrt{2}}{1 - \gamma}$. By substituting $\theta_{\varepsilon,t}$ into (63) and using the inequalities (65) and (66), it holds that

$$\begin{aligned} L(\lambda(\theta_{t+1}), \mu_t) &\geq \max_{\theta \in \mathbb{R}^K} \left\{ L(\lambda(\theta), \mu_t) - (1 + C_0)\ell_\theta \|\theta - \theta_t\|^2 \right\} \\ &\geq L(\lambda(\theta_{\varepsilon,t}), \mu_t) - (1 + C_0)\ell_\theta \|\theta_{\varepsilon,t} - \theta_t\|^2 \\ &\geq (1 - \varepsilon)L(\lambda(\theta_t), \mu_t) + \varepsilon L(\lambda(\theta^*), \mu_t) - \frac{2\varepsilon^2(1 + C_0)\ell_\theta L_\lambda^2}{(1 - \gamma)^2}, \end{aligned}$$

which implies that

$$L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{t+1}), \mu_t) \leq (1 - \varepsilon) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] + \frac{2\varepsilon^2(1 + C_0)\ell_\theta L_\lambda^2}{(1 - \gamma)^2}. \quad (67)$$

Consequently, one can obtain the recursion

$$\begin{aligned} & L(\lambda(\theta^*), \mu_{t+1}) - L(\lambda(\theta_{t+1}), \mu_{t+1}) \\ &= [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{t+1}), \mu_t)] + [L(\lambda(\theta^*), \mu_{t+1}) - L(\lambda(\theta^*), \mu_t)] \\ & \quad + [L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_{t+1}), \mu_{t+1})] \\ &\stackrel{(i)}{\leq} (1 - \varepsilon) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] + \frac{2\varepsilon^2(1 + C_0)\ell_\theta L_\lambda^2}{(1 - \gamma)^2} + [L(\lambda(\theta^*), \mu_{t+1}) - L(\lambda(\theta^*), \mu_t)] \\ & \quad + [L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_{t+1}), \mu_{t+1})] \\ &\stackrel{(ii)}{\leq} (1 - \varepsilon) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] + \frac{2\varepsilon^2(1 + C_0)\ell_\theta L_\lambda^2}{(1 - \gamma)^2} + 2\eta_\mu M^2, \end{aligned} \quad (68)$$

where we use (67) in (i). Step (ii) is due to the bound

$$\begin{aligned} |L(\lambda(\theta), \mu_t) - L(\lambda(\theta), \mu_{t+1})| &= |(\mu_t - \mu_{t+1}) \cdot g(\lambda(\theta))| \\ &= \left| \left[\mu_t - \mathcal{P}_U(\mu_t - \eta_\mu \nabla_\mu L(\lambda(\theta_t), \mu_t)) \right] \cdot g(\lambda(\theta)) \right| \\ &\leq |\eta_\mu \nabla_\mu L(\lambda(\theta_t), \mu_t) \cdot g(\lambda(\theta))| \\ &= |\eta_\mu g(\lambda(\theta_t)) \cdot g(\lambda(\theta))| \\ &\leq \eta_\mu M^2, \quad \forall \theta \in \mathbb{R}^K, \end{aligned} \quad (69)$$

where the two inequalities above result from the non-expansive property of the projection operator and the boundedness of $g(\lambda(\theta))$, i.e., $|g(\lambda(\theta))| \leq M$, respectively. Utilizing the recursion (68), we derive that

$$\begin{aligned} & L(\lambda(\theta^*), \mu_{t+1}) - L(\lambda(\theta_{t+1}), \mu_{t+1}) \\ &\leq (1 - \varepsilon) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] + 2\varepsilon^2(1 + C_0)\ell_\theta L_\lambda^2 + 2\eta_\mu M^2 \\ &\leq (1 - \varepsilon)^2 [L(\lambda(\theta^*), \mu_{t-1}) - L(\lambda(\theta_{t-1}), \mu_{t-1})] + (1 + 1 - \varepsilon) \left[\frac{2\varepsilon^2(1 + C_0)\ell_\theta L_\lambda^2}{(1 - \gamma)^2} + 2\eta_\mu M^2 \right] \\ &\leq (1 - \varepsilon)^{t+1} [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + \sum_{i=0}^t (1 - \varepsilon)^i \left[\frac{2\varepsilon^2(1 + C_0)\ell_\theta L_\lambda^2}{(1 - \gamma)^2} + 2\eta_\mu M^2 \right] \\ &= (1 - \varepsilon)^{t+1} [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + \frac{1 - (1 - \varepsilon)^{t+1}}{\varepsilon} \left[\frac{2\varepsilon^2(1 + C_0)\ell_\theta L_\lambda^2}{(1 - \gamma)^2} + 2\eta_\mu M^2 \right], \end{aligned}$$

which is equivalent to

$$\begin{aligned} & L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t) \\ &\leq (1 - \varepsilon)^t [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + (1 - (1 - \varepsilon)^t) \left(\frac{2\varepsilon(1 + C_0)\ell_\theta L_\lambda^2}{(1 - \gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon} \right), \quad \forall t \geq 0. \end{aligned}$$

Summing the above inequality over $t = 0, 1, \dots, T - 1$ yields that

$$\begin{aligned}
& \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] \\
& \leq \sum_{t=0}^{T-1} (1-\varepsilon)^t [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + (1 - (1-\varepsilon)^T) \left(\frac{2\varepsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon} \right) \\
& = \frac{1 - (1-\varepsilon)^T}{\varepsilon} [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + \left(T - \frac{1 - (1-\varepsilon)^T}{\varepsilon} \right) \left(\frac{2\varepsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon} \right) \\
& \leq \frac{1}{\varepsilon} [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + T \left(\frac{2\varepsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon} \right).
\end{aligned}$$

The proof of (26) is completed by dividing both sides of the inequality by T .

Now we proceed to the case where $f(\cdot)$ is σ -strongly concave with respect to λ . We begin with (63):

$$L(\lambda(\theta_{t+1}), \mu_t) \geq \max_{\theta \in \mathbb{R}^K} \{L(\lambda(\theta), \mu_t) - (1+C_0)\ell_\theta \|\theta - \theta_t\|^2\}, \quad (70)$$

For $\varepsilon \leq \bar{\varepsilon}$, we define $\theta_{\varepsilon,t} := \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}((1-\varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*))$ similarly as (64). Due to the fact that $L(\cdot, \mu)$ is σ -strongly concave in λ , which results from the σ -strongly concavity of $f(\cdot)$ and the concavity of $g(\cdot)$, we have that

$$\begin{aligned}
L(\lambda(\theta_{\varepsilon,t}), \mu_t) & = L\left(\lambda \circ \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}((1-\varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*)), \mu_t\right) \\
& = L((1-\varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*), \mu_t) \\
& \geq (1-\varepsilon)L(\lambda(\theta_t), \mu_t) + \varepsilon L(\lambda(\theta^*), \mu_t) + \frac{\sigma}{2}\varepsilon(1-\varepsilon) \|\lambda(\theta^*) - \lambda(\theta_t)\|^2.
\end{aligned} \quad (71)$$

By Assumption 2, the Lipschitz continuity of $\lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}$ implies that

$$\begin{aligned}
\|\theta_{\varepsilon,t} - \theta_t\|^2 & = \left\| \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}((1-\varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*)) - \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}(\lambda(\theta_t)) \right\|^2 \\
& \leq L_\lambda^2 \|(1-\varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*) - \lambda(\theta_t)\|^2 \\
& \leq \varepsilon^2 L_\lambda^2 \|\lambda(\theta^*) - \lambda(\theta_t)\|^2.
\end{aligned} \quad (72)$$

Substitute $\theta_{\varepsilon,t}$ into the right-hand side of (70), we have that

$$\begin{aligned}
& L(\lambda(\theta_{t+1}), \mu_t) \\
& \geq \max_{\theta \in \mathbb{R}^K} \{L(\lambda(\theta), \mu_t) - (1+C_0)\ell_\theta \|\theta - \theta_t\|^2\} \\
& \geq \max_{0 \leq \varepsilon \leq \bar{\varepsilon}} \left\{ L(\lambda(\theta_{\varepsilon,t}), \mu_t) - (1+C_0)\ell_\theta \|\theta_{\varepsilon,t} - \theta_t\|^2 \right\} \\
& \geq \max_{0 \leq \varepsilon \leq \bar{\varepsilon}} \left\{ (1-\varepsilon)L(\lambda(\theta_t), \mu_t) + \varepsilon L(\lambda(\theta^*), \mu_t) + \left(\frac{\sigma}{2}\varepsilon(1-\varepsilon) - \varepsilon^2(1+C_0)\ell_\theta L_\lambda^2 \right) \|\lambda(\theta^*) - \lambda(\theta_t)\|^2 \right\},
\end{aligned}$$

where we use (71) and (72) in the last inequality. Consequently,

$$\begin{aligned}
L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{t+1}), \mu_t) & \leq \min_{0 \leq \varepsilon \leq \bar{\varepsilon}} \left\{ (1-\varepsilon)[L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] \right. \\
& \quad \left. - \left(\frac{\sigma\varepsilon(1-\varepsilon)}{2} - \varepsilon^2(1+C_0)\ell_\theta L_\lambda^2 \right) \|\lambda(\theta^*) - \lambda(\theta_t)\|^2 \right\}.
\end{aligned} \quad (73)$$

We note that

$$\frac{\sigma\varepsilon(1-\varepsilon)}{2} - \varepsilon^2(1+C_0)\ell_\theta L_\lambda^2 \geq 0, \quad \text{if } 0 \leq \varepsilon \leq \frac{\sigma}{\sigma + 2L_\lambda^2(1+C_0)\ell_\theta}. \quad (74)$$

By letting $\tilde{\varepsilon} := \min\{\bar{\varepsilon}, \sigma/(\sigma + 2L_\lambda^2(1+C_0)\ell_\theta)\} \leq \bar{\varepsilon}$, it follows from (73) that

$$\begin{aligned} & L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{t+1}), \mu_t) \\ & \leq (1 - \tilde{\varepsilon}) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] - \left(\frac{\sigma\tilde{\varepsilon}(1-\tilde{\varepsilon})}{2} - \tilde{\varepsilon}^2(1+C_0)\ell_\theta L_\lambda^2 \right) \|\lambda(\theta^*) - \lambda(\theta_t)\|^2 \\ & \leq (1 - \tilde{\varepsilon}) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)], \end{aligned} \quad (75)$$

where the second inequality results from (74). We rearrange the terms in (75) to obtain

$$L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{t+1}), \mu_t) \leq \frac{1-\tilde{\varepsilon}}{\tilde{\varepsilon}} [L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_t), \mu_t)],$$

which implies that

$$\begin{aligned} L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t) &= [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{t+1}), \mu_t)] + [L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_t), \mu_t)] \\ &\leq \left(\frac{1-\tilde{\varepsilon}}{\tilde{\varepsilon}} + 1 \right) [L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_t), \mu_t)] \\ &= \frac{1}{\tilde{\varepsilon}} [L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_t), \mu_t)]. \end{aligned}$$

Summing it over $t = 0, \dots, T-1$, we have that

$$\begin{aligned} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] &\leq \frac{1}{\tilde{\varepsilon}} \sum_{t=0}^{T-1} [L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_t), \mu_t)] \\ &= \frac{1}{\tilde{\varepsilon}} \left(L(\lambda(\theta_T), \mu_{T-1}) - L(\lambda(\theta_0), \mu_0) + \sum_{t=0}^{T-2} [L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_{t+1}), \mu_{t+1})] \right) \\ &\stackrel{(i)}{\leq} \frac{1}{\tilde{\varepsilon}} [L(\lambda(\theta_T), \mu_{T-1}) - L(\lambda(\theta_0), \mu_0) + (T-1)\eta_\mu M^2], \end{aligned}$$

where we use (69) to bound the difference $L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_{t+1}), \mu_{t+1})$ in (i). The proof is completed by dividing both sides of the inequality by T . \square

C.3 Proof of Lemma 4

PROOF. Lemma 4 is a simplified version of Lemma 6, and its proof simply follows from that of Lemma 6 in Appendix C.6. The only subtlety is that (88) and (90) become

$$\begin{aligned} (\mu_{t+1} - \mu)^2 &\leq (\mu_t - \mu)^2 - 2\eta_\mu(\mu_t - \mu) \cdot g(\lambda(\theta_t)) + (\eta_\mu g(\lambda(\theta_t)))^2 \\ &\leq (\mu_t - \mu)^2 - 2\eta_\mu(\mu_t - \mu) \cdot g(\lambda(\theta_t)) + \eta_\mu^2 M^2. \end{aligned} \quad (76)$$

Therefore, we can show that in the absence of approximation errors, the optimality gap (32a) and constraint violation (32b) reduce to (28a) and (28b), respectively. \square

C.4 Proof of Theorem 1 (Exact Setting)

PROOF. The proof of Theorem 1 relies on Lemmas 3 and 4. Combining (26) in Lemma 3 and (28a) in Lemma 4, we have that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} f(\lambda(\theta^*)) - f(\lambda(\theta_t)) &\leq \frac{L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)}{\varepsilon T} + \frac{2\varepsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon} + \frac{\mu_0^2 - \mu_T^2}{2T\eta_\mu} + \frac{\eta_\mu M^2}{2} \\ &\stackrel{(i)}{\leq} \frac{2M}{\varepsilon T} + \frac{2\varepsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon} + \frac{\mu_0^2 - \mu_T^2}{2T\eta_\mu} + \frac{\eta_\mu M^2}{2} \\ &\stackrel{(ii)}{\leq} \frac{2M + M^2/2}{T^{2/3}} + \frac{2(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2 T^{1/3}} + \frac{2M^2}{T^{1/3}}, \end{aligned}$$

where step (i) follows from the fact that $\mu_0 = 0$ specified in Theorem 1, and thereby $L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0) = f(\lambda(\theta^*)) - f(\lambda(\theta_0)) \leq 2M$ from Lemma 2. Since $\mu_0^2 - \mu_T^2 \leq 0$, the third term in step (i) can be dropped. Then, step (ii) is obtained by specifying $\varepsilon = T^{-1/3} \leq \bar{\varepsilon}$ and $\eta_\mu = T^{-2/3}$. This completes the proof of the optimality gap in (29a).

Similarly, from (28b) in Lemma 4 and Lemma 3, it holds that

$$\begin{aligned} \frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ &\leq \frac{L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)}{\varepsilon T} + \frac{2\varepsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon} \\ &\quad + \frac{\max_{\mu \in U} \{(\mu_0 - \mu)^2 - (\mu_T - \mu)^2\}}{2T\eta_\mu} + \frac{\eta_\mu M^2}{2} \\ &\stackrel{(i)}{\leq} \frac{2M}{\varepsilon T} + \frac{2\varepsilon(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2} + \frac{2\eta_\mu M^2}{\varepsilon} + \frac{C_0^2}{2T\eta_\mu} + \frac{\eta_\mu M^2}{2} \\ &\stackrel{(ii)}{\leq} \frac{2M + M^2/2}{T^{2/3}} + \frac{2(1+C_0)\ell_\theta L_\lambda^2}{(1-\gamma)^2 T^{1/3}} + \frac{2M^2 + C_0^2/2}{T^{1/3}}, \end{aligned}$$

where in step (i), we apply the inequality $\max_{\mu \in U} \{(\mu_0 - \mu)^2 - (\mu_T - \mu)^2\} = \max_{\mu \in U} \{(2\mu - \mu_T)\mu_T\} \leq C_0^2$. The last step (ii) is from letting $\varepsilon = T^{-1/3}$ and $\eta_\mu = T^{-2/3}$. This completes the proof of the constraint violation in (29b).

When $f(\cdot)$ is σ -strongly concave with respect to λ , the optimality gap resulted from Lemmas 3 and 4 can be similarly derived as follows

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} f(\lambda(\theta^*)) - f(\lambda(\theta_t)) &\leq \frac{L(\lambda(\theta_T), \mu_{T-1}) - L(\lambda(\theta_0), \mu_0)}{\tilde{\varepsilon} T} + \frac{\eta_\mu M^2}{\tilde{\varepsilon}} + \frac{\mu_0^2 - \mu_T^2}{2T\eta_\mu} + \frac{\eta_\mu M^2}{2} \\ &\leq \frac{(2+C_0)M}{\tilde{\varepsilon} T} + \left(\frac{M^2}{\tilde{\varepsilon}} + \frac{M^2}{2} \right) \frac{1}{\sqrt{T}}, \end{aligned} \tag{77}$$

where we specify $\eta_\mu = T^{-1/2}$ and apply the inequality $L(\lambda(\theta_T), \mu_{T-1}) - L(\lambda(\theta_0), \mu_0) = L(\lambda(\theta_T), \mu_{T-1}) - f(\lambda(\theta_0)) \leq (1+C_0)M + M = (2+C_0)M$. Finally, we show that the constraint violation satisfies that

$$\begin{aligned} \frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ &\leq \frac{L(\lambda(\theta_T), \mu_{T-1}) - L(\lambda(\theta_0), \mu_0)}{\tilde{\varepsilon} T} + \frac{\eta_\mu M^2}{\tilde{\varepsilon}} + \frac{\max_{\mu \in U} \{(\mu_0 - \mu)^2 - (\mu_T - \mu)^2\}}{2T\eta_\mu} + \frac{\eta_\mu M^2}{2} \\ &\leq \frac{(2+C_0)M}{\tilde{\varepsilon} T} + \left(\frac{M^2}{\tilde{\varepsilon}} + \frac{M^2 + C_0^2}{2} \right) \frac{1}{\sqrt{T}}. \end{aligned}$$

This completes the proof of Theorem 1. \square

C.5 Proof of Lemma 5

PROOF. We begin with showing an ascent property of the Lagrangian in the following lemma.

LEMMA 10. Suppose that Assumptions 2 and 3 hold true, and let $\{(\theta_t, \mu_t)\}_{t=0}^{T-1}$ be the sequence generated by Algorithm 1. Then, for every $0 \leq t \leq T-1$, it holds that

$$L(\lambda(\theta_{t+1}), \mu_t) \geq L(\lambda(\theta_t), \mu_t) + \eta_\theta \|\nabla_\theta L(\lambda(\theta_t), \mu_t)\| - 2\eta_\theta \|e_{L,t}\| - \frac{(1+C_0)\ell_\theta\eta_\theta^2}{2} - 2\eta_\theta(1+C_0)L_{\theta,\lambda}Y^H,$$

where $e_{L,t} := d_{L,t} - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)$.

PROOF. By the smoothness of $L(\lambda(\theta), \mu)$ in θ from Lemma 2, we have that

$$\begin{aligned} L(\lambda(\theta_{t+1}), \mu_t) &\geq L(\lambda(\theta_t), \mu_t) + \langle \nabla_\theta L(\lambda(\theta_t), \mu_t), \theta_{t+1} - \theta_t \rangle - \frac{(1+C_0)\ell_\theta}{2} \|\theta_{t+1} - \theta_t\|^2 \\ &\stackrel{(i)}{=} L(\lambda(\theta_t), \mu_t) + \eta_\theta \left\langle \nabla_\theta L(\lambda(\theta_t), \mu_t), \frac{d_{L,t}}{\|d_{L,t}\|} \right\rangle - \frac{(1+C_0)\ell_\theta}{2} \eta_\theta^2 \\ &= L(\lambda(\theta_t), \mu_t) + \eta_\theta \left\langle \nabla_\theta L(\lambda_H(\theta_t), \mu_t), \frac{d_{L,t}}{\|d_{L,t}\|} \right\rangle - \frac{(1+C_0)\ell_\theta}{2} \eta_\theta^2 \\ &\quad + \eta_\theta \left\langle \nabla_\theta L(\lambda(\theta_t), \mu_t) - \nabla_\theta L(\lambda_H(\theta_t), \mu_t), \frac{d_{L,t}}{\|d_{L,t}\|} \right\rangle \\ &\stackrel{(ii)}{\geq} L(\lambda(\theta_t), \mu_t) + \eta_\theta \left\langle \nabla_\theta L(\lambda_H(\theta_t), \mu_t), \frac{d_{L,t}}{\|d_{L,t}\|} \right\rangle - \frac{(1+C_0)\ell_\theta\eta_\theta^2}{2} - \eta_\theta(1+C_0)L_{\theta,\lambda}Y^H, \end{aligned} \tag{78}$$

where step (i) is due to the update rule of θ in Algorithm 1, and step (ii) is due to (25) in Lemma 2. We then bound the second term in the right-hand side of (78):

$$\begin{aligned} &\left\langle \nabla_\theta L(\lambda_H(\theta_t), \mu_t), \frac{d_{L,t}}{\|d_{L,t}\|} \right\rangle \\ &= \left\langle d_{L,t} - [d_{L,t} - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)], \frac{d_{L,t}}{\|d_{L,t}\|} \right\rangle \\ &\geq \|d_{L,t}\| - \|d_{L,t} - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)\| \\ &\stackrel{(i)}{\geq} \|\nabla_\theta L(\lambda_H(\theta_t), \mu_t)\| - 2\|d_{L,t} - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)\| \\ &\geq \|\nabla_\theta L(\lambda(\theta_t), \mu_t)\| - \|\nabla_\theta L(\lambda(\theta_t), \mu_t) - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)\| - 2\|d_{L,t} - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)\| \\ &\stackrel{(ii)}{\geq} \|\nabla_\theta L(\lambda(\theta_t), \mu_t)\| - (1+C_0)L_{\theta,\lambda}Y^H - 2\|d_{L,t} - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)\|, \end{aligned} \tag{79}$$

where step (i) applies the triangular inequality to the term $d_{L,t} = \nabla_\theta L(\lambda_H(\theta_t), \mu_t) + [d_{L,t} - \nabla_\theta L(\lambda_H(\theta_t), \mu_t)]$, and step (ii) is again due to (25) in Lemma 2. The proof of Lemma 10 is completed by substituting (79) back into (78). \square

Now, we finalize the proof of Lemma 5. Recall that $\theta_{\varepsilon,t} := \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}((1-\varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*))$ from (64), which is always well-defined as long as $\varepsilon \leq \bar{\varepsilon}$ (see Assumption 2). Using the fact that $L(\lambda(\theta), \mu)$ is $(1+C_0)\ell_\theta$ -smooth with respect to θ from Lemma 2, we have that

$$\begin{aligned} L(\lambda(\theta_t), \mu_t) &\geq L(\lambda(\theta_{\varepsilon,t}), \mu_t) - \langle \nabla_\theta L(\lambda(\theta_t), \mu_t), \theta_{\varepsilon,t} - \theta_t \rangle - \frac{(1+C_0)\ell_\theta}{2} \|\theta_{\varepsilon,t} - \theta_t\|^2 \\ &\geq (1-\varepsilon)L(\lambda(\theta_t), \mu_t) + \varepsilon L(\lambda(\theta^*), \mu_t) - \langle \nabla_\theta L(\lambda(\theta_t), \mu_t), \theta_{\varepsilon,t} - \theta_t \rangle - \frac{(1+C_0)\ell_\theta}{2} \|\theta_{\varepsilon,t} - \theta_t\|^2, \end{aligned} \quad (80)$$

where the last line is due to the concavity of $L(\lambda(\theta), \mu)$ in λ . By the local inverse assumption and L_λ -Lipschitz continuity of $\lambda_{\mathcal{V}_{\lambda(\theta)}}^{-1}(\cdot)$ for all $\theta \in \mathbb{R}^K$ (see Assumption 2), it holds that

$$\|\theta_{\varepsilon,t} - \theta_t\| = \left\| \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}((1-\varepsilon)\lambda(\theta_t) + \varepsilon\lambda(\theta^*)) - \lambda_{\mathcal{V}_{\lambda(\theta_t)}}^{-1}(\lambda(\theta_t)) \right\| \leq L_\lambda \cdot \varepsilon \|\lambda(\theta^*) - \lambda(\theta_t)\| \leq \frac{\sqrt{2}\varepsilon L_\lambda}{1-\gamma}. \quad (81)$$

Then, we apply the ascent property in Lemma 10 and substitute (81) into (80) to obtain the following

$$\begin{aligned} L(\lambda(\theta_{t+1}), \mu_t) &\geq L(\lambda(\theta_t), \mu_t) + \eta_\theta \|\nabla_\theta L(\lambda(\theta_t), \mu_t)\| - 2\eta_\theta \|e_{L,t}\| - \frac{(1+C_0)\ell_\theta\eta_\theta^2}{2} - 2\eta_\theta(1+C_0)L_{\theta,\lambda}Y^H \\ &\geq (1-\varepsilon)L(\lambda(\theta_t), \mu_t) + \varepsilon L(\lambda(\theta^*), \mu_t) + \eta_\theta \|\nabla_\theta L(\lambda(\theta_t), \mu_t)\| - 2\eta_\theta \|e_{L,t}\| \\ &\quad - \frac{(1+C_0)\ell_\theta\eta_\theta^2}{2} - 2\eta_\theta(1+C_0)L_{\theta,\lambda}Y^H - \langle \nabla_\theta L(\lambda(\theta_t), \mu_t), \theta_{\varepsilon,t} - \theta_t \rangle - \frac{(1+C_0)\ell_\theta L_\lambda^2 \varepsilon^2}{(1-\gamma)^2} \\ &\stackrel{(i)}{\geq} (1-\varepsilon)L(\lambda(\theta_t), \mu_t) + \varepsilon L(\lambda(\theta^*), \mu_t) + \left(\eta_\theta - \frac{\sqrt{2}\varepsilon L_\lambda}{1-\gamma} \right) \|\nabla_\theta L(\lambda(\theta_t), \mu_t)\| \\ &\quad - 2\eta_\theta \|e_{L,t}\| - (1+C_0) \left(\frac{\ell_\theta\eta_\theta^2}{2} + 2\eta_\theta L_{\theta,\lambda}Y^H + \frac{\ell_\theta L_\lambda^2 \varepsilon^2}{(1-\gamma)^2} \right), \end{aligned} \quad (82)$$

where in step (i) we use (81) again to lower-bound the term $\langle \nabla_\theta L(\lambda(\theta_t), \mu_t), \theta_{\varepsilon,t} - \theta_t \rangle$. By subtracting $L(\lambda(\theta^*), \mu_t)$ from both sides of (82) and rearranging the terms, we derive that

$$\begin{aligned} L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{t+1}), \mu_t) &\leq (1-\varepsilon) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] - \left(\eta_\theta - \frac{\sqrt{2}\varepsilon L_\lambda}{1-\gamma} \right) \|\nabla_\theta L(\lambda(\theta_t), \mu_t)\| \\ &\quad + 2\eta_\theta \|e_{L,t}\| + (1+C_0) \left(\frac{\ell_\theta\eta_\theta^2}{2} + 2\eta_\theta L_{\theta,\lambda}Y^H + \frac{\ell_\theta L_\lambda^2 \varepsilon^2}{(1-\gamma)^2} \right). \end{aligned} \quad (83)$$

Note that by the dual variable update, we have that

$$|L(\lambda(\theta), \mu_{t+1}) - L(\lambda(\theta), \mu_t)| = |(\mu_{t+1} - \mu_t) \cdot g(\lambda(\theta))| \leq \eta_\mu M^2, \quad \forall \theta \in \mathbb{R}^K.$$

Together with (83), we deduce that

$$\begin{aligned}
L(\lambda(\theta^*), \mu_{t+1}) - L(\lambda(\theta_{t+1}), \mu_{t+1}) &= L(\lambda(\theta^*), \mu_{t+1}) - L(\lambda(\theta^*), \mu_t) + L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_{t+1}), \mu_t) \\
&\quad + L(\lambda(\theta_{t+1}), \mu_t) - L(\lambda(\theta_{t+1}), \mu_{t+1}) \\
&\leq (1 - \varepsilon) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] - \left(\eta_\theta - \frac{\sqrt{2\varepsilon}L_\lambda}{1 - \gamma} \right) \|\nabla_\theta L(\lambda(\theta_t), \mu_t)\| + 2\eta_\theta \|e_{L,t}\| \\
&\quad + (1 + C_0) \left(\frac{\ell_\theta \eta_\theta^2}{2} + 2\eta_\theta L_{\theta,\lambda} \gamma^H + \frac{\ell_\theta L_\lambda^2 \varepsilon^2}{(1 - \gamma)^2} \right) + 2\eta_\mu M^2.
\end{aligned}$$

When we choose $\eta_\theta - \frac{\sqrt{2\varepsilon}L_\lambda}{(1-\gamma)} \geq 0$, i.e., $\varepsilon \leq \frac{(1-\gamma)\eta_\theta}{\sqrt{2}L_\lambda}$, then it holds that

$$\begin{aligned}
L(\lambda(\theta^*), \mu_{t+1}) - L(\lambda(\theta_{t+1}), \mu_{t+1}) &\leq (1 - \varepsilon) [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] + 2\eta_\theta \|e_{L,t}\| \\
&\quad + (1 + C_0) \left(\frac{\ell_\theta \eta_\theta^2}{2} + 2\eta_\theta L_{\theta,\lambda} \gamma^H + \frac{\ell_\theta L_\lambda^2 \varepsilon^2}{(1 - \gamma)^2} \right) + 2\eta_\mu M^2.
\end{aligned}$$

We recursively apply the above inequality to obtain that

$$\begin{aligned}
L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t) &\leq (1 - \varepsilon)^t [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + \sum_{i=0}^{t-1} (1 - \varepsilon)^i \cdot 2\eta_\theta \|e_{L,t-1-i}\| \\
&\quad + \sum_{i=0}^{t-1} (1 - \varepsilon)^i \left[(1 + C_0) \left(\frac{\ell_\theta \eta_\theta^2}{2} + 2\eta_\theta L_{\theta,\lambda} \gamma^H + \frac{\ell_\theta L_\lambda^2 \varepsilon^2}{(1 - \gamma)^2} \right) + 2\eta_\mu M^2 \right] \\
&\leq (1 - \varepsilon)^t [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + \sum_{i=0}^{t-1} (1 - \varepsilon)^i \cdot 2\eta_\theta \|e_{L,t-1-i}\| \\
&\quad + \frac{1}{\varepsilon} \left[(1 + C_0) \left(\frac{\ell_\theta \eta_\theta^2}{2} + 2\eta_\theta L_{\theta,\lambda} \gamma^H + \frac{\ell_\theta L_\lambda^2 \varepsilon^2}{(1 - \gamma)^2} \right) + 2\eta_\mu M^2 \right],
\end{aligned} \tag{84}$$

where we use $\sum_{i=0}^{t-1} (1 - \varepsilon)^i = \frac{1 - (1 - \varepsilon)^t}{\varepsilon} \leq \frac{1}{\varepsilon}$ in the last line. Summing the inequality (84) over $t = 0, \dots, T - 1$ yields that

$$\begin{aligned}
\sum_{t=0}^{T-1} L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t) &\leq \frac{1}{\varepsilon} [L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)] + \frac{2\eta_\theta}{\varepsilon} \sum_{t=0}^{T-1} \|e_{L,t}\| \\
&\quad + \frac{T}{\varepsilon} \left[(1 + C_0) \left(\frac{\ell_\theta \eta_\theta^2}{2} + 2\eta_\theta L_{\theta,\lambda} \gamma^H + \frac{\ell_\theta L_\lambda^2 \varepsilon^2}{(1 - \gamma)^2} \right) + 2\eta_\mu M^2 \right],
\end{aligned} \tag{85}$$

where the second term $\frac{2\eta_\theta}{\varepsilon} \sum_{t=0}^{T-1} \|e_{L,t}\|$ is due to

$$\sum_{t=0}^{T-1} \sum_{i=0}^{t-1} (1 - \varepsilon)^i \cdot 2\eta_\theta \|e_{L,t-1-i}\| = 2\eta_\theta \sum_{t=0}^{T-1} \left(\sum_{i=0}^{T-1-t} (1 - \varepsilon)^i \right) \|e_{L,t}\| \leq \frac{2\eta_\theta}{\varepsilon} \sum_{t=0}^{T-1} \|e_{L,t}\|. \tag{86}$$

The proof of Lemma 5 is completed by dividing both sides of (85) by T . \square

C.6 Proof of Lemma 6

PROOF. By definition of the Lagrangian function $L(\lambda(\theta), \mu)$, we have that

$$\begin{aligned}
\frac{1}{T} \sum_{t=0}^{T-1} f(\lambda(\theta^*)) - f(\lambda(\theta_t)) &= \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - L(\lambda(\theta_t), \mu_t)] + \frac{1}{T} \sum_{t=0}^{T-1} \mu_t g(\lambda(\theta_t)) \\
&\stackrel{(i)}{=} \frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu^*) - L(\lambda(\theta_t), \mu_t)] + \frac{1}{T} \sum_{t=0}^{T-1} \mu_t g(\lambda(\theta_t)) \\
&\stackrel{(ii)}{\leq} \frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] + \frac{1}{T} \sum_{t=0}^{T-1} \mu_t g(\lambda(\theta_t)) \\
&\leq C + \frac{1}{T} \sum_{t=0}^{T-1} \mu_t g(\lambda(\theta_t)),
\end{aligned} \tag{87}$$

where the equality in step (i) holds due to complementary slackness that $\mu^* g(\lambda(\theta^*)) = 0$ (see Lemma 1), and step (ii) is due to the fact that $\mu^* = \operatorname{argmin}_{\mu \geq 0} L(\lambda(\theta^*), \mu)$.

Then, we upper-bound the second term in the last line of (87). By the update rule of μ_t in Algorithm 1 and the fact that $\nabla_{\mu} L(\lambda(\theta_t), \mu_t) = g(\lambda_t)$, we obtain that for any $\mu \in U$

$$\begin{aligned}
(\mu_{t+1} - \mu)^2 &= [\mathcal{P}_u(\mu_t - \eta_{\mu} g(\lambda_t)) - \mu]^2 \\
&\stackrel{(i)}{\leq} (\mu_t - \eta_{\mu} g(\lambda_t) - \mu)^2 \\
&= (\mu_t - \mu)^2 - 2\eta_{\mu}(\mu_t - \mu) \cdot g(\lambda_t) + (\eta_{\mu} g(\lambda_t))^2,
\end{aligned} \tag{88}$$

where step (i) follows from the non-expansive property of the projection operator. Then, using the inequality $(x + y)^2 \leq 2x^2 + 2y^2$ and the Lipschitz continuity of $g(\cdot)$, we have that

$$[g(\lambda_t)]^2 \leq 2[g(\lambda(\theta_t))]^2 + 2[g(\lambda(\theta_t)) - g(\lambda_t)]^2 \leq 2M^2 + 2M_{\lambda}^2 \|\lambda(\theta_t) - \lambda_t\|^2, \tag{89}$$

where we use the boundedness of $g(\cdot)$ on Λ . We substitute (89) into (88) to obtain that

$$(\mu_{t+1} - \mu)^2 \leq (\mu_t - \mu)^2 - 2\eta_{\mu}(\mu_t - \mu) \cdot g(\lambda_t) + 2\eta_{\mu}^2 M^2 + 2M_{\lambda}^2 \|\lambda(\theta_t) - \lambda_t\|^2 \eta_{\mu}^2. \tag{90}$$

Then, by setting $\mu = 0$ and rearranging the terms, we have that

$$\mu_t g(\lambda_t) \leq \frac{1}{2\eta_{\mu}} \left(\mu_t^2 - \mu_{t+1}^2 + 2\eta_{\mu}^2 M^2 + 2M_{\lambda}^2 \|\lambda(\theta_t) - \lambda_t\|^2 \eta_{\mu}^2 \right). \tag{91}$$

Finally, using (91) and the Lipschitz continuity of $g(\cdot)$, we can bound $\mu_t g(\lambda(\theta_t))$ as follows

$$\begin{aligned}
\mu_t g(\lambda(\theta_t)) &= \mu_t g(\lambda_t) + \mu_t (g(\lambda(\theta_t)) - g(\lambda_t)) \\
&\leq \frac{1}{2\eta_{\mu}} \left(\mu_t^2 - \mu_{t+1}^2 + 2\eta_{\mu}^2 M^2 + 2M_{\lambda}^2 \|\lambda(\theta_t) - \lambda_t\|^2 \eta_{\mu}^2 \right) + C_0 M_{\lambda} \|\lambda(\theta_t) - \lambda_t\| \\
&\stackrel{(i)}{\leq} \frac{1}{2\eta_{\mu}} \left(\mu_t^2 - \mu_{t+1}^2 \right) + \eta_{\mu} M^2 + M_{\lambda}^2 \eta_{\mu} (2 \|\lambda(\theta_t) - \lambda_H(\theta_t)\|^2 + 2 \|\lambda_H(\theta_t) - \lambda_t\|^2) \\
&\quad + C_0 M_{\lambda} (\|\lambda(\theta_t) - \lambda_H(\theta_t)\| + \|\lambda_H(\theta_t) - \lambda_t\|),
\end{aligned} \tag{92}$$

where (i) is again due to the inequality $(x + y)^2 \leq 2x^2 + 2y^2$. We sum both sides of (92) from $t = 0$ to $T - 1$ and adopt the notation $e_{\lambda,t} := \lambda_H(\theta_t) - \lambda_t$ to obtain that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mu_t g(\lambda(\theta_t)) &\leq \frac{\mu_0^2 - \mu_T^2}{2T\eta_\mu} + \eta_\mu M^2 + \frac{2\eta_\mu M_\lambda^2}{T} \sum_{t=0}^{T-1} \left(\|\lambda(\theta_t) - \lambda_H(\theta_t)\|^2 + \|e_{\lambda,t}\|^2 \right) \\ &\quad + \frac{C_0 M_\lambda}{T} \sum_{t=0}^{T-1} (\|\lambda(\theta_t) - \lambda_H(\theta_t)\| + \|e_{\lambda,t}\|) \\ &\stackrel{(i)}{\leq} \frac{\mu_0^2 - \mu_T^2}{2T\eta_\mu} + \eta_\mu M^2 + \frac{C_0 M_\lambda \gamma^H}{1-\gamma} + \frac{2\eta_\mu M_\lambda^2 \gamma^{2H}}{(1-\gamma)^2} + \frac{C_0 M_\lambda}{T} \sum_{t=0}^{T-1} \|e_{\lambda,t}\| + \frac{2\eta_\mu M_\lambda^2}{T} \sum_{t=0}^{T-1} \|e_{\lambda,t}\|^2, \end{aligned} \quad (93)$$

where we further apply the upper bound $\|\lambda(\theta_t) - \lambda_H(\theta_t)\| \leq \gamma^H / (1 - \gamma)$ in (i). It is clear that the optimality gap in (32a) can be obtained by substituting the upper bound in (93) into (87).

Now, we proceed to the constraint violation. If $\sum_{t=0}^{T-1} g(\lambda(\theta_t)) \geq 0$, then there is no constraint violation and the bound in (32b) is trivially satisfied. Therefore, from now on, we assume $[\sum_{t=0}^{T-1} -g(\lambda(\theta_t))]_+ > 0$, which implies that $\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) = [\sum_{t=0}^{T-1} -g(\lambda(\theta_t))]_+$. Define $\widehat{\mu} := \mu^* + 1 \geq 1$, as $\mu^* \geq 0$. By the boundedness of μ^* (see Lemma 1), we have that

$$\widehat{\mu} = \mu^* + 1 \leq \frac{f(\lambda(\theta^*)) - f(\lambda(\widetilde{\theta}))}{\xi} + 1 \leq \frac{M - f(\lambda(\widetilde{\theta}))}{\xi} + 1 = C_0,$$

which implies that $\widehat{\mu} \in U$. Thus, it follows that

$$\frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ = (\widehat{\mu} - \mu^*) \cdot \frac{1}{T} \sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \leq \max_{\mu \in U} \left\{ (\mu^* - \mu) \cdot \frac{1}{T} \sum_{t=0}^{T-1} g(\lambda(\theta_t)) \right\}. \quad (94)$$

To upper-bound the last line in (94), we note that

$$\begin{aligned} (\mu^* - \mu) \cdot g(\lambda(\theta_t)) &= (\mu^* - \mu_t) \cdot g(\lambda(\theta_t)) + (\mu_t - \mu) \cdot g(\lambda(\theta_t)) \\ &= (\mu^* - \mu_t) \cdot \nabla_\mu L(\lambda(\theta_t), \mu_t) + (\mu_t - \mu) \cdot g(\lambda(\theta_t)) \\ &\stackrel{(i)}{=} L(\lambda(\theta_t), \mu^*) - L(\lambda(\theta_t), \mu_t) + (\mu_t - \mu) \cdot g(\lambda(\theta_t)) \\ &\stackrel{(ii)}{\leq} L(\lambda(\theta^*), \mu^*) - L(\lambda(\theta_t), \mu_t) + (\mu_t - \mu) \cdot g(\lambda(\theta_t)) \\ &\stackrel{(iii)}{\leq} L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t) + (\mu_t - \mu) \cdot g(\lambda(\theta_t)), \end{aligned} \quad (95)$$

where we use the linearity of $L(\lambda(\theta), \mu)$ w.r.t. μ in step (i). Step (ii) follows from the fact that θ^* maximizes $L(\lambda(\cdot), \mu^*)$, and step (iii) is due to the fact that μ^* minimize $L(\lambda(\theta^*), \cdot)$ from (13). Then, by (90) and (92), we can control the last term in step (iii) of (95) as

$$\begin{aligned} &(\mu_t - \mu) \cdot g(\lambda(\theta_t)) \\ &= (\mu_t - \mu) \cdot g(\lambda_t) + (\mu_t - \mu) \cdot [g(\lambda(\theta_t)) - g(\lambda_t)] \\ &\leq \frac{1}{2\eta_\mu} [(\mu_t - \mu)^2 - (\mu_{t+1} - \mu)^2 + 2\eta_\mu^2 M^2 + 2M_\lambda^2 \|\lambda(\theta_t) - \lambda_t\|^2 \eta_\mu^2] + C_0 M_\lambda \|\lambda(\theta_t) - \lambda_t\| \\ &\leq \frac{1}{2\eta_\mu} [(\mu_t - \mu)^2 - (\mu_{t+1} - \mu)^2] + \eta_\mu M^2 + 2\eta_\mu M_\lambda^2 \left(\frac{\gamma^{2H}}{(1-\gamma)^2} + \|e_{\lambda,t}\|^2 \right) + C_0 M_\lambda \left(\frac{\gamma^H}{1-\gamma} + \|e_{\lambda,t}\| \right), \end{aligned} \quad (96)$$

where the last inequality is due to similar bounds on $\|\lambda(\theta_t) - \lambda_t\|$ and $\|\lambda(\theta_t) - \lambda_t\|^2$ as (92) and (93). Combining (95) and (96), we derive the constraint violation as follows

$$\begin{aligned} \frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ &\leq \frac{1}{T} \sum_{t=0}^{T-1} [L(\lambda(\theta^*), \mu_t) - L(\lambda(\theta_t), \mu_t)] + \eta_\mu M^2 + \frac{C_0 M_\lambda \gamma^H}{1-\gamma} + \frac{2\eta_\mu M_\lambda^2 \gamma^{2H}}{(1-\gamma)^2} \\ &+ \frac{C_0 M_\lambda}{T} \sum_{t=0}^{T-1} \|e_{\lambda,t}\| + \frac{2\eta_\mu M_\lambda^2}{T} \sum_{t=0}^{T-1} \|e_{\lambda,t}\|^2 + \frac{\max_{\mu \in U} \{(\mu_0 - \mu)^2 - (\mu_T - \mu)^2\}}{2T\eta_\mu}, \end{aligned} \quad (97)$$

which directly implies (32b). \square

C.7 Proof of Theorem 2 (Sample-based Setting)

PROOF. The proof primarily utilizes Lemma 5, Lemma 6, and the boundedness of approximation errors $\|e_{L,t}\|$ and $\|e_{\lambda,t}\|$ (see Appendix C.8). Let $\varepsilon = \frac{(1-\gamma)\eta_\theta}{\sqrt{2}L_\lambda} = \frac{1-\gamma}{\sqrt{2}L_\lambda} T^{-1/2}$. For $T \geq \frac{1}{2} \left(\frac{1-\gamma}{L_\lambda \varepsilon}\right)^2$, this choice of ε ensures that $\varepsilon \leq \bar{\varepsilon}$, and therefore it is consistent with Assumption 2. Besides, by viewing the step-sizes as variables, one can also verify that the above parameter choices obtain the optimal convergence rate for both the optimality gap and constraint violation. Under this set of parameter, we have that $\mathbb{E}[\|e_{\lambda,t}\|^2] \lesssim \frac{C_\lambda}{\sqrt{t+2}}$ from (118). Then, summing it over $t = 0, \dots, T-1$ yields that

$$\sum_{t=0}^{T-1} \mathbb{E}[\|e_{\lambda,t}\|^2] \lesssim C_\lambda \int_0^T \frac{1}{\sqrt{t+2}} dt \lesssim C_\lambda T^{1/2}. \quad (98)$$

Furthermore, since $\mathbb{E}[\|e_{\lambda,t}\|] \leq \sqrt{\mathbb{E}[\|e_{\lambda,t}\|^2]}$, it holds that $\mathbb{E}[\|e_{\lambda,t}\|] \lesssim \frac{\sqrt{C_\lambda}}{(t+2)^{1/4}}$. Hence, summing it over $t = 0, \dots, T-1$ yields that

$$\sum_{t=0}^{T-1} \mathbb{E}[\|e_{\lambda,t}\|] \lesssim \sqrt{C_\lambda} \int_0^T \frac{1}{(t+2)^{1/4}} dt \lesssim \sqrt{C_\lambda} T^{3/4}. \quad (99)$$

Then, from (155) in Lemma 14, we have that $\mathbb{E}[\|e_{L,t}\|] \lesssim \frac{C_e(1+C_0)}{(t+1)^{1/4}}$. Summing the term over $t = 0, \dots, T-1$ results in that

$$\sum_{t=0}^{T-1} \mathbb{E}[\|e_{L,t}\|] \lesssim C_e(1+C_0) \int_0^T \frac{1}{(t+1)^{1/4}} dt \lesssim C_e(1+C_0) T^{3/4}. \quad (100)$$

Therefore, by (31) in Lemma 5 and (32a) in Lemma 6, the optimality gap in the sample-based setting can be upper-bounded as follows

$$\begin{aligned}
& \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} f(\lambda(\theta^*)) - f(\lambda(\theta_t)) \right] \\
& \leq \frac{L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)}{\varepsilon T} + \frac{2\eta_\theta}{\varepsilon T} \sum_{t=0}^{T-1} \mathbb{E} [\|e_{L,t}\|] + \frac{1}{\varepsilon} \left[(1+C_0) \left(\frac{\ell_\theta \eta_\theta^2}{2} + 2L_{\theta,\lambda} \eta_\theta \gamma^H \right) + 2\eta_\mu M^2 \right] \\
& \quad + \frac{(1+C_0)\ell_\theta L_\lambda^2 \varepsilon}{(1-\gamma)^2} + \frac{\mu_0^2 - \mu_T^2}{2T\eta_\mu} + \eta_\mu M^2 + \frac{C_0 M_\lambda \gamma^H}{1-\gamma} + \frac{2\eta_\mu M_\lambda^2 \gamma^{2H}}{(1-\gamma)^2} + \frac{C_0 M_\lambda}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|e_{\lambda,t}\|] + \frac{2\eta_\mu M_\lambda^2}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|e_{\lambda,t}\|^2] \\
& \stackrel{(i)}{\lesssim} \frac{L_\lambda M}{(1-\gamma)T^{1/2}} + \frac{L_\lambda C_e (1+C_0)}{(1-\gamma)T^{1/4}} + (1+C_0) \left(\frac{\ell_\theta L_\lambda}{(1-\gamma)T^{1/2}} + \frac{L_{\theta,\lambda} L_\lambda \gamma^H}{1-\gamma} \right) + \frac{L_\lambda M^2}{(1-\gamma)T^{1/4}} + \frac{(1+C_0)\ell_\theta L_\lambda}{(1-\gamma)T^{1/2}} + \frac{M^2}{T^{3/4}} \\
& \quad + \frac{C_0 M_\lambda \gamma^H}{1-\gamma} + \frac{M_\lambda^2 \gamma^{2H}}{(1-\gamma)^2 T^{3/4}} + \frac{C_0 M_\lambda \sqrt{C_\lambda}}{T^{1/4}} + \frac{M_\lambda^2 C_\lambda}{T^{5/4}} \\
& \stackrel{(ii)}{\lesssim} \left[\frac{C_0 L_\lambda C_e + L_\lambda M^2}{1-\gamma} + C_0 M_\lambda \sqrt{C_\lambda} \right] \cdot \frac{1}{T^{1/4}} \\
& = \mathcal{O} \left(\frac{T^{-1/4}}{(1-\gamma)^4} \right).
\end{aligned}$$

For the first term in step (i), since $\mu_0 = 0$, $\varepsilon = \frac{(1-\gamma)\eta_\theta}{\sqrt{2}L_\lambda} = \frac{1-\gamma}{\sqrt{2}L_\lambda} T^{-1/2}$, and $|f(\lambda)| \leq M$ for all $\lambda \in \Lambda$ from Lemma 2, we have that

$$\frac{L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)}{\varepsilon T} = \frac{f(\lambda(\theta^*)) - f(\lambda(\theta_0))}{\varepsilon T} \lesssim \frac{L_\lambda M}{(1-\gamma)T^{1/2}}. \quad (101)$$

The second term in step (i) is obtained from (100) under the choice of $\eta_\theta = T^{-1/2}$, i.e.,

$$\frac{2\eta_\theta}{\varepsilon T} \sum_{t=0}^{T-1} \|e_{L,t}\| \lesssim \frac{\sqrt{2}L_\lambda}{(1-\gamma)T} \cdot C_e (1+C_0) T^{3/4} \lesssim \frac{L_\lambda C_e (1+C_0)}{(1-\gamma)T^{1/4}}. \quad (102)$$

The remaining terms in step (i) can be similarly derived by substituting the specified parameters into the equation and using the upper bounds derived in (98) and (99). By only collecting the terms associated with $T^{-1/4}$ and ignoring the problem-independent constants, we arrive at the final bound in (ii). Finally, we remark that since $C_e = \mathcal{O}\left(\frac{1}{(1-\gamma)^3}\right)$, the bound in (ii) has the order $\mathcal{O}\left(\frac{T^{-1/4}}{(1-\gamma)^4}\right)$.

Next, the constraint violation of Algorithm 1 can be derived from (31) in Lemma 5 and (32b) in Lemma 6. Since the derivation closely resembles that of the optimality gap, we will not reiterate some intermediate steps and

only write the main inequalities:

$$\begin{aligned}
\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ &\leq \frac{L(\lambda(\theta^*), \mu_0) - L(\lambda(\theta_0), \mu_0)}{\varepsilon T} + \frac{2\eta_\theta}{\varepsilon T} \sum_{t=0}^{T-1} \mathbb{E} [\|e_{L,t}\|] + \frac{(1+C_0)\ell_\theta L_\lambda^2 \varepsilon}{(1-\gamma)^2} \\
&+ \frac{1}{\varepsilon} \left[(1+C_0) \left(\frac{\ell_\theta \eta_\theta^2}{2} + 2L_{\theta,\lambda} \eta_\theta \gamma^H \right) + 2\eta_\mu M^2 \right] + \frac{\max_{\mu \in U} \{(\mu_0 - \mu)^2 - (\mu_T - \mu)^2\}}{2T\eta_\mu} \\
&+ \eta_\mu M^2 + \frac{C_0 M_\lambda \gamma^H}{1-\gamma} + \frac{2\eta_\mu M_\lambda^2 \gamma^{2H}}{(1-\gamma)^2} + \frac{C_0 M_\lambda}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|e_{\lambda,t}\|] + \frac{2\eta_\mu M_\lambda^2}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|e_{\lambda,t}\|^2] \\
&\stackrel{(i)}{\lesssim} \frac{L_\lambda M}{(1-\gamma)T^{1/2}} + \frac{L_\lambda C_e (1+C_0)}{(1-\gamma)T^{1/4}} + (1+C_0) \left(\frac{\ell_\theta L_\lambda}{(1-\gamma)T^{1/2}} + \frac{L_{\theta,\lambda} L_\lambda \gamma^H}{1-\gamma} \right) + \frac{L_\lambda M^2}{(1-\gamma)T^{1/4}} \\
&+ \frac{(1+C_0)\ell_\theta L_\lambda}{(1-\gamma)T^{1/2}} + \frac{C_0^2}{T^{1/4}} + \frac{M^2}{T^{3/4}} + \frac{C_0 M_\lambda \gamma^H}{1-\gamma} + \frac{M_\lambda^2 \gamma^{2H}}{(1-\gamma)^2 T^{3/4}} + \frac{C_0 M_\lambda \sqrt{C_\lambda}}{T^{1/4}} + \frac{M_\lambda^2 C_\lambda}{T^{5/4}} \\
&\lesssim \left[\frac{C_0 L_\lambda C_e}{1-\gamma} + C_0^2 + C_0 M_\lambda \sqrt{C_\lambda} \right] \cdot \frac{1}{T^{1/4}} \\
&= O \left(\frac{T^{-1/4}}{(1-\gamma)^4} \right).
\end{aligned}$$

where in step (i), we again use the inequality $\max_{\mu \in U} \{(\mu_0 - \mu)^2 - (\mu_T - \mu)^2\} = \max_{\mu \in U} \{(2\mu - \mu_T)\mu_T\} \leq C_0^2$. This completes the proof of Theorem 2. \square

C.8 Control the Sizes of $\|e_{L,t}\|$ and $\|e_{\lambda,t}\|$

In this section, we provide upper bounds for the approximation error terms $e_{L,t}$ and $e_{\lambda,t}$, which are modified from that of [7]. Throughout the section, we assume that Assumption 3 holds true and will adhere to the choices of $\alpha_t = (t+1)^{-1/2}$ and $\eta_\theta = T^{-1/2}$. The following technical lemmas will be frequently used in this section.

LEMMA 11. *Suppose that the sequence $\{x_t\}_{t \geq 0}$ satisfies $x_t \leq (1 - \alpha_t)^2 x_{t-1} + \beta_t$, where $\alpha_t = (t+1)^{-1/2}$ and $\beta_t \leq C(t+1)^{-1}$. Then, it holds that*

$$x_t \leq \frac{x_0}{t+1} + \frac{12C}{\sqrt{t+2}} = O \left(\frac{1}{\sqrt{t+2}} \right). \quad (103)$$

PROOF. We first unroll the recursion for x_t to obtain that

$$x_t \leq x_0 \prod_{t'=1}^t (1 - \alpha_{t'})^2 + \sum_{t'=1}^t \beta_{t'} \cdot \prod_{t''=t'+1}^t (1 - \alpha_{t''})^2 \stackrel{(i)}{\leq} x_0 \prod_{t'=1}^t (1 - \alpha_{t'})^2 + \sum_{t'=1}^t \beta_{t'} \cdot \exp \left(-2 \sum_{t''=t'+1}^t \alpha_{t''} \right). \quad (104)$$

where step (i) uses the basic inequality $1 - x \leq \exp(-x)$.

We first focus on the second term on the right-hand side of (104) and begin with bounding $\sum_{t''=t'+1}^t \alpha_{t''}$ by integration. Since $\alpha_t = (t+1)^{-1/2}$ is monotone decreasing in t , it holds that

$$\sum_{t''=t'+1}^t \alpha_{t''} \geq \int_1^{t+1} (t'+1)^{-\frac{1}{2}} dt' = 2(t+2)^{\frac{1}{2}} - 2\sqrt{3}, \quad (105a)$$

$$\sum_{t''=t'+1}^t \alpha_{t''} \leq \int_0^t (t'+1)^{-\frac{1}{2}} dt' = 2(t+1)^{\frac{1}{2}} - 2. \quad (105b)$$

Then, we have that

$$\begin{aligned}
\sum_{t'=1}^t \beta_{t'} \cdot \exp\left(-2 \sum_{t''=t'+1}^t \alpha_{t''}\right) &= \exp\left(-2 \sum_{t'=1}^t \alpha_{t'}\right) \cdot \sum_{t'=1}^t \beta_{t'} \cdot \exp\left(2 \sum_{t''=1}^{t'} \alpha_{t''}\right) \\
&\stackrel{(i)}{\leq} \exp\left(-2 \sum_{t'=1}^t \alpha_{t'}\right) \cdot \sum_{t'=1}^t \beta_{t'} \cdot \exp\left(4(t'+1)^{\frac{1}{2}} - 4\right) \\
&\leq C \exp\left(-2 \sum_{t'=1}^t \alpha_{t'} - 4\right) \cdot \underbrace{\sum_{t'=1}^t (t'+1)^{-1} \cdot \exp\left(4(t'+1)^{\frac{1}{2}}\right)}_{\mathcal{T}},
\end{aligned} \tag{106}$$

where step (i) uses the upper bound in (105b). To further bound \mathcal{T} , we consider the function $f(t') = (t'+1)^{-1} \cdot \exp(4(t'+1)^{1/2})$, whose derivative can be computed as follows

$$f'(t') = (t'+1)^{-\frac{5}{2}} \cdot \exp\left(4(t'+1)^{\frac{1}{2}}\right) \cdot \left(2t' - (t'+1)^{\frac{1}{2}} + 2\right). \tag{107}$$

Hence, $f(y)$ is monotone increasing when $y \geq 0$. Therefore, the summation term \mathcal{T} can be again upper-bounded by the integration:

$$\begin{aligned}
\mathcal{T} &= \sum_{t'=1}^t (t'+1)^{-1} \cdot \exp\left(4(t'+2)^{\frac{1}{2}}\right) \leq \int_1^{t+1} (t'+1)^{-1} \cdot \exp\left(4(t'+1)^{\frac{1}{2}}\right) dt' \\
&\stackrel{(i)}{\leq} \int_2^{t+2} y^{-1} \exp\left(4y^{\frac{1}{2}}\right) dy =: I,
\end{aligned} \tag{108}$$

where in step (i), we perform a change of variables and let $y := t' + 1$, and then we define the entire term as I . Next, our goal is to provide an upper bound for I .

$$\begin{aligned}
I &= \int_2^{t+2} y^{-1} \cdot \frac{1}{2} y^{\frac{1}{2}} \cdot d\left(\exp(4y^{\frac{1}{2}})\right) \\
&\stackrel{(i)}{=} \frac{1}{2} y^{-\frac{1}{2}} \exp\left(4y^{\frac{1}{2}}\right) \Big|_2^{t+2} + \frac{1}{4} \int_2^{t+2} y^{-\frac{3}{2}} \exp\left(4y^{\frac{1}{2}}\right) dy \\
&\stackrel{(ii)}{\leq} \frac{1}{2} (t+2)^{-\frac{1}{2}} \cdot \exp\left(4(t+2)^{\frac{1}{2}}\right) + \underbrace{\frac{1}{4\sqrt{2}} \int_2^{t+2} y^{-1} \exp\left(4y^{\frac{1}{2}}\right) dy}_{=I}.
\end{aligned} \tag{109}$$

where step (i) involves integration by parts, and step (ii) is derived as follows

$$\int_2^{t+2} y^{-\frac{3}{2}} \exp\left(4y^{\frac{1}{2}}\right) dy = \int_2^{t+2} y^{-\frac{1}{2}} \cdot y^{-1} \exp\left(4y^{\frac{1}{2}}\right) dy \leq 2^{-\frac{1}{2}} \int_2^{t+2} y^{-1} \exp\left(4y^{\frac{1}{2}}\right) dy. \tag{110}$$

By rearranging the terms, (109) can be further simplified as

$$\left(1 - \frac{1}{4\sqrt{2}}\right) I \leq \frac{1}{2} (t+2)^{-\frac{1}{2}} \cdot \exp\left(4(t+2)^{\frac{1}{2}}\right), \tag{111}$$

and therefore

$$I \leq \frac{1}{1 - \frac{1}{4\sqrt{2}}} \cdot \frac{1}{2} (t+2)^{-\frac{1}{2}} \cdot \exp\left(4(t+2)^{\frac{1}{2}}\right) = \frac{2\sqrt{2}}{4\sqrt{2} - 1} (t+2)^{-\frac{1}{2}} \cdot \exp\left(4(t+2)^{\frac{1}{2}}\right). \tag{112}$$

Combining (106) and (112), we obtain that

$$\begin{aligned} \sum_{t'=1}^t \beta_{t'} \cdot \exp\left(-2 \sum_{t''=t'+1}^t \alpha_{t''}\right) &\leq C \exp\left(-2 \sum_{t'=1}^t \alpha_{t'} - 4\right) \cdot \sum_{t'=1}^t (t'+1)^{-1} \cdot \exp\left(4(t'+1)^{\frac{1}{2}}\right) \\ &\stackrel{(i)}{\leq} C \exp\left(-4(t+2)^{\frac{1}{2}} + 4\sqrt{3} - 4\right) \cdot \frac{2\sqrt{2}}{4\sqrt{2}-1} (t+2)^{-\frac{1}{2}} \exp\left(4(t+2)^{\frac{1}{2}}\right) \\ &\leq \frac{12C}{\sqrt{t+2}}, \end{aligned} \quad (113)$$

where in step (i), we use (105a) to upper-bound the term $\exp(-2 \sum_{t'=1}^t \alpha_{t'})$.

Finally, we consider the first term $x_0 \prod_{t'=1}^t (1 - \alpha_{t'})^2$ in (104). It holds that

$$1 - \alpha_{t'} = 1 - (t'+1)^{-\frac{1}{2}} \leq 1 - \frac{1}{2}(t')^{-1} \stackrel{(i)}{\leq} (1 + (t')^{-1})^{-\frac{1}{2}} = \left(\frac{t'+1}{t'}\right)^{-\frac{1}{2}} = \frac{\alpha_{t'}}{\alpha_{t'-1}}, \quad (114)$$

where step (i) is due to Bernoulli's inequality that $(1+z)^r \geq 1+rz$ for all $z \geq -1$ and $r \leq 0$. Multiplying the above term over $t' = 1, \dots, t$, we have that

$$\prod_{t'=1}^t (1 - \alpha_{t'})^2 \leq \left[\prod_{t'=1}^t \frac{\alpha_{t'}}{\alpha_{t'-1}} \right]^2 = \alpha_t^2 = \frac{1}{t+1}. \quad (115)$$

Combining (104), (113) and (115), we arrive at

$$x_t \leq x_0 \prod_{t'=1}^t (1 - \alpha_{t'})^2 + \sum_{t'=1}^t \beta_{t'} \cdot \exp\left(-2 \sum_{t''=t'+1}^t \alpha_{t''}\right) \leq \frac{x_0}{t+1} + \frac{12C}{\sqrt{t+2}}. \quad (116)$$

This completes the derivation of Lemma 11. \square

LEMMA 12. Under Algorithm 1, let $\{e_{\lambda,t}\}_{t=0}^{T-1}$ be the sequence defined as $e_{\lambda,t} := \lambda_H(\theta_t) - \lambda_t$, for all $t \geq 0$. Then, it holds that

$$\mathbb{E} \left[\|e_{\lambda,t}\|^2 \right] \leq (1 - \alpha_t)^2 \mathbb{E} \left[\|e_{\lambda,t-1}\|^2 \right] + \frac{2C_w \eta_\theta^2}{(1-\gamma)^2} + \frac{2\alpha_t^2}{(1-\gamma)^2}, \quad (117)$$

where $C_w := 2H \left[(8H+2)M_{\pi,1}^2 + 2M_{\pi,2} \right]$. Furthermore, by letting $\alpha_t = (t+1)^{-1/2}$, $\eta_\theta = T^{-1/2}$, and $C_\lambda := \frac{C_w+1}{(1-\gamma)^2}$, it holds that

$$\mathbb{E} \left[\|e_{\lambda,t}\|^2 \right] \leq \frac{25C_\lambda}{\sqrt{t+2}} \lesssim \frac{C_\lambda}{\sqrt{t+2}}. \quad (118)$$

PROOF. To show the recursion in (117), we build a connection between $e_{\lambda,t}$ and $e_{\lambda,t-1}$:

$$\begin{aligned} e_{\lambda,t} &= \lambda_H(\theta_t) - \lambda_t \\ &\stackrel{(i)}{=} \lambda_H(\theta_t) - \widehat{\lambda}(\tau_t) - (1 - \alpha_t) \left[\lambda_{t-1} - w(\tau_t | \theta_{t-1}, \theta_t) \widehat{\lambda}(\tau_t) \right] \\ &\stackrel{(ii)}{=} (1 - \alpha_t) (\lambda_H(\theta_{t-1}) - \lambda_{t-1}) + \alpha_t y_t - (1 - \alpha_t) z_t \\ &= (1 - \alpha_t) e_{\lambda,t-1} + \alpha_t y_t - (1 - \alpha_t) z_t, \end{aligned} \quad (119)$$

where step (i) follows from the update rule of λ_t in (20) of Algorithm 1. In step (ii), we let

$$y_t := \lambda_H(\theta_t) - \widehat{\lambda}(\tau_t), \quad z_t := \lambda_H(\theta_{t-1}) - \lambda_H(\theta_t) + \widehat{\lambda}(\tau_t) (1 - w(\tau_t | \theta_{t-1}, \theta_t)). \quad (120)$$

Hence, it holds from (119) that

$$\begin{aligned} \mathbb{E}[\|e_{\lambda,t}\|^2] &= (1 - \alpha_t)^2 \mathbb{E}[\|e_{\lambda,t-1}\|^2] + \mathbb{E}[\|\alpha_t y_t - (1 - \alpha_t) z_t\|^2] \\ &\quad + \mathbb{E}[\langle (1 - \alpha_t) e_{\lambda,t-1}, \alpha_t y_t - (1 - \alpha_t) z_t \rangle]. \end{aligned} \quad (121)$$

We first show that the third expectation term on the right-hand side of (121) is equal to zero. Consider the filtration $\{\mathcal{F}_t\}_{t \geq 0}$, where the σ -algebra \mathcal{F}_t is generated by all trajectories of length H sampled at $0, 1, \dots, t$ -th periods. Therefore, for each t , the error term $e_{\lambda,t}$ is \mathcal{F}_t -measurable. Since τ_t is sampled from policy π_{θ_t} and $\mathbb{E}[w(\tau_t | \theta_{t-1}, \theta_t) \widehat{\lambda}(\tau_t)] = \lambda_H(\theta_{t-1})$ by Lemma 8, we observe that $\mathbb{E}[y_t] = \mathbb{E}[z_t] = \mathbb{E}[y_t | \mathcal{F}_{t-1}] = \mathbb{E}[z_t | \mathcal{F}_{t-1}] = 0$. By the tower rule of conditional expectation, it holds that

$$\begin{aligned} \mathbb{E}[\langle (1 - \alpha_t) e_{\lambda,t-1}, \alpha_t y_t - (1 - \alpha_t) z_t \rangle] &= \mathbb{E}[\mathbb{E}[\langle (1 - \alpha_t) e_{\lambda,t-1}, \alpha_t y_t - (1 - \alpha_t) z_t \rangle | \mathcal{F}_{t-1}]] \\ &= \mathbb{E}[\langle (1 - \alpha_t) e_{\lambda,t-1}, \mathbb{E}[\alpha_t y_t - (1 - \alpha_t) z_t | \mathcal{F}_{t-1}] \rangle] \\ &= 0. \end{aligned} \quad (122)$$

Then, by (122), we can further simplify (121) as

$$\mathbb{E}[\|e_{\lambda,t}\|^2] \leq (1 - \alpha_t)^2 \mathbb{E}[\|e_{\lambda,t-1}\|^2] + 2\alpha_t^2 \mathbb{E}[\|y_t\|^2] + 2(1 - \alpha_t)^2 \mathbb{E}[\|z_t\|^2], \quad (123)$$

where we apply the inequality $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$.

Now, we proceed to bound the last two terms in (123):

$$\begin{aligned} \mathbb{E}[\|y_t\|^2] &= \mathbb{E}[\|\lambda_H(\theta_t) - \widehat{\lambda}(\tau_t)\|^2] \\ &= \mathbb{E}[\|\widehat{\lambda}(\tau_t)\|^2 - 2\langle \widehat{\lambda}(\tau_t), \lambda_H(\theta_t) \rangle + \|\lambda_H(\theta_t)\|^2] \\ &= \mathbb{E}[\|\widehat{\lambda}(\tau_t)\|^2] - 2\langle \mathbb{E}[\widehat{\lambda}(\tau_t)], \lambda_H(\theta_t) \rangle + \|\lambda_H(\theta_t)\|^2 \\ &\leq \mathbb{E}[\|\widehat{\lambda}(\tau_t)\|^2] \\ &\leq \frac{1}{(1 - \gamma)^2}, \end{aligned} \quad (124)$$

where the last inequality stems from the fact that

$$\|\widehat{\lambda}(\tau_t)\| \leq \sum_{t=0}^{H-1} \gamma^t \|\delta_{s_t, a_t}\| = \sum_{t=0}^{H-1} \gamma^t \leq \frac{1}{1 - \gamma}. \quad (125)$$

Next, we bound the last term in (123). By the definition z_t from (120) and the fact that $\mathbb{E}[\widehat{\lambda}(\tau_t) (1 - w(\tau_t|\theta_{t-1}, \theta_t))] = \lambda_H(\theta_t) - \lambda_H(\theta_{t-1})$, we have that

$$\begin{aligned}
\mathbb{E}[\|z_t\|^2] &\stackrel{(i)}{\leq} \mathbb{E}\left[\|\widehat{\lambda}(\tau_t) (1 - w(\tau_t|\theta_{t-1}, \theta_t))\|^2\right] \\
&= \mathbb{E}\left[(1 - w(\tau_t|\theta_{t-1}, \theta_t))^2 \|\widehat{\lambda}(\tau_t)\|^2\right] \\
&\stackrel{(ii)}{\leq} \frac{1}{(1-\gamma)^2} \mathbb{E}\left[(1 - w(\tau_t|\theta_{t-1}, \theta_t))^2\right] \\
&= \frac{1}{(1-\gamma)^2} \left[\text{Var}(w(\tau_t|\theta_{t-1}, \theta_t)) + (\mathbb{E}[w(\tau_t|\theta_{t-1}, \theta_t)] - 1)^2\right] \\
&= \frac{1}{(1-\gamma)^2} \text{Var}(w(\tau_t|\theta_{t-1}, \theta_t)) \\
&\leq \frac{C_w \eta_\theta^2}{(1-\gamma)^2},
\end{aligned} \tag{126}$$

where step (i) is due to the basic fact that $\mathbb{E}[\|X - \mathbb{E}[X]\|^2] \leq \mathbb{E}[\|X\|^2]$ for any random variable X . Then, inequality (i) stems from (125), and the last two steps follow from Lemma 9. Incorporating (123), (124) and (126) completes the proof of (117).

When $\alpha_t = (t+1)^{-1/2}$ and $\eta_\theta = T^{-1/2}$, we can apply Lemma 11 to (117) with $\beta_t = \frac{2C_w \eta_\theta^2}{(1-\gamma)^2} + \frac{2\alpha_t^2}{(1-\gamma)^2} \leq \frac{2C_w+2}{(1-\gamma)^2} \cdot \frac{1}{t+1}$ to derive that

$$\begin{aligned}
\mathbb{E}[\|e_{\lambda,t}\|^2] &\leq \frac{\mathbb{E}[\|e_{\lambda,0}\|^2]}{t+1} + 12 \cdot \frac{2C_w+2}{(1-\gamma)^2} \cdot \frac{1}{\sqrt{t+2}} \\
&\stackrel{(i)}{\leq} \frac{2}{(1-\gamma)^2(t+1)} + \frac{24(C_w+1)}{(1-\gamma)^2} \cdot \frac{1}{\sqrt{t+2}} \\
&\leq \frac{25(C_w+1)}{(1-\gamma)^2} \cdot \frac{1}{\sqrt{t+2}} \\
&\lesssim \frac{C_w+1}{(1-\gamma)^2} \cdot \frac{1}{\sqrt{t+2}} \\
&= \frac{C_\lambda}{\sqrt{t+2}},
\end{aligned} \tag{127}$$

where inequality (i) uses the fact that $\|e_{\lambda,0}\|^2 \leq 2\|\lambda_0\|^2 \leq \frac{2}{(1-\gamma)^2}$ with probability 1. This completes the proof of Lemma 12. \square

LEMMA 13. Under Algorithm 1, let $\{e_{f,t}\}_{t=0}^{T-1}$ be the sequence defined as $e_{f,t} := d_{f,t} - \nabla_\theta f(\lambda_H(\theta_t))$, for all $t \geq 0$. Then, it holds that

$$\mathbb{E}[\|e_{f,t}\|] \lesssim \frac{C_e}{(t+1)^{1/4}}, \tag{128}$$

where C_e is defined in (34b). The same upper bound also holds for $\{e_{g,t}\}_{t=0}^{T-1}$, where $e_{g,t} := d_{g,t} - \nabla_\theta g(\lambda_H(\theta_t))$.

PROOF. From the definition of $e_{f,t}$, we can decompose it as

$$e_{f,t} = d_{f,t} - \nabla_\theta f(\lambda_H(\theta_t)) = d_{f,t} - \langle \nabla_\theta \lambda_H(\theta_t), r_{f,t-1} \rangle + \langle \nabla_\theta \lambda_H(\theta_t), r_{f,t-1} - \nabla_\lambda f(\lambda_H(\theta_t)) \rangle, \tag{129}$$

where $r_{f,t} = \nabla_{\lambda} f(\lambda_t)$, defined in Algorithm 1, is the shadow reward estimator at period t . We remark that the quantities $d_{f,t}$ and $r_{f,t-1}$ in (129) are random, whereas other terms are deterministic. By letting $\widehat{e}_{f,t} := d_{f,t} - \langle \nabla_{\theta} \lambda_H(\theta_t), r_{f,t-1} \rangle$, we have from (129) that

$$\mathbb{E} [\|e_{f,t}\|] \leq \mathbb{E} [\|\widehat{e}_{f,t}\|] + \mathbb{E} [\|\langle \nabla_{\theta} \lambda_H(\theta_t), r_{f,t-1} - \nabla_{\lambda} f(\lambda_H(\theta_t)) \rangle\|]. \quad (130)$$

For the second term on the right-hand side of (130), it holds that

$$\begin{aligned} & \langle \nabla_{\theta} \lambda_H(\theta_t), r_{f,t-1} - \nabla_{\lambda} f(\lambda_H(\theta_t)) \rangle \\ & \stackrel{(i)}{\leq} \frac{2M_{\pi,1}}{(1-\gamma)^2} \|r_{f,t-1} - \nabla_{\lambda} f(\lambda_H(\theta_t))\|_{\infty} \\ & \leq \frac{2M_{\pi,1}}{(1-\gamma)^2} (\|r_{f,t-1} - \nabla_{\lambda} f(\lambda_H(\theta_{t-1}))\|_{\infty} + \|\nabla_{\lambda} f(\lambda_H(\theta_{t-1})) - \nabla_{\lambda} f(\lambda_H(\theta_t))\|_{\infty}) \\ & \stackrel{(ii)}{\leq} \frac{2M_{\pi,1}}{(1-\gamma)^2} (\ell_{\lambda} \|\lambda_{t-1} - \lambda_H(\theta_{t-1})\| + \ell_{\lambda} \|\lambda_H(\theta_{t-1}) - \lambda_H(\theta_t)\|) \\ & \stackrel{(iii)}{\leq} \frac{2M_{\pi,1}}{(1-\gamma)^2} (\ell_{\lambda} \|\lambda_{t-1} - \lambda_H(\theta_{t-1})\| + \frac{2\ell_{\lambda} M_{\pi,1}}{(1-\gamma)^2} \|\theta_{t-1} - \theta_t\|) \\ & = \frac{2M_{\pi,1}\ell_{\lambda}}{(1-\gamma)^2} \|e_{\lambda,t-1}\| + \frac{4M_{\pi,1}^2\ell_{\lambda}}{(1-\gamma)^4} \eta_{\theta}. \end{aligned} \quad (131)$$

In (131), step (i) is derived from the fact that $\langle \nabla_{\theta} \lambda_H(\theta_t), r \rangle = \nabla_{\theta} \langle \lambda_H(\theta_t), r \rangle$ is the policy gradient of the truncated value function with the reward $r = r_{f,t-1} - \nabla_{\lambda} f(\lambda_H(\theta_t))$. Then, from Lemma 15, it holds that

$$\langle \nabla_{\theta} \lambda_H(\theta), r \rangle = \mathbb{E} \left[\sum_{k=0}^{H-1} \gamma^k \nabla_{\theta} \log \pi_{\theta}(a_k | s_k) Q^{\pi_{\theta}}(r; s_k, a_k) \middle| \pi_{\theta}, s_0 \sim \rho \right] \leq \frac{2M_{\pi,1}}{(1-\gamma)^2} \|r\|_{\infty}, \quad (132)$$

where the last inequality uses the bound on the score function $\nabla_{\theta} \log \pi_{\theta}(a|s)$ from Lemma 19 and the fact that $|Q^{\pi_{\theta}}(r; s, a)| \leq \frac{\|r\|_{\infty}}{1-\gamma}$. Then, step (ii) in (131) is by the ℓ_{λ} -smoothness of function $f(\cdot)$ from Assumption 3 and the definition $r_{f,t-1} = \nabla_{\lambda} f(\lambda_{t-1})$. Step (iii) results from the Lipschitz continuity of $\lambda_H(\theta)$ with respect to θ stated in Lemma 20. Finally, the last step in (131) is by the definition $e_{\lambda,t-1} = \lambda_H(\theta_{t-1}) - \lambda_{t-1}$ and the update rule for θ in Algorithm 1, i.e., $\|\theta_t - \theta_{t-1}\| = \|\eta_{\theta,t-1} d_{L,t-1}\| = \eta_{\theta}$. Incorporating (130) and (131), we arrive at

$$\mathbb{E} [\|e_{f,t}\|] \leq \mathbb{E} [\|\widehat{e}_{f,t}\|] + \frac{2M_{\pi,1}\ell_{\lambda}}{(1-\gamma)^2} \mathbb{E} [\|e_{\lambda,t-1}\|] + \frac{4M_{\pi,1}^2\ell_{\lambda}\eta_{\theta}}{(1-\gamma)^4}. \quad (133)$$

Next, we focus on upper-bounding the term $\mathbb{E} [\|\widehat{e}_{f,t}\|]$ by deriving a recursion for the sequence $\{\widehat{e}_{f,t}\}_{t=0}^{T-1}$. Similar to the derivation in (119), we establish the relation between $\widehat{e}_{f,t}$ and $\widehat{e}_{f,t-1}$ as follows

$$\begin{aligned} \widehat{e}_{f,t} &= d_{f,t} - \langle \nabla_{\theta} \lambda_H(\theta_t), r_{f,t-1} \rangle \\ & \stackrel{(i)}{=} \widehat{d}(\tau_t, \theta_t, r_{f,t-1}) + (1-\alpha_t) [d_{f,t-1} - w(\tau_t | \theta_{t-1}, \theta_t) \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})] - \langle \nabla_{\theta} \lambda_H(\theta_t), r_{f,t-1} \rangle \\ & \stackrel{(ii)}{=} (1-\alpha_t) (d_{f,t-1} - \langle \nabla_{\theta} \lambda_H(\theta_{t-1}), r_{f,t-2} \rangle) + \alpha_t \widehat{y}_t + (1-\alpha_t) \widehat{z}_t \\ & = (1-\alpha_t) \widehat{e}_{f,t-1} + \alpha_t \widehat{y}_t + (1-\alpha_t) \widehat{z}_t, \end{aligned} \quad (134)$$

where step (i) is by the update rule of $d_{f,t}$ in Algorithm 1, and in step (ii), we define that

$$\widehat{y}_t := \widehat{d}(\tau_t, \theta_t, r_{f,t-1}) - \langle \nabla_{\theta} \lambda_H(\theta_t), r_{f,t-1} \rangle. \quad (135a)$$

$$\widehat{z}_t := \langle \nabla_{\theta} \lambda_H(\theta_{t-1}), r_{f,t-2} \rangle - \langle \nabla_{\theta} \lambda_H(\theta_t), r_{f,t-1} \rangle + \widehat{d}(\tau_t, \theta_t, r_{f,t-1}) - w(\tau_t | \theta_{t-1}, \theta_t) \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2}). \quad (135b)$$

We notice that τ_t is collected under π_{θ_t} , yet both $r_{f,t-1}$ and $r_{f,t-2}$ are computed before period t . Thus, the similar reasoning from (121) to (123) still applies. We have $\mathbb{E}[\widehat{y}_t] = \mathbb{E}[\widehat{z}_t] = \mathbb{E}[\widehat{y}_t | \mathcal{F}_{t-1}] = \mathbb{E}[\widehat{z}_t | \mathcal{F}_{t-1}] = 0$, and thereby (134) further implies that

$$\begin{aligned} \mathbb{E}[\|\widehat{e}_{f,t}\|^2] &\leq (1 - \alpha_t)^2 \mathbb{E}[\|\widehat{e}_{f,t-1}\|^2] + \mathbb{E}[\|\alpha_t \widehat{y}_t + (1 - \alpha_t) \widehat{z}_t\|^2] \\ &\leq (1 - \alpha_t)^2 \mathbb{E}[\|\widehat{e}_{f,t-1}\|^2] + 2\alpha_t^2 \mathbb{E}[\|\widehat{y}_t\|^2] + 2(1 - \alpha_t)^2 \mathbb{E}[\|\widehat{z}_t\|^2]. \end{aligned} \quad (136)$$

We first focus on the term $\mathbb{E}[\|\widehat{y}_t\|^2]$ in (136). By the definition in (135a), we observe that

$$\|\widehat{y}_t\| \leq \|\widehat{d}(\tau_t, \theta_t, r_{f,t-1})\| + \|\langle \nabla_{\theta} \lambda_H(\theta_t), r_{f,t-1} \rangle\| \stackrel{(i)}{\leq} \frac{2M_{\lambda}M_{\pi,1}}{(1 - \gamma)^2} + \frac{2M_{\lambda}M_{\pi,1}}{(1 - \gamma)^2} = \frac{4M_{\lambda}M_{\pi,1}}{(1 - \gamma)^2}, \quad (137)$$

where the second term in (i) follows from the bound in (132) with $\|r\|_{\infty} = \|r_{f,t-1}\|_{\infty} = \|\nabla_{\lambda} f(\lambda_{t-1})\|_{\infty} \leq M_{\lambda}$, i.e., M_{λ} -Lipschitz continuity of $f(\cdot)$ in Assumption 3. The first term in (i) comes from an upper bound on $\|\widehat{d}(\tau_t, \theta_t, r_{f,t-1})\|$, which can be derived as follows

$$\begin{aligned} \|\widehat{d}(\tau_t, \theta_t, r_{f,t-1})\| &\stackrel{(i)}{=} \left\| \sum_{h=0}^{H-1} \left(\sum_{k=h}^{H-1} \gamma^k r_{f,t-1}(s_k, a_k) \right) \nabla \log \pi_{\theta_t}(a_h | s_h) \right\| \\ &\leq \sum_{h=0}^{H-1} \sum_{k=h}^{H-1} \gamma^k \|r_{f,t-1}\|_{\infty} \cdot \|\nabla \log \pi_{\theta_t}(a_h | s_h)\| \\ &\leq M_{\lambda} \sum_{h=0}^{H-1} \sum_{k=h}^{H-1} \gamma^k \|\nabla \log \pi_{\theta_t}(a_h | s_h)\| \\ &\stackrel{(ii)}{\leq} 2M_{\lambda}M_{\pi,1} \sum_{h=0}^{H-1} \sum_{k=h}^{H-1} \gamma^k \\ &= 2M_{\lambda}M_{\pi,1} \sum_{k=0}^{H-1} \sum_{h=0}^k \gamma^k \\ &\leq 2M_{\lambda}M_{\pi,1} \sum_{k=0}^{H-1} (k+1) \gamma^k \\ &\leq \frac{2M_{\lambda}M_{\pi,1}}{(1 - \gamma)^2}. \end{aligned} \quad (138)$$

where (i) is by the definition of the policy gradient estimator in (15), and (ii) applies the bound on the score function $\nabla_{\theta} \log \pi_{\theta}(a|s)$ from Lemma 19. Hence, (137) indicates that

$$\mathbb{E}[\|\widehat{y}_t\|^2] \leq \frac{16(M_{\lambda}M_{\pi,1})^2}{(1 - \gamma)^4}. \quad (139)$$

Next, we focus on $\mathbb{E}[\|\widehat{z}_t\|^2]$. Since $\mathbb{E}[\widehat{z}_t] = 0$, it follows from the same logic as (126) that

$$\begin{aligned}
\mathbb{E}[\|\widehat{z}_t\|^2] &\leq \mathbb{E}\left[\left\|\widehat{d}(\tau_t, \theta_t, r_{f,t-1}) - w(\tau_t|\theta_{t-1}, \theta_t)\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\right\|^2\right] \\
&= \mathbb{E}\left[\left\|\widehat{d}(\tau_t, \theta_t, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1}) + \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\right.\right. \\
&\quad \left.\left.+ (1 - w(\tau_t|\theta_{t-1}, \theta_t))\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\right\|^2\right] \\
&\leq 3\mathbb{E}\left[\left\|\widehat{d}(\tau_t, \theta_t, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1})\right\|^2\right] + 3\mathbb{E}\left[\left\|\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\right\|^2\right] \\
&\quad + 3\mathbb{E}\left[(1 - w(\tau_t|\theta_{t-1}, \theta_t))^2\left\|\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\right\|^2\right].
\end{aligned} \tag{140}$$

where we apply the inequality $\|a + b + c\|^2 \leq 3\|a\|^2 + 3\|b\|^2 + 3\|c\|^2$ in the last step. Below, we derive some upper bounds for the three terms on the right-hand side of (140) separately.

1. The first term of (140):

$$\begin{aligned}
\left\|\widehat{d}(\tau_t, \theta_t, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1})\right\| &\stackrel{(i)}{\leq} \frac{2(M_{\pi,1}^2 + M_{\pi,2})}{(1-\gamma)^2} \|r_{f,t-1}\|_\infty \cdot \|\theta_t - \theta_{t-1}\| \\
&= \frac{2(M_{\pi,1}^2 + M_{\pi,2})}{(1-\gamma)^2} \|\nabla_\lambda f(\lambda_{t-1})\|_\infty \cdot \eta_\theta \\
&\leq \frac{2M_\lambda(M_{\pi,1}^2 + M_{\pi,2})}{(1-\gamma)^2} \eta_\theta,
\end{aligned} \tag{141}$$

where inequality (i) follows from Lemma 20, and last line is due to M_λ -Lipschitz continuity of $f(\cdot)$ from Assumption 3. Therefore, with $\eta_\theta = T^{-1/2}$, we have that

$$\mathbb{E}\left[\left\|\widehat{d}(\tau_t, \theta_t, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1})\right\|^2\right] \leq \frac{4M_\lambda^2(M_{\pi,1}^2 + M_{\pi,2})^2}{(1-\gamma)^4} \eta_\theta^2 = \frac{4M_\lambda^2(M_{\pi,1}^2 + M_{\pi,2})^2}{(1-\gamma)^4 T}. \tag{142}$$

2. The second term of (140):

$$\begin{aligned}
\left\|\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\right\| &= \left\|\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1} - r_{f,t-2})\right\| \\
&\stackrel{(i)}{\leq} \sum_{h=0}^{H-1} \sum_{k=h}^{H-1} \gamma^k \|r_{f,t-1} - r_{f,t-2}\|_\infty \cdot \|\nabla \log \pi_{\theta_{t-1}}(a_h|s_h)\| \\
&\stackrel{(ii)}{\leq} \frac{2M_{\pi,1}}{(1-\gamma)^2} \|r_{f,t-1} - r_{f,t-2}\|_\infty \\
&= \frac{2M_{\pi,1}}{(1-\gamma)^2} \|\nabla_\lambda f(\lambda_{t-1}) - \nabla_\lambda f(\lambda_{t-2})\|_\infty \\
&\stackrel{(iii)}{\leq} \frac{2M_{\pi,1}\ell_\lambda}{(1-\gamma)^2} \|\lambda_{t-1} - \lambda_{t-2}\|,
\end{aligned} \tag{143}$$

where steps (i) and (ii) follow from the same reasoning as (138), and step (iii) uses the ℓ_λ -smoothness of $f(\cdot)$ from Assumption 3. Utilizing (143), we obtain that

$$\mathbb{E}\left[\left\|\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\right\|^2\right] \leq \frac{4M_{\pi,1}^2\ell_\lambda^2}{(1-\gamma)^4} \cdot \mathbb{E}[\|\lambda_{t-1} - \lambda_{t-2}\|^2]. \tag{144}$$

Hence, it remains to show the boundedness of $\mathbb{E}[\|\lambda_{t-1} - \lambda_{t-2}\|^2]$. Based on the update rule of λ_t , we observe that

$$\begin{aligned}\lambda_t - \lambda_{t-1} &= \widehat{\lambda}(\tau_t) + (1 - \alpha_t)[\lambda_{t-1} - w(\tau_t|\theta_{t-1}, \theta_t)\widehat{\lambda}(\tau_t)] - \lambda_{t-1} \\ &= \alpha_t(\widehat{\lambda}(\tau_t) - \lambda_t) + \alpha_t(\lambda_t - \lambda_{t-1}) + (1 - \alpha_t)(1 - w(\tau_t|\theta_{t-1}, \theta_t))\widehat{\lambda}(\tau_t).\end{aligned}\quad (145)$$

By rearranging the terms, we obtain that

$$\begin{aligned}\lambda_t - \lambda_{t-1} &= \frac{\alpha_t}{1 - \alpha_t}(\widehat{\lambda}(\tau_t) - \lambda_t) + (1 - w(\tau_t|\theta_{t-1}, \theta_t))\widehat{\lambda}(\tau_t) \\ &= \frac{\alpha_t}{1 - \alpha_t}(\lambda_H(\theta_t) - \lambda_t) + \frac{\alpha_t}{1 - \alpha_t}(\widehat{\lambda}(\tau_t) - \lambda_H(\theta_t)) + (1 - w(\tau_t|\theta_{t-1}, \theta_t))\widehat{\lambda}(\tau_t) \\ &= \frac{\alpha_t}{1 - \alpha_t}e_{\lambda,t} + \frac{\alpha_t}{1 - \alpha_t}(\widehat{\lambda}(\tau_t) - \lambda_H(\theta_t)) + (1 - w(\tau_t|\theta_{t-1}, \theta_t))\widehat{\lambda}(\tau_t).\end{aligned}\quad (146)$$

Using the above equality, we further arrive at

$$\begin{aligned}\mathbb{E}[\|\lambda_t - \lambda_{t-1}\|^2] &\leq \frac{3\alpha_t^2\mathbb{E}[\|e_{\lambda,t}\|^2]}{(1 - \alpha_t)^2} + \frac{3\alpha_t^2\mathbb{E}[\|\lambda_H(\theta_t) - \widehat{\lambda}(\tau_t)\|^2]}{(1 - \alpha_t)^2} + 3\mathbb{E}[\|(1 - w(\tau_t|\theta_{t-1}, \theta_t))\widehat{\lambda}(\tau_t)\|^2] \\ &\stackrel{(i)}{\leq} \frac{75\alpha_t^2C_\lambda}{(1 - \alpha_t)^2\sqrt{t+2}} + \frac{3\alpha_t^2}{(1 - \alpha_t)^2(1 - \gamma)^2} + \frac{3C_w\eta_\theta^2}{(1 - \gamma)^2} \\ &\stackrel{(ii)}{\leq} \frac{75C_\lambda}{(t+1)^{3/2}} + \frac{3}{(1 - \gamma)^2(t+1)} + \frac{3C_w}{(1 - \gamma)^2T} \\ &\leq \frac{78C_\lambda}{t+1} \\ &\lesssim \frac{C_\lambda}{t+1},\end{aligned}\quad (147)$$

where in step (i), we respectively use the bound for $\mathbb{E}[\|e_{\lambda,t}\|^2]$ from (118) in Lemma 12, the bound for $\mathbb{E}[\|\lambda_H(\theta_t) - \widehat{\lambda}(\tau_t)\|^2]$ in (124), and the bound for $\mathbb{E}[\|(1 - w(\tau_t|\theta_{t-1}, \theta_t))\widehat{\lambda}(\tau_t)\|^2]$ in (126). Then, inequality (ii) comes from the choices of $\alpha_t = (t+1)^{-1/2}$ and $\eta_\theta = T^{-1/2}$. Applying (147) to (144), we obtain the following inequality for all $t \geq 1$:

$$\mathbb{E}[\|\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\|^2] \leq \frac{4(M_{\pi,1}\ell_\lambda)^2}{(1 - \gamma)^4} \cdot \frac{78C_\lambda}{t} \leq \frac{624(M_{\pi,1}\ell_\lambda)^2C_\lambda}{(1 - \gamma)^4(t+1)}.\quad (148)$$

The above upper bound also apply to the case $t = 0$, where the term $\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-1}) - \widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})$ actually does not show up in Eq. (140).

3. The third term of (140):

$$\begin{aligned}\mathbb{E}[(1 - w(\tau_t|\theta_{t-1}, \theta_t))^2\|\widehat{d}(\tau_t, \theta_{t-1}, r_{f,t-2})\|^2] &\stackrel{(i)}{\leq} \frac{4(M_\lambda M_{\pi,1})^2}{(1 - \gamma)^4} \cdot \mathbb{E}[(1 - w(\tau_t|\theta_{t-1}, \theta_t))^2] \\ &\stackrel{(ii)}{\leq} \frac{4(M_\lambda M_{\pi,1})^2C_w\eta_\theta^2}{(1 - \gamma)^4} \\ &= \frac{4(M_\lambda M_{\pi,1})^2C_w}{(1 - \gamma)^4T},\end{aligned}\quad (149)$$

where step (i) follows from (138), and step (ii) applies Lemma 9 in a similar way as (126).

Combining the upper bounds for the three terms shown in (142), (148), and (149), we finalize the bound for $\mathbb{E}[\|\widehat{z}_t\|^2]$ in (140) as follows

$$\mathbb{E}[\|\widehat{z}_t\|^2] \leq \frac{12M_\lambda^2(M_{\pi,1}^2 + M_{\pi,2})^2}{(1-\gamma)^4 T} + \frac{1872(M_{\pi,1}\ell_\lambda)^2 C_\lambda}{(1-\gamma)^4(t+1)} + \frac{12(M_\lambda M_{\pi,1})^2 C_w}{(1-\gamma)^4 T}. \quad (150)$$

Finally, by substituting the bounds for $\mathbb{E}[\|\widehat{y}_t\|^2]$ in (139) and $\mathbb{E}[\|\widehat{z}_t\|^2]$ in (150) into (136), we conclude that

$$\begin{aligned} \mathbb{E}[\|\widehat{e}_{f,t}\|^2] &\leq (1-\alpha_t)^2 \mathbb{E}[\|\widehat{e}_{f,t-1}\|^2] + 2\alpha_t^2 \cdot \frac{16(M_\lambda M_{\pi,1})^2}{(1-\gamma)^4} \\ &\quad + 2(1-\alpha_t)^2 \left(\frac{12M_\lambda^2(M_{\pi,1}^2 + M_{\pi,2})^2}{(1-\gamma)^4 T} + \frac{1872(M_{\pi,1}\ell_\lambda)^2 C_\lambda}{(1-\gamma)^4(t+1)} + \frac{12(M_\lambda M_{\pi,1})^2 C_w}{(1-\gamma)^4 T} \right) \\ &\leq (1-\alpha_t)^2 \mathbb{E}[\|\widehat{e}_{\lambda,t-1}\|^2] + \frac{32(M_\lambda M_{\pi,1})^2}{(1-\gamma)^4} \cdot \frac{1}{t+1} \\ &\quad + \frac{24}{(1-\gamma)^4} \left(\frac{M_\lambda^2 [C_w M_{\pi,1}^2 + (M_{\pi,1}^2 + M_{\pi,2})^2]}{T} + \frac{156(M_{\pi,1}\ell_\lambda)^2 C_\lambda}{t+1} \right). \end{aligned} \quad (151)$$

Therefore, $\mathbb{E}[\|\widehat{e}_{f,t}\|^2]$ satisfies the recursion $\mathbb{E}[\|\widehat{e}_{f,t}\|^2] \leq (1-\alpha_t)^2 \mathbb{E}[\|\widehat{e}_{f,t-1}\|^2] + \beta_t$ with

$$\beta_t \leq \underbrace{\frac{32(M_\lambda M_{\pi,1})^2 + 24 \{M_\lambda^2 [C_w M_{\pi,1}^2 + (M_{\pi,1}^2 + M_{\pi,2})^2] + 156(M_{\pi,1}\ell_\lambda)^2 C_\lambda\}}{(1-\gamma)^4}}_{=:C_\beta} \cdot \frac{1}{t+1}, \quad (152)$$

where we define C_β to be the coefficient of $1/(t+1)$ in the above upper bound. Hence, by applying Lemma 11, we have that

$$\begin{aligned} \mathbb{E}[\|\widehat{e}_{f,t}\|^2] &\leq \frac{\mathbb{E}[\|\widehat{e}_{f,0}\|^2]}{t+1} + \frac{12C_\beta}{\sqrt{t+2}} \\ &\stackrel{(i)}{\leq} \frac{16(M_\lambda M_{\pi,1})^2}{(1-\gamma)^4(t+1)} + \frac{12C_\beta}{\sqrt{t+2}} \\ &\leq \underbrace{\frac{M_\lambda^2 [(1+C_w)M_{\pi,1}^2 + (M_{\pi,1}^2 + M_{\pi,2})^2] + (M_{\pi,1}\ell_\lambda)^2 C_\lambda}{(1-\gamma)^4}}_{=:C_{\widehat{e}}} \cdot \frac{1}{\sqrt{t+2}}, \end{aligned} \quad (153)$$

where step (i) uses the upper bound $\|\widehat{e}_{f,0}\| \leq \|d_{f,0}\| + \|\langle \nabla_{\theta} \lambda_H(\theta_0), r_{f,-1} \rangle\| \leq \frac{4M_\lambda M_{\pi,1}}{(1-\gamma)^2}$ with the derivations identical to (137).

Lastly, combining (127), (133), and (153), we conclude that

$$\begin{aligned}
\mathbb{E} [\|e_{f,t}\|] &\lesssim \frac{\sqrt{C_{\tilde{e}}}}{(t+2)^{1/4}} + \frac{M_{\pi,1}\ell_\lambda\sqrt{C_\lambda}}{(1-\gamma)^2} \cdot \frac{1}{(t+1)^{1/4}} + \frac{M_{\pi,1}^2\ell_\lambda}{(1-\gamma)^4} \cdot \frac{1}{\sqrt{T}} \\
&\lesssim \left(\sqrt{C_{\tilde{e}}} + \frac{M_{\pi,1}\ell_\lambda\sqrt{C_\lambda}}{(1-\gamma)^2} \right) \frac{1}{(t+1)^{1/4}} \\
&=: \underbrace{\frac{\sqrt{M_\lambda^2 [(1+C_w)M_{\pi,1}^2 + (M_{\pi,1}^2 + M_{\pi,2})^2] + (M_{\pi,1}\ell_\lambda)^2 C_\lambda + M_{\pi,1}\ell_\lambda\sqrt{C_\lambda}}}{(1-\gamma)^2}}_{=:C_e} \cdot \frac{1}{(t+1)^{1/4}}.
\end{aligned} \tag{154}$$

Since function $g(\lambda(\theta))$ satisfies the same property as $f(\lambda(\theta))$, we know that the same upper bound also applies to $\|e_{g,t}\|$. This completes the proof of Lemma 13. \square

LEMMA 14. Under Algorithm 1, let $\{e_{L,t}\}_{t=0}^{T-1}$ be the sequence defined as $e_{L,t} := d_{L,t} - \nabla_{\theta} L(\lambda_H(\theta_t), \mu_t)$, for all $t \geq 0$. Then, it holds that

$$\mathbb{E} [\|e_{L,t}\|] = \mathbb{E} [\|e_{f,t} + \mu_t e_{g,t}\|] \leq \mathbb{E} [\|e_{f,t}\|] + \mu_t \mathbb{E} [\|e_{g,t}\|] \lesssim \frac{C_e(1+C_0)}{(t+1)^{1/4}}, \tag{155}$$

where C_e is defined in (34b).

PROOF. To control $e_{L,t}$, we notice the following decomposition:

$$\begin{aligned}
e_{L,t} &= d_{L,t} - \nabla_{\theta} L(\lambda_H(\theta_t), \mu_t) \\
&= (d_{f,t} + \mu_t d_{g,t}) - [\nabla_{\theta} f(\lambda_H(\theta_t)) + \mu_t \nabla_{\theta} g(\lambda_H(\theta_t))] \\
&= [d_{f,t} - f(\lambda_H(\theta_t))] + \mu_t [d_{g,t} - g(\lambda_H(\theta_t))] \\
&= e_{f,t} + \mu_t e_{g,t}.
\end{aligned} \tag{156}$$

Then, according to Lemma 13, we have that

$$\mathbb{E} [\|e_{L,t}\|] = \mathbb{E} [\|e_{f,t} + \mu_t e_{g,t}\|] \leq \mathbb{E} [\|e_{f,t}\|] + \mu_t \mathbb{E} [\|e_{g,t}\|] \stackrel{(i)}{\lesssim} \frac{C_e(1+C_0)}{(t+1)^{1/4}}, \tag{157}$$

where step (i) is by the fact that $\mu_t \in [0, C_0]$. This completes the proof of Lemma 14. \square

D Supplementary Materials for Section 5

This section is dedicated to the proof of the zero constraint violation in Section 5.

D.1 Proof of Theorem 3

PROOF. We begin with general arguments that apply to both concave and strongly concave cases. Firstly, by the assumption $\delta < \xi$, it holds that

$$g_\delta(\lambda(\theta)) = g(\lambda(\theta)) - \delta \implies |g_\delta(\lambda(\theta))| \leq M + \xi \text{ and } g_\delta(\lambda(\tilde{\theta})) \geq \xi - \delta > 0, \tag{158}$$

which gives the constraint upper bound and strict feasibility for problem (36).

Let θ_δ^* be an optimal solution to the pessimistic problem (36). Then,

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_t))] &= \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_\delta^*))] + \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta_\delta^*)) - f(\lambda(\theta_t))] \\ &= [f(\lambda(\theta^*)) - f(\lambda(\theta_\delta^*))] + \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta_\delta^*)) - f(\lambda(\theta_t))]. \end{aligned} \quad (159)$$

To upper-bound the first term in (159), let θ_δ be a policy parameter that satisfies

$$\lambda(\theta_\delta) := \frac{\xi - \delta}{\xi} \lambda(\theta^*) + \frac{\delta}{\xi} \lambda(\tilde{\theta}), \quad (160)$$

where we assume $0 < \delta < \xi$. It is easy to verify that θ_δ is a feasible point to (36):

$$\begin{aligned} g_\delta(\lambda(\theta_\delta)) &= g(\lambda(\theta_\delta)) - \delta \\ &= g\left(\frac{\xi - \delta}{\xi} \lambda(\theta^*) + \frac{\delta}{\xi} \lambda(\tilde{\theta})\right) - \delta \\ &\stackrel{(i)}{\geq} \frac{\xi - \delta}{\xi} g(\lambda(\theta^*)) + \frac{\delta}{\xi} g(\lambda(\tilde{\theta})) - \delta \\ &\stackrel{(ii)}{\geq} 0 + \frac{\delta}{\xi} \cdot \xi - \delta \\ &= 0, \end{aligned}$$

where (i) follows from the concavity of $g(\cdot)$ and (ii) uses the feasibility of θ^* to (11) as well as Assumption 1. We remark that the policy corresponds to $\lambda(\theta_\delta)$ is unique and given by

$$\pi_{\theta_\delta}(a|s) = \frac{\lambda(\theta_\delta; s, a)}{\sum_{a' \in \mathcal{A}} \lambda(\theta_\delta; s, a')}.$$

On the other hand, due to the assumption of over-parameterization (see Assumption 2), θ_δ may not be unique, and it suffices to choose one such θ_δ that satisfies (160).

The feasibility of θ_δ to (36) implies that $f(\lambda(\theta_\delta^*)) \geq f(\lambda(\theta_\delta))$. Consequently,

$$\begin{aligned} f(\lambda(\theta^*)) - f(\lambda(\theta_\delta^*)) &\leq f(\lambda(\theta^*)) - f(\lambda(\theta_\delta)) \\ &= f(\lambda(\theta^*)) - f\left(\frac{\xi - \delta}{\xi} \lambda(\theta^*) + \frac{\delta}{\xi} \lambda(\tilde{\theta})\right) \\ &\leq f(\lambda(\theta^*)) - \left(\frac{\xi - \delta}{\xi} f(\lambda(\theta^*)) + \frac{\delta}{\xi} f(\lambda(\tilde{\theta}))\right) \\ &= \frac{\delta}{\xi} [f(\lambda(\theta^*)) - f(\lambda(\tilde{\theta}))] \\ &\leq \frac{2\delta M}{\xi}. \end{aligned} \quad (161)$$

By (158), for any $\delta < \xi$, choosing $C_{0,\delta} = 1 + \frac{M - f(\lambda(\tilde{\theta}))}{(\xi - \delta)}$ ensures that the optimal dual variable of problem (36) belongs to the dual feasible region $U = [0, C_{0,\delta}]$ (see Lemma 1).

Now, we show the first part of Theorem 3. When $f(\cdot)$ is a general concave function, we apply (29a) and (29b) in Theorem 1 to obtain that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta_t^*)) - f(\lambda(\theta_t))] &\leq \frac{2M + (M + \xi)^2/2}{T^{2/3}} + \frac{2(1 + C_{0,\delta})\ell_\theta L_\lambda^2}{(1 - \gamma)^2 T^{1/3}} + \frac{2(M + \xi)^2}{T^{1/3}}, \\ \frac{1}{T} \left[\sum_{t=0}^{T-1} -g_\delta(\lambda(\theta_t)) \right]_+ &\leq \frac{2M + (M + \xi)^2/2}{T^{2/3}} + \frac{2(1 + C_{0,\delta})\ell_\theta L_\lambda^2}{(1 - \gamma)^2 T^{1/3}} + \frac{2(M + \xi)^2 + C_{0,\delta}^2/2}{T^{1/3}}, \end{aligned} \quad (162)$$

where we replace the upper bound M on $|g(\cdot)|$ by the new upper bound $M + \xi$ for $|g_\delta(\cdot)|$. Therefore, together with (159) and (161), we have the following optimality gap for the original problem (11):

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta_t^*)) - f(\lambda(\theta_t))] &\leq \frac{2\delta M}{\xi} + \frac{2M + (M + \xi)^2/2}{T^{2/3}} + \frac{2(1 + C_{0,\delta})\ell_\theta L_\lambda^2}{(1 - \gamma)^2 T^{1/3}} + \frac{2(M + \xi)^2}{T^{1/3}} \\ &= \mathcal{O}(\delta) + \mathcal{O}(T^{-1/3}). \end{aligned} \quad (163)$$

For the constraint violation, we have that

$$\begin{aligned} \frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ &= \frac{1}{T} \left[\sum_{t=0}^{T-1} -(g(\lambda(\theta_t)) - \delta) - \delta \right]_+ \\ &= \left[\frac{1}{T} \left[\sum_{t=0}^{T-1} -(g(\lambda(\theta_t)) - \delta) \right]_+ - \delta \right]_+ \\ &= \left[\frac{1}{T} \left[\sum_{t=0}^{T-1} -g_\delta(\lambda(\theta_t)) \right]_+ - \delta \right]_+ \\ &\leq \left[\frac{2M + (M + \xi)^2/2}{T^{2/3}} + \frac{2(1 + C_{0,\delta})\ell_\theta L_\lambda^2}{(1 - \gamma)^2 T^{1/3}} + \frac{2(M + \xi)^2 + C_{0,\delta}^2/2}{T^{1/3}} - \delta \right]_+. \end{aligned} \quad (164)$$

By choosing δ such that

$$\frac{2M + (M + \xi)^2/2}{T^{2/3}} + \frac{2(1 + C_{0,\delta})\ell_\theta L_\lambda^2}{(1 - \gamma)^2 T^{1/3}} + \frac{2(M + \xi)^2 + C_{0,\delta}^2/2}{T^{1/3}} - \delta = 0, \quad (165)$$

the constraint violation (164) becomes 0. Such constant δ must exist since the summation of the first three terms in (165) decreases in the order of $\mathcal{O}(T^{-1/3})$ as T increases. As (165) implies $\delta = \mathcal{O}(T^{-1/3})$, the convergence rate of the optimality gap (163) remains $\mathcal{O}(T^{-1/3})$. Finally, we remark that the requirement $\delta < \xi$ is naturally satisfied when T is large.

When $f(\cdot)$ is σ -strongly concave, we apply (30a) and (30b) in Theorem 1 to obtain that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta_t^*)) - f(\lambda(\theta_t))] &\leq \frac{2M + C_{0,\delta}(M + \xi)}{\tilde{\epsilon}T} + \left(\frac{(M + \xi)^2}{\tilde{\epsilon}} + \frac{(M + \xi)^2}{2} \right) \frac{1}{\sqrt{T}}, \\ \frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ &\leq \frac{2M + C_{0,\delta}(M + \xi)}{\tilde{\epsilon}T} + \left(\frac{(M + \xi)^2}{\tilde{\epsilon}} + \frac{(M + \xi)^2 + C_{0,\delta}^2}{2} \right) \frac{1}{\sqrt{T}}. \end{aligned}$$

We note that, similar to (77), the numerator term $2M + C_{0,\delta}(M + \xi)$ comes from the new upper bound on the term $L_\delta(\lambda(\theta_T), \mu_{T-1}) - L_\delta(\lambda(\theta_0), \mu_0)$. Together with (159) and (161), we derive the optimality gap as

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} [f(\lambda(\theta^*)) - f(\lambda(\theta_t))] &\leq \frac{2\delta M}{\xi} + \frac{2M + C_{0,\delta}(M + \xi)}{\tilde{\epsilon}T} + \left(\frac{(M + \xi)^2}{\tilde{\epsilon}} + \frac{(M + \xi)^2}{2} \right) \frac{1}{\sqrt{T}} \\ &= \mathcal{O}(\delta) + \mathcal{O}(T^{-1/2}). \end{aligned} \quad (166)$$

For the constraint violation, similarly as (164), we have that

$$\frac{1}{T} \left[\sum_{t=0}^{T-1} -g(\lambda(\theta_t)) \right]_+ \leq \left[\frac{2M + C_{0,\delta}(M + \xi)}{\tilde{\epsilon}T} + \left(\frac{(M + \xi)^2}{\tilde{\epsilon}} + \frac{(M + \xi)^2 + C_{0,\delta}^2}{2} \right) \frac{1}{\sqrt{T}} - \delta \right]_+. \quad (167)$$

We choose δ such that

$$\frac{2M + C_{0,\delta}(M + \xi)}{\tilde{\epsilon}T} + \left(\frac{(M + \xi)^2}{\tilde{\epsilon}} + \frac{(M + \xi)^2 + C_{0,\delta}^2}{2} \right) \frac{1}{\sqrt{T}} - \delta = 0, \quad (168)$$

which guarantees the constraint violation (167) to be zero. As (168) implies $\delta = \mathcal{O}(T^{-1/2})$, the convergence rate of the optimality gap (166) is $\mathcal{O}(T^{-1/2})$. This completes the proof. \square

E Auxiliary Lemmas

In this section, we present a few auxiliary lemmas used to prove main results in this paper. These lemmas are standard results on Markov decision processes. We refer the reader to Section 2 for necessary definitions and [4] for the proofs of these results.

LEMMA 15 (POLICY GRADIENT UNDER GENERAL PARAMETERIZATION). *Let $V^{\pi_\theta}(r)$ be the value function under policy π_θ with reward $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. The gradient of $V^{\pi_\theta}(r)$ with respect to θ can be given by the following three equivalent forms:*

$$\begin{aligned} \nabla_\theta V^{\pi_\theta}(r) &= \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^{\pi_\theta}} \mathbb{E}_{a \sim \pi_\theta(\cdot|s)} [\nabla_\theta \log \pi_\theta(a|s) \cdot Q^{\pi_\theta}(r; s, a)] \\ &= \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k \nabla_\theta \log \pi_\theta(a_k | s_k) \cdot Q^{\pi_\theta}(r; s_k, a_k) \middle| \pi_\theta, s_0 \sim \rho \right] \\ &= \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k \cdot r(s_k, a_k) \cdot \left(\sum_{k'=0}^k \nabla_\theta \log \pi_\theta(a_{k'} | s_{k'}) \right) \middle| \pi_\theta, s_0 \sim \rho \right] \end{aligned}$$

LEMMA 16 (POLICY GRADIENT UNDER DIRECT PARAMETERIZATION). *Let $V^\pi(r)$ be the value function under policy π with reward $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. The gradient of $V^\pi(r)$ with respect to π is*

$$\frac{\partial V^\pi(r)}{\partial \pi(a|s)} = \frac{1}{1-\gamma} d^\pi(s) \cdot Q^\pi(r; s, a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

LEMMA 17 (SMOOTHNESS OF $V^\pi(r)$ W.R.T. π). *Let $V^\pi(r)$ be the value function under policy π with reward $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. For every two policies π and π' , it holds that*

$$\left\| \nabla_\pi V^\pi(r) - \nabla_{\pi'} V^{\pi'}(r) \right\| \leq \frac{4\gamma|\mathcal{A}|}{(1-\gamma)^3} \cdot \|r\|_\infty \cdot \|\pi - \pi'\|.$$

LEMMA 18 (PERFORMANCE DIFFERENCE). Let $V^\pi(r)$ be the value function under policy π with reward $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. For every two policies π and π' , it holds that

$$\begin{aligned} V^{\pi'}(r) - V^\pi(r) &= \langle r, \lambda^{\pi'} - \lambda^\pi \rangle \\ &= \frac{1}{1-\gamma} \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} (\pi'(a|s) - \pi(a|s)) \cdot Q^{\pi'}(r; s, a) \\ &= \frac{1}{1-\gamma} \sum_{s \in \mathcal{S}} d^{\pi'}(s) \sum_{a \in \mathcal{A}} \pi'(a|s) \cdot A^\pi(r; s, a), \end{aligned}$$

where $A^\pi(r; s, a)$ denotes the advantage function with reward r under policy π , defined as

$$A^\pi(r; s, a) := Q^\pi(r; s, a) - \mathbb{E}_{a \sim \pi(\cdot|s)} [Q^\pi(r; s, a)], \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (169)$$

LEMMA 19. Under the general soft-max parameterization (10), let (24) of Assumption 2 hold. Then, for every $\theta \in \Theta$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have that

$$\|\nabla_\theta \log \pi_\theta(a|s)\| \leq 2M_{\pi,1}; \quad (170)$$

$$\|\nabla_\theta^2 \log \pi_\theta(a|s)\| \leq 2(M_{\pi,1}^2 + M_{\pi,2}). \quad (171)$$

LEMMA 20. Let Assumptions 2 and 3 hold. Then, for every $\theta_1, \theta_2 \in \mathbb{R}^k$, we have that

$$\|\lambda_H(\theta_1) - \lambda_H(\theta_2)\| \leq \frac{2M_{\pi,1}}{(1-\gamma)^2} \|\theta_1 - \theta_2\|; \quad (172)$$

$$\|\widehat{d}(\tau, \theta_1, r) - \widehat{d}(\tau, \theta_2, r)\| \leq \frac{2(M_{\pi,1}^2 + M_{\pi,2})}{(1-\gamma)^2} \|r\|_\infty \cdot \|\theta_1 - \theta_2\|, \quad \forall r \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}. \quad (173)$$

F Experiment Details

This section elaborates on the details of the environment and experiments from Section 6.

F.1 Details on Environments

The environments are 8×8 and 20×20 tabular MDPs where $(0, 0)$ is a fixed initial state (marked as S) and they have a fixed goal point at $(7, 7)$ and $(19, 19)$, respectively (See Figure 1a for the coordinate). For the coordination, we set the upper left point as the origin point, i.e., $(0, 0)$, and each block has a length of 1. The agent executes one of the four actions of up, left, right, and down. For every step, the agent receives a -0.2 reward, and if it reaches the goal state, it receives a $+100$ reward.

F.2 Details on Experiments

We elaborate on the experiment's hyperparameters in this subsection.

F.2.1 8×8 Grid World.

- Iteration number $T = 10000$.
- Trajectory length $H = 14$.
- Gamma $\gamma = 0.9$.
- Initial primal step-size $\eta_\theta \in \{0.23, 0.24, 0.25\}$.
- Initial dual step-size $\eta_\mu = 0.1$.
- Step size $\alpha_t = 0.1$ for all $t \geq 1$.
- Initial dual variable $\mu_0 = 1$.

- Dual variable interval upper bound $C_0 = 10$ (Line 14 of Algorithm 1).
- Constraint right-hand side value $d_0 \in \{0.001, 0.005, 0.01, 0.05\}$ (see (42)).

F.2.2 20×20 Grid World.

- Iteration number $T = 10000$.
- Trajectory length $H = 38$.
- Gamma $\gamma = 0.9$.
- Initial primal step-size $\eta_\theta = 0.22$.
- Initial dual step-size $\eta_\mu \in \{0.01, 0.02, 0.03, 0.04, 0.05\}$.
- Step size $\alpha_t \in \{0.03, 0.05, 0.07, 0.09\}$ for all $t \geq 1$.
- Initial dual variable $\mu_0 = 1$.
- Dual variable interval upper bound $C_0 = 10$ (Line 14 of Algorithm 1).
- Constraint right-hand side value $d_0 \in \{0.001, 0.005, 0.01\}$ (see (42)).

G Reproducibility Checklist for JAIR

Select the answers that apply to your research – one per item.

All articles:

- (1) All claims investigated in this work are clearly stated. [yes]
- (2) Clear explanations are given how the work reported substantiates the claims. [yes]
- (3) Limitations or technical assumptions are stated clearly and explicitly. [yes]
- (4) Conceptual outlines and/or pseudo-code descriptions of the AI methods introduced in this work are provided, and important implementation details are discussed. [yes]
- (5) Motivation is provided for all design choices, including algorithms, implementation choices, parameters, data sets and experimental protocols beyond metrics. [yes]

Articles containing theoretical contributions:

Does this paper make theoretical contributions? [yes]

If yes, please complete the list below.

- (1) All assumptions and restrictions are stated clearly and formally. [yes]
- (2) All novel claims are stated formally (e.g., in theorem statements). [yes]
- (3) Proofs of all non-trivial claims are provided in sufficient detail to permit verification by readers with a reasonable degree of expertise (e.g., that expected from a PhD candidate in the same area of AI). [yes]
- (4) Complex formalism, such as definitions or proofs, is motivated and explained clearly. [yes]
- (5) The use of mathematical notation and formalism serves the purpose of enhancing clarity and precision; gratuitous use of mathematical formalism (i.e., use that does not enhance clarity or precision) is avoided. [yes]
- (6) Appropriate citations are given for all non-trivial theoretical tools and techniques. [yes]

Articles reporting on computational experiments:

Does this paper include computational experiments? [yes]

If yes, please complete the list below.

- (1) All source code required for conducting experiments is included in an online appendix or will be made publicly available upon publication of the paper. The online appendix follows best practices for source code readability and documentation as well as for long-term accessibility. [yes]

- (2) The source code comes with a license that allows free usage for reproducibility purposes. [yes]
- (3) The source code comes with a license that allows free usage for research purposes in general. [yes]
- (4) Raw, unaggregated data from all experiments is included in an online appendix or will be made publicly available upon publication of the paper. The online appendix follows best practices for long-term accessibility. [NA]
- (5) The unaggregated data comes with a license that allows free usage for reproducibility purposes. [NA]
- (6) The unaggregated data comes with a license that allows free usage for research purposes in general. [NA]
- (7) If an algorithm depends on randomness, then the method used for generating random numbers and for setting seeds is described in a way sufficient to allow replication of results. [yes]
- (8) The execution environment for experiments, the computing infrastructure (hardware and software) used for running them, is described, including GPU/CPU makes and models; amount of memory (cache and RAM); make and version of operating system; names and versions of relevant software libraries and frameworks. [no]
- (9) The evaluation metrics used in experiments are clearly explained and their choice is explicitly motivated. [yes]
- (10) The number of algorithm runs used to compute each result is reported. [yes]
- (11) Reported results have not been “cherry-picked” by silently ignoring unsuccessful or unsatisfactory experiments. [yes]
- (12) Analysis of results goes beyond single-dimensional summaries of performance (e.g., average, median) to include measures of variation, confidence, or other distributional information. [yes]
- (13) All (hyper-) parameter settings for the algorithms/methods used in experiments have been reported, along with the rationale or method for determining them. [yes]
- (14) The number and range of (hyper-) parameter settings explored prior to conducting final experiments have been indicated, along with the effort spent on (hyper-) parameter optimisation. [NA]
- (15) Appropriately chosen statistical hypothesis tests are used to establish statistical significance in the presence of noise effects. [NA]

Articles using data sets:

Does this work rely on one or more data sets (possibly obtained from a benchmark generator or similar software artifact)? [no]

Received 15 January 2025; revised 16 June 2025; accepted 18 June 2025