
Risk-Averse Constrained Reinforcement Learning with Optimized Certainty Equivalents

Jane H. Lee

Department of Computer Science
Yale University
jane.h.lee@yale.edu

Baturay Saglam

Department of Electrical Engineering
Yale University
baturay.saglam@yale.edu

Spyridon Pougkakiotis

Department of Mathematics
King's College London
spyridon.pougkakiotis@kcl.ac.uk

Amin Karbasi*

Foundation AI
Cisco Systems Inc.
karbasi@cisco.com

Dionysis Kalogerias

Department of Electrical Engineering
Yale University
dionysis.kalogerias@yale.edu

Abstract

Constrained optimization provides a common framework for dealing with conflicting objectives in reinforcement learning (RL). In most of these settings, the objectives (and constraints) are expressed through the expected accumulated reward. However, this formulation neglects risky or even possibly catastrophic events at the tails of the reward distribution, and is often insufficient for high-stakes applications in which the risk involved in outliers is critical. In this work, we propose a framework for risk-aware constrained RL, which exhibits per-stage robustness properties jointly in reward values and time using optimized certainty equivalents (OCEs). Our framework ensures an exact equivalent to the original constrained problem within a parameterized strong Lagrangian duality framework under appropriate constraint qualifications, and yields a simple algorithmic recipe which can be wrapped around standard RL solvers, such as PPO. Lastly, we establish the convergence of the proposed algorithm under common assumptions, and verify the risk-aware properties of our approach through several numerical experiments.

1 Introduction

Autonomous agents are often used in settings where they must handle several conflicting requirements while maximizing a main objective, such as navigating a maze without hitting any walls. These conflicting requirements are often modeled with constrained reinforcement learning (RL) (see Sutton and Barto [42]) and are either solved as a multi-objective problem with conflicting requirements controlled by weights, chosen manually or through hyperparameter tuning (e.g., see Borkar [10], Achiam et al. [2], Mannor and Shimkin [28], Tessler et al. [45], Mania et al. [27]), or by using primal-dual methods which optimize over the weights [45, 5]. Others have approached constrained RL through reward-free techniques, e.g., [30], or strategic exploration [54]. The work of Paternain et al. [33] has given theoretical support to primal-dual approaches by proving that, despite their

*Work done while affiliated with Yale University.

inherent nonconvexity, constrained RL models exhibit zero duality gap under mild assumptions. However, in the standard (risk-neutral) setting (including that of Paternain et al. [33]), the objectives and constraints are often formulated as expected values of a (discounted) sum of rewards. Nonetheless, in many practical scenarios, especially in the context of safety [24] and other high-stakes applications, the standard expectation may not be strict enough to describe the desired behavior. For example, instead of maximizing the on-average return (e.g., of an investment), one may want to learn policies that mitigate risk (e.g., of losing money).

Stochastic Optimization with Risk Measures Risk measures are often used to quantify this risk in problems with uncertain (random) outcomes (see, e.g., Shapiro et al. [41]). More broadly, stochastic optimization with risk measures in the objectives and/or constraints has been studied in the online setting [26], bandits [11, 37, 44], statistical learning [3, 15, 32, 48], non-convex resource allocation problems [21], and others [20, 55]. These methods and analyses are not readily applicable in our setting either due to non-convexity/concavity in the objective and constraints or due to the fact that the risk envelope for our problem has dependence on the policy π .

Risk-Averse/Risk-Aware RL In reinforcement learning specifically, there is a deep literature on risk-averse methods [31, 53, 13, 14, 29, 35, 49, 51, 50] with many applications in the area of robust control [52, 23]. While the most conservative approach consists of learning policies robust to worst-case scenarios [1, 18], we often prefer to take into account a set of events with significant probability. To do this, Huang et al. [17], Tamar et al. [43] consider the policy gradient method under coherent risk measures, where the risk measure replaces the expectation over the discounted sum of rewards. The work of Chow et al. [12] handles similar CVaR risk objectives with CVaR risk constraints by utilizing Lagrangian relaxation. Xu et al. [51], Wang et al. [50] study risk-sensitive RL with OCEs. Recently, Bonetti et al. [7] propose an alternative risk-aware framework for (unconstrained) reinforcement learning which places the risk measure in the occupancy measure to derive what the authors call the reward-based conditional value at risk (RCVaR) and mean-mean absolute deviation (Mean-RMAD), in contrast to the standard return-based formulations (like that of [12, 17, 43]).

Contributions This work approaches risk-aware constrained RL through employing reward-based risk measures in both the objective and the constraints. Our contributions are as follows:

- We extend the setting of Bonetti et al. [7] and, to the best of our knowledge, our work is the first to handle reward-based constraints, covering a large class of risk measures (OCEs);
- We establish a parameterized strong duality relation resulting in a computationally tractable partial Lagrangian relaxation (which is exact under certain constraint qualifications);
- We propose an online algorithm which is modular, offering the user flexibility in modeling (allowing for a mix of risk-neutral or risk-averse objectives and/or constraints) and implementation (enabling the use of existing RL algorithms as a black-box). By rigorously reducing the underlying problem to a certain instance of stochastic minimax optimization, we then establish convergence under common assumptions.
- We demonstrate practical usefulness of our approach in extensive numerical experiments on standard benchmarks, showcasing its effectiveness in reducing risk constraint violations and improving stochastic stability through explicit risk management.²

The setting of Bonetti et al. [7] is closely related to this work but is restricted to the unconstrained case, while our setting handles reward-based risk constraints as well. Our algorithm differs from the methodology of [7], which is based on a block-coordinate descent scheme, the analysis of which does not readily extend to the type of optimization problems arising in the constrained setting. Further, [7] requires exact policy solvers but we work with inexact ones (cf. Assumption 3.6). Compared to other risk-averse constrained RL methods, our partial Lagrangian relaxation is exact under a certain constraint qualification, whereas no meaningful (let alone exact) relation between the original primal problem and the employed Lagrangian relaxation has been shown in other works (e.g., [12]).

2 Background

Notation Given $m \in \mathbb{N}$, we let $[m] := \{1, \dots, m\}$. Given $x \in \mathbb{R}$, we write $(x)_+ \equiv \max\{x, 0\}$. We denote by $\mathcal{P}(\mathcal{S})$ the space of probability distributions over $\mathcal{S}$, equipped with its Borel σ -algebra.

²We provide the code at <https://github.com/baturaysaglam/risk-averse-constrained-RL>.

2.1 Risk Measures

In the context of stochastic optimization, it is well-known that expectations are unable to capture “risky” events (related to statistical variability, dispersion, or fat-tail behavior of the randomness). *Risk-averse optimization* aims to minimize the risk associated with such events, which is captured by certain functionals known as *risk measures*. For an extended discussion on risk-averse optimization, we refer the reader to Shapiro et al. [41, Chapter 6].

Optimized Certainty Equivalents (OCEs) Among different classes of risk measures, in this work we are interested in those which are in *infimal convolution* form, i.e., $\rho(Z) = \inf_t \{t + \mathbb{E}[g(Z - t)]\}$, assuming that Z represents losses. A notable example is the so-called *Conditional Value-at-Risk* at level β (denoted by $\text{CVaR}^\beta(Z)$), which is retrieved by setting $g(\cdot) = \frac{1}{\beta}(\cdot)_+$, for some $\beta \in (0, 1)$.

If Z represents rewards (as often in RL), we consider risk measures in the *supremal convolution* form, i.e., $\tilde{\rho}(Z) = \sup_t \{t + \mathbb{E}[g(Z - t)]\}$. For example, we can consider the reflected $-\text{CVaR}^\beta(-Z)$ (for reward maximization) by setting $g(u) = -\frac{1}{\beta}(-u)_+$. In this notation, g is called a *utility function*. Under certain conditions on the utility (e.g., satisfied by $g(u) = -(-u)_+$), these risk measures are known as *optimized certainty equivalents* (OCEs) (see Ben-Tal and Teboulle [4, Definition 2.2], where the reader can find many practical examples and properties). The class of OCEs is very rich (including many risk measures of practical value) and satisfies several important properties.

2.2 Risk-Neutral RL

Let $\tau \in \mathbb{N} \cup \{0\}$ denote the time instant and let $\mathcal{S} \subset \mathbb{R}^n$ and $\mathcal{A} \subset \mathbb{R}^d$ be compact sets describing the possible states and actions of an agent. Reinforcement learning is often modeled as a Markovian dynamical system where given a state and action pair at time τ , say (s_τ, a_τ) , the next state distribution only depends on the current state-action pair (and is independent of $(s_{<\tau}, a_{<\tau})$). The agent chooses actions at each time instant using a policy $\pi \in \mathcal{P}(\mathcal{S})$, and the action taken by the agent results in rewards defined by the uniformly bounded function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. We next describe the standard RL objective, that is, to maximize the (discounted, infinite) expected sum of rewards.

Problem (Primal RL Formulation).

$$V^* := \sup_{\pi \in \mathcal{P}(\mathcal{S})} \left\{ V(\pi) := \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau r(s_\tau^\pi, a_\tau^\pi) \right] \right\}, \quad (1)$$

where (s_τ^π, a_τ^π) denotes the state-action vector evolving under policy π . It is well-known that the above can be reformulated in terms of the (discounted) *occupancy (or occupation) measure* of a policy, defined as the discounted mixture $d\nu^\pi(\cdot, \bullet) = (1 - \gamma) \sum_{\tau=0}^{\infty} \gamma^\tau dp_\tau^\pi(\cdot, \bullet)$, where p_τ^π denotes the Borel probability measure induced by vector (s_τ^π, a_τ^π) . Let $\mathcal{R}$ denote the convex and weakly compact set of all occupation measures [8] induced by the policies $\pi \in \mathcal{P}(\mathcal{S})$. We can then rewrite the RL problem as follows.

Problem (Convex Analytic Dual Form³).

$$V^* := \sup_{\nu^\pi \in \mathcal{R}} \left\{ V(\nu^\pi) := \frac{1}{1 - \gamma} \cdot \mathbb{E}_{\nu^\pi}[r(s, a)] \right\}. \quad (2)$$

It is known that the occupation measure ν^π is in a one-to-one relationship with a policy π , i.e., if two policies have the same occupation measure, they must be the same (Sutton and Barto [42]).

2.3 Risk-Averse RL Formulations

There are several different frameworks for modeling “risk-awareness” and “safety” in RL. Existing works which explore risk-averse RL formulations have often taken the approach of replacing the

³This notion of casting dynamic optimization problems into abstract “static” optimization problems over a closed convex set of measures (as above) is referred to as the *convex analytic* approach by Borkar [9]. Thus, we will refer to these equivalent problems (1) and (2) as the “primal” and “convex analytic dual” formulations (not to be confused with duality in the Lagrangian sense) throughout.

Table 1: Comparison of the reward-based objective and the standard RL objective for general OCEs.

FORMULATION	RISK-AVERSE (REWARD-BASED)	RISK-NEUTRAL
PRIMAL	$\sup_{\pi, t} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} (t + g(r(s_{\tau}^{\pi}, a_{\tau}^{\pi}) - t)) \right]$	$\sup_{\pi} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} r(s_{\tau}^{\pi}, a_{\tau}^{\pi}) \right]$
DUAL	$\sup_{\nu^{\pi}, t} \frac{1}{1 - \gamma} \mathbb{E}_{\nu^{\pi}} [t + g(r(s, a) - t)]$	$\sup_{\nu^{\pi}} \frac{1}{1 - \gamma} \mathbb{E}_{\nu^{\pi}} [r(s, a)]$

expectations in the “primal” formulation (1) with risk measures, such as in the studies of Huang et al. [17], Tamar et al. [43], Chow et al. [12], among others:

$$\sup_{\pi \in \mathcal{P}(\mathcal{S})} -\rho \left[- \sum_{\tau=0}^{\infty} \gamma^{\tau} r(s_{\tau}^{\pi}, a_{\tau}^{\pi}) \right]. \quad (3)$$

Alternatively, one could generalize the tower property of expectations and iteratively evaluate a risk measure at each decision stage (e.g., see Ruszczyński [36]). A less common formulation –of interest herein– *applies a risk measure on the occupancy measure*, rather than the original probability space, as in Bonetti et al. [7]:

$$R^* = \sup_{\nu^{\pi} \in \mathcal{R}} \frac{1}{1 - \gamma} \cdot -\rho_{\nu^{\pi}(s, a)}(-r(s, a)), \quad (4)$$

where $\rho : \mathcal{L}_1(\nu^{\pi}, \mathbb{R}) \rightarrow \mathbb{R}$ is a finite-valued risk measure. This choice (4) enforces robustness over both space and time, capturing *per-stage risk* in contrast to (3); for details see Appendix A.1.

OCE Formulation By utilizing the relationship between (1) and (2) for risk measures induced by OCEs with a utility g , given that the reward is bounded, we obtain⁴

$$\frac{1}{1 - \gamma} \sup_{\nu^{\pi}} \sup_t \mathbb{E}_{\nu^{\pi}} [t + g(r(s, a) - t)] \iff \sup_t \sup_{\pi} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} (t + g(r(s_{\tau}^{\pi}, a_{\tau}^{\pi}) - t)) \right].$$

This relationship is summarized in Table 1. As one can see, the risk-averse problem involves a maximization over an additional variable t which depends on π (in this sequence). Importantly, given a fixed t , maximization over π resembles a problem of the form of (1) with a modified reward function. Thus one can take advantage of several existing algorithms which solve (1). This has been explored in the unconstrained setting of Bonetti et al. [7, Algorithm 2].

2.4 Constraints

Given the discussion in the previous section for risk-aware objectives, we can similarly describe constraints also of this form. It is common to address solving constrained problems through Lagrangian relaxation, which is amenable to numerical optimization, but (generally) only offers an approximate solution to the original constrained problem. We nonetheless show that the (partial) Lagrangian relaxation employed in this paper is *exact* under certain constraint qualifications.

Here, we highlight that Bonetti et al. [7] do *not* consider the constrained RL setting (which requires a dedicated analysis of the relation between the Lagrangian relaxation and the original constrained problem), while others using return-based risk objectives and constraints in the convex analytic “primal” space (e.g., Chow et al. [12]) are not able to show any meaningful relation between the primal problem and their employed Lagrangian relaxation.

3 Main Results (Risk Constrained Learning)

The framework we propose has favorable robustness properties, better capturing per-stage risk (see Appendix A.1) for applications which are also time-sensitive. Now, we: (1) show that this formulation

⁴Going forward, the expectation or risk without subscripts is taken over the initial state and sample path distribution induced by the state transitions of the system and the policy π .

allows us to solve the original constrained problem through a partial Lagrangian relaxation which is practically implementable and (2) reduce the problem to an instance of stochastic minimax optimization which lends itself to non-asymptotic convergence analysis.

3.1 Problem Description and Equivalent Reformulations

To avoid generalities and to facilitate exposition, hereafter we exclusively consider the CVaR^β . However, all results presented hold for general OCEs. We also use the same β for all constraints and the objective, but these need not be the same. We first write a constrained version of the problem (4):

$$\sup_{\pi \in \mathcal{P}} \frac{1}{1-\gamma} \cdot -\text{CVaR}_{\nu^\pi}^\beta(-r_0(s, a)) \quad \text{s.t.} \quad \frac{1}{1-\gamma} \cdot -\text{CVaR}_{\nu^\pi}^\beta(-r_i(s, a)) \geq c_i, \quad \forall i \in [m], \quad (5)$$

or, equivalently, as:

$$\sup_{\pi \in \mathcal{P}(\mathcal{S})} \sup_{t_0 \in \mathbb{R}} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau (r'_0(s_\tau^\pi, a_\tau^\pi, t_0)) \right] \quad \text{s.t.} \quad \sup_{t_i \in \mathbb{R}} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau (r'_i(s_\tau^\pi, a_\tau^\pi, t_i)) \right] \geq c_i, \quad \forall i \in [m], \quad (6)$$

where $r'_i(s, a, t) := t - \frac{1}{\beta}(t - r_i(s, a))_+$, and $r_i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ are the reward functions for $i \in \{0\} \cup [m]$. Note that r'_i is written specifically for the CVaR^β here, but for general OCEs we simply replace $-\frac{1}{\beta}(\cdot)_+$ with $g(\cdot)$. Recall that, by standard convention, r_i is bounded for all i .

Remark 3.1. Consider any OCE, and assume that r_i are bounded for all $i \in \{0\} \cup [m]$. Then, by Ben-Tal and Teboulle [4, Proposition 2.1], we have that

$$\sup_{t \in \mathbb{R}} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau (t + g(r_i(s_\tau^\pi, a_\tau^\pi) - t)) \right] = \sup_{t \in \mathcal{T}_i} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau (t + g(r_i(s_\tau^\pi, a_\tau^\pi) - t)) \right],$$

where $\mathcal{T}_i$ is a bounded interval containing the uniformly smallest and largest values of the reward $r_i(s_\tau, a_\tau)$. Thus, we can (equivalently) cast problem (6), by restricting $t := [t_0, t_1, \dots, t_m] \in \mathcal{T}$, with $\mathcal{T} \subset \mathbb{R}^{m+1}$ some convex and compact set. This will be utilized throughout.

Lemma 3.2. *Problem (6) is equivalent to*

$$\sup_{\substack{\pi \in \mathcal{P}(\mathcal{S}) \\ t \in \mathcal{T}}} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau (r'_0(s_\tau^\pi, a_\tau^\pi, t_0)) \right] \quad \text{s.t.} \quad \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau (r'_i(s_\tau^\pi, a_\tau^\pi, t_i)) \right] \geq c_i, \quad \forall i \in [m]. \quad (7)$$

The proof is based on Chow et al. [12], and given in Appendix A.2. Although it may seem somewhat redundant, (7) is a maximization problem jointly over the variables π, t , subject to functional inequality constraints. Thus, the Lagrangian associated to (7) reads $\mathcal{L}(\pi, t, \lambda) := \mathbb{E}[\hat{\mathcal{L}}(\pi, t, \lambda)]$, with

$$\hat{\mathcal{L}}(\pi, t, \lambda) := \sum_{\tau=0}^{\infty} \gamma^\tau \left[(r'_0(s_\tau^\pi, a_\tau^\pi, t_0)) - \sum_{i=1}^m \lambda_i c_i + \sum_{i=1}^m \lambda_i (r'_i(s_\tau^\pi, a_\tau^\pi, t_i)) \right]. \quad (8)$$

Then, the primal problem is $P^* = \sup_{\pi, t} \inf_{\lambda \geq 0} \mathcal{L}(\pi, t, \lambda)$. Note that the above, for fixed t, λ , is the Lagrangian for a risk-neutral constrained problem (similar to those considered by Paternain et al. [33]), parameterized by t , which controls risk-aversion (in harmony with the choice of β).

3.2 Partial Lagrangian Relaxation

Next, we derive a partial Lagrangian relaxation for (7) (which is exact under a constraint qualification; cf. Assumption 3.4). In turn, this allows us to cast the problem in a stochastic minimax optimization framework and derive an effective online algorithm for its solutions (cf. Section 3.3).

Partial Maximization Problem Let us now consider problem (7) for fixed $t \in \mathcal{T}$, i.e.,

$$\sup_{\pi \in \mathcal{P}(\mathcal{S})} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau (r'_0(s_\tau^\pi, a_\tau^\pi, t_0)) \right] \quad \text{s.t.} \quad \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau (r'_i(s_\tau^\pi, a_\tau^\pi, t_i)) \right] \geq c_i, \quad \forall i \in [m]. \quad (9)$$

This problem exhibits strong duality under usual conditions.

Proposition 3.3. *Let r_i be bounded functions for all $i \in \{0\} \cup [m]$. Assume that Slater’s condition⁵ holds for (9). Then, (9) exhibits strong duality, and thus*

$$\sup_{\pi} \inf_{\lambda \geq 0} \mathcal{L}(\pi, t, \lambda) = \inf_{\lambda \geq 0} \sup_{\pi} \mathcal{L}(\pi, t, \lambda).$$

The proof of Proposition 3.3 is given in Appendix A.3.

Maximizing over t From Remark 3.1, we know that $t \in \mathcal{T}$ and $\mathcal{T}$ is a convex and compact set (since we assume bounded rewards). Note that Proposition 3.3 holds for each fixed $t \in \mathcal{T}$, assuming Slater’s condition is satisfied at this point. To proceed, it suffices to introduce another constraint qualification, as follows.

Assumption 3.4 (Constraint Qualification). *Let t^* be an argument which attains the maximum over the variable $t = [t_0, t_1, \dots, t_m]$ in (7). There is a (non-singleton) convex and compact set $\mathcal{I} \subset \mathcal{T}$, with $t^* \in \mathcal{I}$, such that Slater’s condition holds for (9), for every $t \in \mathcal{I}$.*

Under Assumption 3.4, we have that

$$\sup_{t \in \mathcal{I}} \sup_{\pi} \inf_{\lambda \geq 0} \mathcal{L}(\pi, t, \lambda) = \sup_{t \in \mathcal{I}} \inf_{\lambda \geq 0} \sup_{\pi} \mathcal{L}(\pi, t, \lambda).$$

Parametrized Policies and Almost-zero Duality Gap In order to solve the (inner) policy optimization problem, we parametrize the policy by a vector $\theta \in \Theta \subset \mathbb{R}^p$, representing the coefficients of, say, a neural network, assuming that Θ is a compact set (which is the case in practice). We have shown that under Assumption 3.4, the proposed partial Lagrangian relaxation is exact. We argue that this remains almost true (in the sense described below) even if we utilize parametrized policies. In what follows, we assume that π_{θ} is an ϵ -universal parametrization of measures in $\mathcal{P}(\mathcal{S})$, according to Paternain et al. [33, Definition 1].

Theorem 3.5 (Almost-zero Duality Gap). *Let Assumption 3.4 hold, assume that π_{θ} is ϵ -universal for measures in $\mathcal{P}(\mathcal{S})$, and that the policy-parametrized version of (7) is feasible. Then,*

$$P^*(t^*) := \sup_{t \in \mathcal{I}} \sup_{\pi} \inf_{\lambda \geq 0} \mathcal{L}(\pi, t, \lambda) \geq \sup_{t \in \mathcal{I}} \inf_{\lambda \geq 0} \sup_{\theta} \mathcal{L}(\pi_{\theta}, t, \lambda) \geq P^*(t^*) - \mathcal{O}\left(\frac{\epsilon}{1-\gamma}\right).$$

The proof is given in Appendix A.4. Theorem 3.5 states that solving the parameterized (partial) dual problem incurs negligible cost (depending on the universal parametrization error $\epsilon > 0$) under mild assumptions and our usual constraint qualification (see Assumption 3.4).

The Practical Parametrized Partial Dual Problem In practice, we do not have access to the set $\mathcal{I}$ (assuming that Assumption 3.4 is satisfied), and thus, we may heuristically search over all $t \in \mathcal{T}$, noting that a reasonable initial guess can maximize our chances of landing in $\mathcal{I}$ (and thus solving the original constrained problem (7) (almost) exactly). Nonetheless, the proposed algorithmic framework does not rely on Assumption 3.4. Indeed, in the absence of this, we recover a partial Lagrangian relaxation which yields an approximate solution to problem (7). However, we highlight that, unlike Lagrangian relaxations considered in the literature, our approach is exact under Assumption 3.4.

Additionally, from Proposition 3.3, we know that the optimal Lagrange multiplier is attained, assuming that Slater’s condition holds for (7). Thus, there is an optimal Lagrange multiplier associated to (7) within some sufficiently large convex and compact set $\Lambda \subset \mathbb{R}_+^m$. In practice, the size of this set can be adjusted dynamically, if necessary. Hence, in what follows, we restrict the minimization over $\lambda \in \Lambda$. In light of the previous discussion, we focus on the following parametrized partial dual formulation:

$$D_{\theta}^* := \sup_{t \in \mathcal{T}} \underbrace{\inf_{\lambda \in \Lambda} \sup_{\theta} \mathcal{L}(\pi_{\theta}, t, \lambda)}_{\text{blackbox RL}}. \quad (10)$$

A Wrapper for Blackbox Unconstrained RL Algorithms Notice that $\mathcal{L}(\pi_{\theta}, t, \lambda)$ is a linear combination of the expected discounted sum of adjusted reward functions which take the original r_i to $r'_i(s, a, t_i) = t_i - \frac{1}{\beta}(t_i - r_i(s, a))_+$ for $i = 0, 1, \dots, m$. Thus, the innermost problem in

⁵Slater’s condition requires that there exists a strictly feasible policy.

Algorithm 1 Reward-Based SGDA with Risk Constraints

- 1: **Input:** Bounded reward functions $r_i : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$, discount factor $\gamma \in (0, 1)$, step sizes $\eta_\lambda, \eta_t > 0$, batch size n .
 - 2: **Initialize:** $t^{(0)}, \lambda^{(0)}$
 - 3: **for** each iteration $j = 1, 2, \dots, J$ **do**
 - 4: Call the inexact oracle to obtain $\pi_{\theta^{(j)}}$ satisfying Assumption 3.6.
 - 5: Collect a batch of trajectories $\{x_k\}_{k=1}^n$ using $\pi_{\theta^{(j)}}$ and compute sample (sub)gradients:
$$\widehat{\nabla}_\lambda \widehat{\mathcal{L}}(\pi_{\theta^{(j)}}, t^{(j)}, \lambda^{(j)}), \quad \widehat{\nabla}_t \widehat{\mathcal{L}}(\pi_{\theta^{(j)}}, t^{(j)}, \lambda^{(j)}).$$
 - 6: Update dual (λ) and auxiliary (t) variables:
$$\lambda^{(j+1)} \leftarrow \Pi_\Lambda \left(\lambda^{(j)} - \eta_\lambda \widehat{\nabla}_\lambda \widehat{\mathcal{L}}(\pi_{\theta^{(j)}}, t^{(j)}, \lambda^{(j)}) \right),$$
$$t^{(j+1)} \leftarrow \Pi_{\mathcal{T}} \left(t^{(j)} + \eta_t \widehat{\nabla}_t \widehat{\mathcal{L}}(\pi_{\theta^{(j)}}, t^{(j)}, \lambda^{(j)}) \right).$$
 - 7: **end for**
 - 8: **Return:** $\pi_{\theta^{(j^*)}}$, where we sample $j^* \sim \text{Unif}([J])$.
-

(10) reduces to an unconstrained RL problem with a linear combination of terms defined by reward functions r'_i , given λ and t . Hence, we have a practical algorithm for solving (10) which, if solved exactly for a starting point $t \in \mathcal{I}$, is arbitrarily close to the true primal solution P^* of (6), under Assumption 3.4. In practice, this problem is non-convex so we may only hope to find a locally optimal solution. Convergence to such a solution is established next.

3.3 Algorithm and Analysis

Next, we present the algorithmic procedure for solving problem (10), in Algorithm 1. Consider a policy oracle that precisely solves the problem $\sup_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$, such that for every $(t, \lambda) \in \mathcal{T} \times \Lambda$, it returns a policy $\pi_{\theta^*(t, \lambda)} \in \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$. Using such an oracle, problem (10) reads

$$D_\theta^* = \sup_{t \in \mathcal{T}} \inf_{\lambda \in \Lambda} \underbrace{\mathcal{L}(\pi_{\theta^*(t, \lambda)}, t, \lambda)}_{:= -f(t, \lambda)} = - \inf_{t \in \mathcal{T}} \sup_{\lambda \in \Lambda} f(t, \lambda). \quad (11)$$

Having access to a solver which returns $\pi_{\theta^*(t, \lambda)}$ exactly is difficult. Instead, we assume the availability of an inexact oracle, as defined next. We justify the generality of this assumption in Appendix A.5 to save space in the main body.

Assumption 3.6 (Local Solutions). For any $(t, \lambda) \in \mathcal{T} \times \Lambda$, let $\theta^*(t, \lambda)$ be a maximum of $\mathcal{L}(\pi_\theta, t, \lambda)$ and $\theta^\dagger(t, \lambda)$ be returned by a generic RL algorithm. Then, there exists a $\delta > 0$ such that $\forall (t, \lambda) \in \mathcal{T} \times \Lambda$, $\widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^\dagger(t, \lambda)}, t, \lambda) = \widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^*(t, \lambda)}, t, \lambda) + b(\theta^*, \theta^\dagger, t, \lambda)$, with $\|b(\theta^*, \theta^\dagger, t, \lambda)\| \leq \delta$.

In what follows, we show that (10) can be reduced to a stochastic minimax optimization problem satisfying appropriate assumptions, and thus we can obtain a non-asymptotic convergence rate for Algorithm 1. For notational convenience, we define the function $\Phi(t) := \sup_{\lambda \in \Lambda} f(t, \lambda)$. First, we show that f is Lipschitz with respect to $t \in \mathbb{R}^{m+1}$ uniformly over $\lambda \in \Lambda$.

Lemma 3.7. *There exists a constant $C > 0$ such that*

$$|\sup_\pi \mathcal{L}(\pi, t_1, \lambda) - \sup_\pi \mathcal{L}(\pi, t_2, \lambda)| \leq C \|t_1 - t_2\|_2$$

for all pairs $t_1, t_2 \in \mathbb{R}^{m+1}$ and $\lambda \in \Lambda$.

We prove Lemma 3.7 (which readily applies to parametrized policies as well) in Appendix A.6. Following Lin et al. [25], we analyze our method under the assumption of Lipschitz smoothness of the function $\mathcal{L}(\pi_{\theta^*(t, \lambda)}, t, \lambda)$, for $\theta^*(t, \lambda) \in \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$, which we state below.

Assumption 3.8 (Lipschitz Smoothness). The function $f(t, \lambda) \equiv -\mathcal{L}(\pi_{\theta^*(t, \lambda)}, t, \lambda)$ is ℓ -smooth over $\mathcal{T} \times \Lambda$, where $\theta^*(t, \lambda) \in \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$ is an arbitrary selection.

Remark 3.9. We note that Assumption 3.8 is not particularly restrictive in our setting, and it holds under several general conditions, as outlined in Appendix A.7. Nonetheless, if not satisfied, we can still enforce this assumption by making a slight algorithmic adjustment involving the appropriate addition of small uniform noise. For a comprehensive discussion we again refer the reader to Appendix A.7, where we also present an explicit formula for computing the (exact) sample gradients.

As appears in prior work [25], we make the standard (and, in our case, mild) assumption that the exact stochastic (sub)gradient oracles are unbiased and have at most σ^2 variance.

Assumption 3.10 ((Sub)gradient oracles). For all $(t, \lambda) \in \mathcal{T} \times \Lambda$, we assume that

$$\mathbb{E}[\widehat{\nabla}_{t,\lambda}\widehat{\mathcal{L}}(\pi^*(t, \lambda), t, \lambda) + \nabla f(\lambda, t)] = 0, \quad \mathbb{E}[\|\widehat{\nabla}_{t,\lambda}\widehat{\mathcal{L}}(\pi^*(t, \lambda), t, \lambda) + \nabla f(\lambda, t)\|_2^2] \leq \sigma^2.$$

Given the imposed assumptions, we now present the convergence guarantees of Algorithm 1. In the context of nonsmooth weakly convex optimization, the Moreau envelope serves as a measure of closeness to stationarity. With this in mind, we introduce the definition of an ϵ -stationary point.

Definition 3.11 (ϵ -stationary point). A point x is ϵ -stationary if $\|\nabla\Phi_{1/2\ell}(x)\| \leq \epsilon$, where, given some $\delta > 0$, we define the Moreau envelope as $\Phi_\delta(x) := \inf_{w \in \mathbb{R}^{m+1}} \{\Phi(w) + \frac{1}{2\delta}\|w - x\|^2\}$.

When $\epsilon = 0$, a stationary point of the Moreau envelope $\Phi_\delta(\cdot)$ is also stationary for $\Phi(\cdot)$. For $\epsilon > 0$, one can argue that an ϵ -stationary point for $\Phi_{1/2\ell}$ is close to an ϵ -stationary point of $\Phi(\cdot)$, as long as $\Phi(\cdot)$ is ℓ -weakly convex (noting that, under our assumptions, $\Phi(\cdot)$ is ℓ -weakly convex). We proceed with a bound on the iteration complexity of Algorithm 1.

Theorem 3.12. *Let the step-sizes η_t and η_λ in Algorithm 1 be small enough as in Eq. (30) and (29), respectively, and batch size $n = 1$. Define the quantities $\widehat{\Delta}_\Phi := \Phi_{1/2\ell}(t^{(0)}) - \min_t \Phi_{1/2\ell}(t)$ and $\widehat{\Delta}_0 := \Phi(t^{(0)}) - f(\lambda^{(0)}, t^{(0)})$. If Assumptions 3.6, 3.8 and 3.10 hold, the iteration complexity for recovering an $\mathcal{O}(\sqrt{\epsilon^2 + \delta\ell(\text{diam}(\mathcal{T}) + \text{diam}(\Lambda))})$ -stationary point (Definition 3.11) is of order*

$$\mathcal{O}\left(\left(\frac{\ell^3(C^2 + \sigma^2 + \delta^2)(\text{diam}(\Lambda))^2 \cdot \widehat{\Delta}_\Phi}{\epsilon^6} + \frac{\ell^3(\text{diam}(\Lambda))^2 \cdot \widehat{\Delta}_0}{\epsilon^4}\right) \max\left\{1, \frac{\sigma^2 + \delta^2}{\epsilon^2}\right\}\right).$$

Note that this result says we only need single trajectories ($n = 1$). Thus our algorithm can be run *online*. The proof, which is detailed in Appendix A.8, extends the developments of [25], by incorporating biased (sub)gradient samples, occurring due to the use of inexact oracles. If $\delta = \mathcal{O}(\epsilon^2)$, then we recover an ϵ -stationary point and the complexity of [25].

4 Experiments

We conduct experiments on locomotion tasks from the Safety-Gymnasium benchmark [19].

4.1 Setup

We consider two constraints in the environment: safe navigation and safe velocity. The first involves simple problems with low-dimensional state and action spaces, where the agent must navigate under rules (e.g., along a path or without hitting obstacles). The second is more challenging, involving standard high-dimensional control tasks based on MuJoCo-v4 agents [46]. At each timestep, the agent receives a cost of +1 if it violates a constraint and 0 otherwise. These violations do not affect environment dynamics—they neither terminate episodes nor alter transitions. Importantly, our experiments do not access this cost directly during training; instead, the agent manages risk by being risk-aware in its velocities.

To simulate risk and uncertainty in an otherwise deterministic setting (where identical actions yield identical outcomes), we inject zero-mean Gaussian noise with a standard deviation of 0.05 (5% of the action range) into all agent actions during both training and evaluation. This adds stochasticity to the agent’s actions and introduces risk.

Experimental Goals and Motivation Our objective is to evaluate whether the proposed risk-sensitive method ensures stable convergence of the dual variable (λ) and the auxiliary variable (t), while also inducing safe and meaningful behavioral changes to the agent. To emphasize convergence, we train agents for extended durations (5–18M steps) to provide a thorough *proof-of-work*. Specifically, we aim to: (1) verify that t converges to the empirical CVaR $^\beta$ of the post-training constraint violation distribution (e.g., safe velocity violations), (2) confirm that λ stabilizes to enforce the constraint, and (3) assess whether the policy maximizes reward while maintaining safe velocities.

Table 2: Cumulative episodic evaluation costs and rewards of the converged agents, computed as the mean of the last 10 episodes. The PPO baseline is unconstrained and risk-neutral. We use the simplest Safety-Gymnasium agent, Point at level 1 difficulty, to avoid additional challenges and isolate the effects of the constraints.

Environment	# Training Steps	Cost ↓		Reward ↑	
		PPO	Ours	PPO	Ours
Button	5M	150.76	0.0	24.29	2.58
Circle	5M	206.74	0.0	60.18	39.19
Goal	10M	45.09	0.0	21.89	13.56
Push	5M	38.48	0.0	0.93	2.42

Optimization with CVaR Constraints Let $v : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ denote the constraint-quantifying function (e.g., velocity). We formulate the following constrained optimization problem:

$$\sup_{\pi \in \mathcal{P}(\mathcal{S})} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} r(s_{\tau}, \pi(s_{\tau})) \right] \quad \text{s.t.} \quad \text{CVaR}_{\nu^{\pi}}^{\beta}(v(s, a)) \leq c,$$

where c is a user-defined threshold and β is the CVaR parameter. Our framework is general and allows risk-neutral objectives to be combined with risk-aware constraints. To solve this, we transform the reward at each timestep using the CVaR-based Lagrangian (see Appendix B.3.1 for the derivation):

$$r(s_{\tau}, \pi(s_{\tau})) + \lambda_i \left(c + t - \frac{1}{\beta} (t + v(s_{\tau}, \pi(s_{\tau})))_+ \right).$$

In safe navigation, the constraint variable reduces to $v(s, a) \in \{0, 1\}$, since navigation is not quantifiable in the same way as velocity. This makes the optimization problem more difficult, as the penalty term (the λ multiplier) becomes sparse.

Proximal Policy Optimization (PPO) [39] is used as a black-box solver to optimize the policy π , implementing line 4 of Algorithm 1. The dual variable λ and auxiliary t are updated via stochastic gradient steps following the rest of Algorithm 1. For further experimental details, see Appendix B.

4.2 Results

Results are shown in Table 2 for safe navigation and in Figure 1 for safe velocity.

Constraint Handling The performance in the navigation tasks shows us that our method can strictly prevent any constraint violations, whether quantified continuously or not. In fact, it is the only method achieving strictly zero violations compared to other PPO-based constraint learning algorithms (see Table 5 in [19]). In some cases, constraint satisfaction also leads to higher rewards than the unconstrained vanilla agent (e.g., in the Push environment). This shows that constraints are not arbitrary but can guide the agent toward optimal task behavior.

Risk-Management On the other hand, the velocity tasks provide insights into the risk management capabilities. With properly tuned step sizes, the dual and CVaR variables stabilize and oscillate around consistent values. The post-training evaluations (the last column in Figure 1) extract the distribution of the trained agent’s velocity over several evaluation episodes. We observe that the converged t -value matches the β -upper quantile of the velocity distribution. This is what the CVaR aims to do: it captures the expected cost in the worst β fraction of outcomes.

Further, the evaluation reward becomes increasingly stable as training progresses despite the uncertainty in the evaluation environment. Considering the number of data points in the curves—1000 per 1M steps (e.g., about 14,000 points in Hopper)—the robustness of our method in managing risk is evident. This is also confirmed in post-training simulations: the agent moves forward cautiously compared to the unconstrained (risk-neutral) PPO agent, demonstrating that the learned policy adheres to the task while remaining interpretable from the perspective of the constraint.

We recognize the extended training duration and the need to tune the step sizes of λ and t . Nonetheless, once these variables converge, the algorithm is effective not only in toy tasks but also in realistic control settings, making patience during training an important practical consideration.




Figure 1: Learning curves of the agent trained to convergence. We report episodic evaluation rewards, the progression of optimization variables t and λ , and the velocity distribution of the trained agent evaluated over 100 post-training episodes without further learning. Curves are *not* smoothed.

5 Conclusion

We introduce a training framework for solving constrained reinforcement learning (RL) problems with risk-awareness which has desirable properties. The problem exhibits a parameterized strong duality under constraint qualifications, which lends itself to an algorithm that is theoretically-supported and flexible to implement; the inner problem can be solved using any black-box RL algorithm, while the other variables are updated using SGDA. The framework handles general OCEs beyond the CVaR, and captures both risk-neutral (e.g., by setting $\beta = 1$ for the CVaR^β) and risk-aware objectives and constraints, or a combination of the two, with empirical evidence that, in the context of risk-aware constrained RL, it effectively manages risk.

Limitations An open problem resulting from this work asks whether or not full strong duality holds unconditionally. Also, while our framework is inherently risk-aware, it is more computationally intensive compared to risk-neutral methods (as each update of λ and t variables requires solving for an approximately optimal policy), as is also true for other risk-aware methods, e.g., [7, 12].

References

- [1] Mohammed Amin Abdullah, Hang Ren, Haitham Bou Ammar, Vladimir Milenkovic, Rui Luo, Mingtian Zhang, and Jun Wang. Wasserstein robust reinforcement learning. *arXiv preprint arXiv:1907.13196*, 2019.
- [2] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. *CoRR*, abs/1705.10528, 2017. URL <http://arxiv.org/abs/1705.10528>.
- [3] Amrit Singh Bedi, Alec Koppel, Ketan Rajawat, and Panchajanya Sanyal. Nonparametric compositional stochastic optimization for risk-sensitive kernel learning, 2020.
- [4] Aharon Ben-Tal and Marc Teboulle. An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance*, 17(3):449–476, 2007. doi: <https://doi.org/10.1111/j.1467-9965.2007.00311.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9965.2007.00311.x>.
- [5] Shalabh Bhatnagar. An online actor-critic algorithm with function approximation for constrained markov decision processes. *Journal of Optimization Theory and Applications*, 153, 06 2012. doi: 10.1007/s10957-012-9989-5.
- [6] Lorenzo Bisi, Luca Sabbioni, Edoardo Vittori, Matteo Papini, and Marcello Restelli. Risk-averse trust region optimization for reward-volatility reduction. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4583–4589. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi: 10.24963/ijcai.2020/632. URL <https://doi.org/10.24963/ijcai.2020/632>. Special Track on AI in FinTech.
- [7] Massimiliano Bonetti, Lorenzo Bisi, and Marcello Restelli. Risk-averse optimization of reward-based coherent risk measures. *Artificial Intelligence*, page 103845, 2023. ISSN 0004-3702. doi: <https://doi.org/10.1016/j.artint.2022.103845>. URL <https://www.sciencedirect.com/science/article/pii/S0004370222001850>.
- [8] Vivek S. Borkar. A convex analytic approach to markov decision processes. *Probability Theory and Related Fields*, 78:583–602, 1988.
- [9] Vivek S. Borkar. *Convex Analytic Methods in Markov Decision Processes*, pages 347–375. Springer US, Boston, MA, 2002. ISBN 978-1-4615-0805-2. doi: 10.1007/978-1-4615-0805-2_11. URL https://doi.org/10.1007/978-1-4615-0805-2_11.
- [10] V.s Borkar. An actor-critic algorithm for constrained markov decision processes. *Systems & Control Letters*, 54:207–213, 03 2005. doi: 10.1016/j.sysconle.2004.08.007.
- [11] Adrian Rivera Cardoso and Huan Xu. Risk-averse stochastic convex bandit. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 39–47. PMLR, 16–18 Apr 2019. URL <https://proceedings.mlr.press/v89/cardoso19a.html>.
- [12] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *Journal of Machine Learning Research*, 18, 12 2015.
- [13] Yingjie Fei, Zhuoran Yang, Yudong Chen, Zhaoran Wang, and Qiaomin Xie. Risk-sensitive reinforcement learning: Near-optimal risk-sample tradeoff in regret. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 22384–22395. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/fdc42b6b0ee16a2f866281508ef56730-Paper.pdf>.
- [14] Peter Geibel and Fritz Wysotzki. Risk-sensitive reinforcement learning applied to control under constraints. *Journal of Artificial Intelligence Research*, 24:81–108, 2005.

- [15] Jun-ya Gotoh and Stan Uryasev. Support vector machines based on convex risk functions and general norms. *Annals of Operations Research*, 249, 02 2017. doi: 10.1007/s10479-016-2326-x.
- [16] Hassaan Hashmi, Spyridon Pougkakiotis, and Dionysis Kalogieras. Model-free learning of two-stage beamformers for passive irs-aided network design. *IEEE Transactions on Signal Processing*, 72:652–669, 2024. doi: 10.1109/TSP.2023.3346182.
- [17] Audrey Huang, Liu Leqi, Zachary C Lipton, and Kamyar Azizzadenesheli. On the convergence and optimality of policy gradient for markov coherent risk. *arXiv preprint arXiv:2103.02827*, 2021.
- [18] Wenjie Huang and William B. Haskell. Risk-aware q-learning for markov decision processes. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 4928–4933, 2017. doi: 10.1109/CDC.2017.8264388.
- [19] Jiaming Ji, Borong Zhang, Jiayi Zhou, Xuehai Pan, Weidong Huang, Ruiyang Sun, Yiran Geng, Yifan Zhong, Josef Dai, and Yaodong Yang. Safety gymnasium: A unified safe reinforcement learning benchmark. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=WZmlxIuIGR>.
- [20] Daniel R. Jiang and Warren B. Powell. Risk-averse approximate dynamic programming with quantile-based risk measures. *Mathematics of Operations Research*, 43(2):554 – 579, 2018. doi: 10.1287/moor.2017.0872.
- [21] Dionysios Kalogieras and Spyridon Pougkakiotis. Strong duality in risk-constrained nonconvex functional programming, 2023. URL <https://arxiv.org/abs/2206.11948v3>.
- [22] Drew P. Kouri and Thomas M. Surowiec. Epi-regularization of risk measures. *Mathematics of Operations Research*, 45(2):774–795, 2020. doi: 10.1287/moor.2019.1013. URL <https://doi.org/10.1287/moor.2019.1013>.
- [23] Thanh Lam, Arun Verma, Bryan Kian Hsiang Low, and Patrick Jaillet. Risk-aware reinforcement learning with coherent risk measures and non-linear function approximation. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=-RwZ0Vybbyj>.
- [24] Jan Leike, Miljan Martic, Viktoriya Krakovna, Pedro Ortega, Tom Everitt, Andrew Lefrancq, Laurent Orseau, and Shane Legg. Ai safety gridworlds. 11 2017.
- [25] Tianyi Lin, Chi Jin, and Michael Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6083–6093. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/lin20a.html>.
- [26] Avinash N. Madavan and Subhonmesh Bose. A stochastic primal-dual method for optimization with conditional value at risk constraints. *Journal of Optimization Theory and Applications*, 190(2):428–460, August 2021. ISSN 0022-3239. doi: 10.1007/s10957-021-01888-x.
- [27] Horia Mania, Aurelia Guy, and Benjamin Recht. Simple random search of static linear policies is competitive for reinforcement learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/7634ea65a4e6d9041cfd3f7de18e334a-Paper.pdf>.
- [28] Shie Mannor and Nahum Shimkin. A geometric approach to multi-criterion reinforcement learning. *J. Mach. Learn. Res.*, 5:325–360, dec 2004. ISSN 1532-4435.
- [29] Oliver Mihatsch and Ralph Neuneier. Risk-sensitive reinforcement learning. *Machine learning*, 49(2):267–290, 2002.

- [30] Sobhan Miryoosefi and Chi Jin. A simple reward-free approach to constrained reinforcement learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 15666–15698. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/miryoosefi22a.html>.
- [31] Somayeh Moazeni, Warren B. Powell, Boris Defourny, and Belgacem Bouzaïene-Ayari. Parallel nonstationary direct policy search for risk-averse stochastic optimization. *INFORMS Journal on Computing*, 29(2):332–349, 2017. doi: 10.1287/ijoc.2016.0733.
- [32] Matthew Norton, Alexander Mafusalov, and Stan Uryasev. Soft margin support vector classification as buffered probability minimization. *Journal of Machine Learning Research*, 18(68): 1–43, 2017. URL <http://jmlr.org/papers/v18/15-566.html>.
- [33] Santiago Paternain, Luiz Chamon, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained reinforcement learning has zero duality gap. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/c1aeb6517a1c7f33514f7ff69047e74e-Paper.pdf>.
- [34] Spyridon Pougkakiotis and Dionysis Kalogerias. A zeroth-order proximal stochastic gradient method for weakly convex stochastic optimization. *SIAM Journal on Scientific Computing*, 45(5):A2679–A2702, 2023. doi: 10.1137/22M1494270. URL <https://doi.org/10.1137/22M1494270>.
- [35] A. PrashanthL. and Michael C. Fu. Risk-sensitive reinforcement learning: A constrained optimization viewpoint. *ArXiv*, abs/1810.09126, 2018.
- [36] Andrzej Ruszczyński. Risk-averse dynamic programming for markov decision processes. *Mathematical programming*, 125:235–261, 2010.
- [37] Amir Sani, Alessandro Lazaric, and Rémi Munos. Risk-aversion in multi-armed bandits. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL https://proceedings.neurips.cc/paper_files/paper/2012/file/83f2550373f2f19492aa30fbd5b57512-Paper.pdf.
- [38] John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1506.02438>.
- [39] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL <https://arxiv.org/abs/1707.06347>.
- [40] Alexander Shapiro. Second-order derivatives of extremal-value functions and optimality conditions for semi-infinite programs. *Mathematics of Operations Research*, 10(2):207–219, 1985. ISSN 0364765X, 15265471. URL <http://www.jstor.org/stable/3689374>.
- [41] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory, Third Edition*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2021. doi: 10.1137/1.9781611976595. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611976595>.
- [42] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018. URL <http://incompleteideas.net/book/the-book-2nd.html>.
- [43] Aviv Tamar, Yinlam Chow, Mohammad Ghavamzadeh, and Shie Mannor. Policy gradient for coherent risk measures. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL <https://proceedings.neurips.cc/paper/2015/file/024d7f84fff11dd7e8d9c510137a2381-Paper.pdf>.

- [44] Aviv Tamar, Yinlam Chow, Mohammad Ghavamzadeh, and Shie Mannor. Sequential decision making with coherent risk. *IEEE Transactions on Automatic Control*, PP:1–1, 12 2016. doi: 10.1109/TAC.2016.2644871.
- [45] Chen Tessler, Daniel J. Mankowitz, and Shie Mannor. Reward constrained policy optimization. *CoRR*, abs/1805.11074, 2018. URL <http://arxiv.org/abs/1805.11074>.
- [46] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.
- [47] Mark Towers, Ariel Kwiatkowski, Jordan Terry, John U. Balis, Gianluca De Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Markus Krimmel, Arjun KG, Rodrigo Perez-Vicente, Andrea Pierré, Sander Schulhoff, Jun Jet Tai, Hannah Tan, and Omar G. Younis. Gymnasium: A standard interface for reinforcement learning environments, 2024. URL <https://arxiv.org/abs/2407.17032>.
- [48] Constantine Alexander Vitt, D. Dentcheva, and Hui Xiong. Risk-averse classification. *ArXiv*, abs/1805.00119, 2018.
- [49] Kaiwen Wang, Nathan Kallus, and Wen Sun. Near-minimax-optimal risk-sensitive reinforcement learning with cvar. In *Proceedings of the 40th International Conference on Machine Learning*, ICML’23. JMLR.org, 2023.
- [50] Kaiwen Wang, Dawen Liang, Nathan Kallus, and Wen Sun. A reductions approach to risk-sensitive reinforcement learning with optimized certainty equivalents, 2025.
- [51] Wenhao Xu, Xuefeng Gao, and Xuedong He. Regret bounds for markov decision processes with recursive optimized certainty equivalents. In *Proceedings of the 40th International Conference on Machine Learning*, ICML’23. JMLR.org, 2023.
- [52] Chen Yu and Andre Rosendo. Risk-aware model-based control. *Frontiers in Robotics and AI*, 8, 2021. ISSN 2296-9144. doi: 10.3389/frobt.2021.617839. URL <https://www.frontiersin.org/articles/10.3389/frobt.2021.617839>.
- [53] Pengqian Yu, William B. Haskell, and Huan Xu. Approximate value iteration for risk-aware markov decision processes. *IEEE Transactions on Automatic Control*, 63(9):3135–3142, 2018. doi: 10.1109/TAC.2018.2790261.
- [54] Tiancheng Yu, Yi Tian, Jingzhao Zhang, and Suvrit Sra. Provably efficient algorithms for multi-objective competitive rl. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 12167–12176. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/yu21b.html>.
- [55] Lifeng Zhou and Pratap Tokekar. *An Approximation Algorithm for Risk-Averse Submodular Optimization*, pages 144–159. 05 2020. ISBN 978-3-030-44050-3. doi: 10.1007/978-3-030-44051-0_9.

A Supporting Details for Theoretical Results

A.1 Robustness in Value and Time

The so-called *reward-based* formulation in (4), as coined by Bisi et al. [6], Bonetti et al. [7], captures *per-stage risk*, in contrast to the return-based approach (3), which captures aggregate discounted risk. To see this, first note that the occupation measure can be reexpressed as

$$d\nu^\pi = \sum_{\tau=0}^{\infty} (1-\gamma)\gamma^\tau dp_\pi^\tau =: \sum_{\tau=0}^{\infty} dp_\pi^{T=\tau} p^T(\tau),$$

where T is a $\mathbb{N} \cup \{0\}$ -valued geometric random variable such that $P(T = \tau) =: p^T(\tau) := (1-\gamma)\gamma^\tau$, $\tau \in \mathbb{N} \cup \{0\}$, and p_π^T is now the conditional state-action distribution relative to T . In other words, ν^π may be seen as the marginalization of the disintegration $dp_\pi^{T=\tau}(s, a)p^T(\tau)$ relative to T , the latter interpreted as a random time. One can show that

$$-\text{CVaR}_{\nu^\pi}^\beta(-r(s, a)) = -\text{CVaR}^\beta(-r(s_T^\pi, a_T^\pi)),$$

where (s_T^π, a_T^π) denotes the state-action vector under policy π and evaluated at random time T . Specifically for the case of CVaR (and similarly for other OCEs), it can be readily seen that for $\beta = 1$ we recover the objective of (1) which can be equivalently written as $V(\pi) = (1-\gamma)^{-1}\mathbb{E}\{r(s_T^\pi, a_T^\pi)\}$, whereas if $\beta \downarrow 0$, we obtain

$$\lim_{\beta \downarrow 0} -\text{CVaR}_{\nu^\pi}^\beta(-r(s, a)) = \inf_{\tau} \text{ess inf}_{(s,a) \sim p_\pi^\tau} r(s, a),$$

which shows in particular that such reward-based RL formulations enforce *joint aversion in reward value and the time* for which the reward takes each value. For instance, for sufficiently small β , the functional $-\text{CVaR}^\beta(-r(s_T^\pi, a_T^\pi))$ will be more sensitive to very small rewards happening far in the future (and with low probability $(1-\gamma)\gamma^\tau$), as compared with (1) and (3), which both consider reward time averaging, heavily discounting future reward contributions.

A.2 Proof of Lemma 3.2

We will work with the convex analytic dual form for this proof, but replacing the occupancy measure with the expected discount sum will give the same result.

Statement The problem

$$\begin{aligned} P^* = \sup_{\pi \in \mathcal{R}} \sup_{t_0 \in \mathbb{R}} \mathbb{E}_{\nu^\pi(s,a)} \left[t_0 - \frac{1}{\beta}(t_0 - r_0(s, a))_+ \right] \\ \text{s.t. } \sup_{t_i \in \mathbb{R}} \mathbb{E}_{\nu^\pi(s,a)} \left[t_i - \frac{1}{\beta}(t_i - r_i(s, a))_+ \right] \geq c_i \quad \forall i = 1, 2, \dots, m \end{aligned} \quad (12)$$

and that of

$$\begin{aligned} \sup_{\pi \in \mathcal{R}, t_0 \in \mathbb{R}, t_i \in \mathbb{R}} \mathbb{E}_{\nu^\pi(s,a)} \left[t_0 - \frac{1}{\beta}(t_0 - r_0(s, a))_+ \right] \\ \text{s.t. } \mathbb{E}_{\nu^\pi(s,a)} \left[t_i - \frac{1}{\beta}(t_i - r_i(s, a))_+ \right] \geq c_i \quad \forall i = 1, 2, \dots, m \end{aligned} \quad (13)$$

are equivalent.

Proof. We need to prove that it is valid to move the supremum over each t_i to the objective. The proof follows that of Chow et al. [12].

- (Case (6) $\leq$ (7).) Let π^1 be a feasible policy for the problem (6). Then define

$$t_i^1 := \arg \max_{t_i \in \mathbb{R}} \mathbb{E}_{\nu^{\pi^1}(s,a)} \left[t_i - \frac{1}{\beta}(t_i - r_i(s, a))_+ \right] \geq c_i, \quad \forall i = 1, 2, \dots, m,$$

so that (π^1, t_i^1) is feasible for (6). Then this tuple (π^1, t_i^1) is clearly also feasible for (7).

- (Case (7) $\leq$ (6).) On the other hand, for any π , we clearly have that

$$-\text{CVaR}_{\nu^\pi}^\beta[-Z] = \sup_t \mathbb{E}_{\nu^\pi} \left[t - \frac{1}{\beta}(t - Z)_+ \right] \geq \mathbb{E}_{\nu^\pi} \left[\hat{t} - \frac{1}{\beta}(\hat{t} - Z)_+ \right],$$

for any $\hat{t}$. So, if the quantity on the RHS of the inequality is lower bounded by c_i (that is, feasible for (7)), then the LHS is clearly also lower bounded by c_i (thus, feasible for (6)).

□

A.3 Proof of Proposition 3.3

Statement Let r_i be bounded functions for all $i \in \{0\} \cup [m]$. Assume that Slater's condition holds for (9). Then, (9) exhibits strong duality, and thus

$$\sup_\pi \inf_{\lambda \geq 0} \mathcal{L}(\pi, t, \lambda) = \inf_{\lambda \geq 0} \sup_\pi \mathcal{L}(\pi, t, \lambda).$$

Proof. Paternain et al. [33, Theorem 1] proved strong duality holds for problems of the form

$$\sup_{\pi \in \mathcal{P}(S)} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau r_0(s_\tau^\pi, a_\tau^\pi) \right] \text{ s.t. } \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau r_i(s_\tau^\pi, a_\tau^\pi) \right] \geq c_i, \quad \forall i \in [m], \quad (14)$$

which involve the expectation of the discounted sums of some bounded reward functions r_i , for $i \in \{0\} \cup [m]$. Thus, the problem (9) is exactly in the form of Paternain et al. [33], with the reward functions $r'_i(s, a, t_i) = t_i - \frac{1}{\beta}(t_i - r_i(s, a))_+$, for $i \in \{0\} \cup [m]$. Since the original r_i are bounded and t_i are fixed, we have that r'_i are bounded and we can readily apply Paternain et al. [33, Theorem 1].

□

A.4 Proof of Theorem 3.5

Under Assumption 3.4, we have shown that the proposed partial Lagrangian relaxation is exact, i.e., that

$$\sup_{t \in \mathcal{I}} \sup_\pi \inf_{\lambda \geq 0} \mathcal{L}(\pi, t, \lambda) = \sup_{t \in \mathcal{I}} \inf_{\lambda \geq 0} \sup_\pi \mathcal{L}(\pi, t, \lambda).$$

We assume that π_θ , for $\theta \in \Theta \subset \mathbb{R}^p$ (with Θ some compact set), is an ϵ -universal parametrization of measures in $\mathcal{P}(S)$, according to Paternain et al. [33, Definition 1] and that the parametrized-policy version of (7) is feasible. We proceed to show that there is almost no price to pay for the policy parametrization, in terms of duality gap.

Indeed, assuming that the parametrized optimization problem is feasible, we can utilize Paternain et al. [33, Theorem 2] and Assumption 3.4 to deduce that, for all $t \in \mathcal{I}$, the following inequalities hold:

$$P^*(t) := \sup_\pi \inf_{\lambda \geq 0} \mathcal{L}(\pi, t, \lambda) \geq \inf_{\lambda \geq 0} \sup_\theta \mathcal{L}(\pi_\theta, t, \lambda) \geq P^*(t) - G(t)\epsilon/(1 - \gamma),$$

where $G(t) = \mathcal{O}(1)$ is a constant depending on t , and $\epsilon > 0$ is an arbitrarily small constant (depending on the employed parametrization). Let $t^* \in \arg \max_{t \in \mathcal{I}} P^*(t)$ and $t_\theta^* \in \arg \max_{t \in \mathcal{I}} \inf_{\lambda \geq 0} \sup_\theta \mathcal{L}(\pi_\theta, t, \lambda)$. Let also $G = \max\{G(t^*), G(t_\theta^*)\}$. Then, we obtain that

$$\begin{aligned} P^*(t^*) &\geq \inf_{\lambda \geq 0} \sup_\pi \mathcal{L}(\pi, t^*, \lambda) \geq P^*(t^*) - G\epsilon/(1 - \gamma), \\ P^*(t_\theta^*) &\geq \inf_{\lambda \geq 0} \sup_\theta \mathcal{L}(\pi_\theta, t_\theta^*, \lambda) \geq P^*(t_\theta^*) - G\epsilon/(1 - \gamma). \end{aligned}$$

Since, by definition, $P^*(t^*) \geq P^*(t_\theta^*)$, while $\sup_\theta \inf_{\lambda \geq 0} \mathcal{L}(\pi_\theta, t^*, \lambda) \leq \sup_\theta \inf_{\lambda \geq 0} \mathcal{L}(\pi_\theta, t_\theta^*, \lambda)$, we obtain that

$$P^*(t^*) \equiv \sup_{t \in \mathcal{I}} \sup_\pi \inf_{\lambda \geq 0} \mathcal{L}(\pi, t, \lambda) \geq \sup_{t \in \mathcal{I}} \inf_{\lambda \geq 0} \sup_\theta \mathcal{L}(\pi_\theta, t, \lambda) \geq P^*(t^*) - \mathcal{O}\left(\frac{\epsilon}{1 - \gamma}\right),$$

which completes the proof.

□

A.5 Discussion on Assumption 3.6

Below, we provide an insight as to why the inexact oracle utilized in Assumption 3.6 is reasonable under minimal conditions (under which the sample gradients used in Algorithm 1 satisfy Assumption 3.6). Specifically, building upon the discussion given later in Appendix A.7, we note that the most general condition guaranteeing Assumption 3.8 is given in the second scenario of Appendix A.7, which requires that the function $\mathcal{L}(\pi_{\theta^*}(t, \lambda), t, \lambda)$ satisfies the Lojasiewicz inequality with uniform constants over (t, λ) , for any selection $\pi_{\theta^*}(t, \lambda) \in \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_{\theta}, t, \lambda)$ (again, see Appendix A.7 for the definition of the Lojasiewicz inequality in this context). Since the (parametrized) policy optimization problem $\max_{\theta \in \Theta} \mathcal{L}(\pi_{\theta}, t, \lambda)$ is nonconvex (and is approximately solved using a solver like PPO), we can only expect to obtain an approximate solution $\theta^\dagger(t, \lambda)$ such that

$$\mathcal{L}(\pi_{\theta^*}(t, \lambda), t, \lambda) - \mathcal{L}(\pi_{\theta^\dagger}(t, \lambda), t, \lambda) \leq \tilde{\varepsilon}(\theta^\dagger, \theta^*, t, \lambda) \leq \varepsilon,$$

for some $\varepsilon > 0$ (uniformly bounded in (t, λ)). Then, if $\mathcal{L}(\pi_{\theta^*}(t, \lambda), t, \lambda)$ satisfies the Lojasiewicz inequality for some uniform positive constants C, η , we obtain that

$$\|\theta^\dagger - \theta^*\| \leq C\varepsilon^\eta.$$

Thus, under the mere assumption that the sample gradient function $\hat{\nabla}_{t, \lambda} \hat{\mathcal{L}}(\pi_{\theta}, t, \lambda)$ is $\hat{L}$ -Lipschitz continuous with respect to θ (which can be enforced, if necessary, by appropriate smoothing of the utility function g ; a general smoothing strategy for generating smooth OCEs from non-smooth utilities can be found in Kouri and Surowiec [22, Example 2]), we obtain that

$$\left\| \hat{\nabla}_{t, \lambda} \hat{\mathcal{L}}(\pi_{\theta^\dagger}(t, \lambda), t, \lambda) - \hat{\nabla}_{t, \lambda} \hat{\mathcal{L}}(\pi_{\theta^*}(t, \lambda), t, \lambda) \right\| \leq \hat{L} \|\theta^\dagger - \theta^*\| \leq \hat{L} C \varepsilon^\eta := \delta.$$

Upon noting that $\hat{\nabla}_{t, \lambda} \hat{\mathcal{L}}(\pi_{\theta^*}(t, \lambda), t, \lambda)$ is an unbiased sample of $-\nabla_{t, \lambda} f(t, \lambda)$ (cf. Appendix A.7), we can easily deduce that under the said minor conditions, the sample gradients utilized in Algorithm 1 satisfy Assumption 3.6. Thus, we can see that our oracle condition in Assumption 3.6 is justified and can be shown to hold under minimal conditions (without requiring convexity or uniqueness of the solution of $\max_{\theta \in \Theta} \mathcal{L}(\pi_{\theta}, t, \lambda)$).

A.6 Lipschitzness of the Lagrangian

First we show a technical lemma, and then we proceed with the proof of Lemma 3.7.

Lemma A.1. *For any pair $(\tilde{t}, \bar{t}) \in \mathbb{R}^{2m+2}$, and for any $\lambda \in \mathbb{R}^m$, it is true that*

$$\left| \sup_{\pi} \mathcal{L}(\pi, \tilde{t}, \lambda) - \sup_{\pi} \mathcal{L}(\pi, \bar{t}, \lambda) \right| \leq \sup_{\pi} |\mathcal{L}(\pi, \tilde{t}, \lambda) - \mathcal{L}(\pi, \bar{t}, \lambda)| \quad (15)$$

Proof of Lemma A.1 For any pair $\tilde{t}, \bar{t} \in \mathbb{R}$ it is true that

$$\sup_{\pi} \mathcal{L}(\pi, \tilde{t}, \lambda) = \sup_{\pi} \mathcal{L}(\pi, \tilde{t}, \lambda) - L(\pi, \bar{t}, \lambda) + L(\pi, \bar{t}, \lambda) \quad (16)$$

$$\leq \sup_{\pi} \mathcal{L}(\pi, \tilde{t}, \lambda) - \mathcal{L}(\pi, \bar{t}, \lambda) + \sup_{\pi} \mathcal{L}(\pi, \bar{t}, \lambda) \implies \quad (17)$$

$$\sup_{\pi} \mathcal{L}(\pi, \tilde{t}, \lambda) - \sup_{\pi} \mathcal{L}(\pi, \bar{t}, \lambda) \leq \sup_{\pi} \mathcal{L}(\pi, \tilde{t}, \lambda) - \mathcal{L}(\pi, \bar{t}, \lambda) \leq \sup_{\pi} |\mathcal{L}(\pi, \tilde{t}, \lambda) - \mathcal{L}(\pi, \bar{t}, \lambda)| \quad (18)$$

Similarly, we can show that

$$\sup_{\pi} \mathcal{L}(\pi, \bar{t}, \lambda) - \sup_{\pi} \mathcal{L}(\pi, \tilde{t}, \lambda) \leq \sup_{\pi} \mathcal{L}(\pi, \bar{t}, \lambda) - \mathcal{L}(\pi, \tilde{t}, \lambda) \leq \sup_{\pi} |\mathcal{L}(\pi, \bar{t}, \lambda) - \mathcal{L}(\pi, \tilde{t}, \lambda)|. \quad (19)$$

The last two displays give the inequality (15).

Proof of Lemma 3.7 For any pair $(\tilde{t}, \bar{t}) \in \mathbb{R}^{2m+2}$, we find an upper bound on the term $|\mathcal{L}(\pi, \tilde{t}, \lambda) - \mathcal{L}(\pi, \bar{t}, \lambda)|$ as follows

$$\begin{aligned} & |\mathcal{L}(\pi, \tilde{t}, \lambda) - \mathcal{L}(\pi, \bar{t}, \lambda)| \\ & \leq \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau \left| \left(\tilde{t}_0 - \frac{1}{\beta} (\tilde{t}_0 - r_0(s_\tau, a_\tau))_+ \right) - \left(\bar{t}_0 - \frac{1}{\beta} (\bar{t}_0 - r_0(s_\tau, a_\tau))_+ \right) \right| \right] \end{aligned}$$

$$\begin{aligned}
& + \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} \left| \sum_{i=1}^m \lambda_i \left(\tilde{t}_i - \frac{1}{\beta_i} (\tilde{t}_i - r_i(s_{\tau}, a_{\tau}))_+ \right) \right. \right. \\
& \quad \left. \left. - \sum_{i=1}^m \lambda_i \left(\bar{t}_i - \frac{1}{\beta_i} (\bar{t}_i - r_i(s_{\tau}, a_{\tau}))_+ \right) \right| \right] \\
& = \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} \left| \left(\tilde{t}_0 - \frac{1}{\beta} \max \{ \tilde{t}_0 - r_0(s_{\tau}, a_{\tau}), 0 \} \right) - \left(\bar{t}_0 - \frac{1}{\beta} \max \{ \bar{t}_0 - r_0(s_{\tau}, a_{\tau}), 0 \} \right) \right| \right] \\
& \quad + \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} \left| \sum_{i=1}^m \lambda_i \left(\tilde{t}_i - \frac{1}{\beta_i} \max \{ \tilde{t}_i - r_i(s_{\tau}, a_{\tau}), 0 \} \right) \right. \right. \\
& \quad \left. \left. - \sum_{i=1}^m \lambda_i \left(\bar{t}_i - \frac{1}{\beta_i} \max \{ \bar{t}_i - r_i(s_{\tau}, a_{\tau}), 0 \} \right) \right| \right] \\
& \leq \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} (|\tilde{t}_0 - \bar{t}_0| + \frac{1}{\beta} |\max \{ \tilde{t}_0 - r_0(s_{\tau}, a_{\tau}), 0 \} - \max \{ \bar{t}_0 - r_0(s_{\tau}, a_{\tau}), 0 \}|) \right] \\
& \quad + \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} \sum_{i=1}^m \lambda_i (|\tilde{t}_i - \bar{t}_i| + \frac{1}{\beta_i} |\max \{ \tilde{t}_i - r_i(s_{\tau}, a_{\tau}), 0 \} - \max \{ \bar{t}_i - r_i(s_{\tau}, a_{\tau}), 0 \}|) \right] \\
& \leq \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} (|\tilde{t}_0 - \bar{t}_0| + \frac{1}{\beta} \max \{ |\tilde{t}_0 - r_0(s_{\tau}, a_{\tau}) - \bar{t}_0 - r_0(s_{\tau}, a_{\tau})|, 0 \}) \right] \\
& \quad + \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} \sum_{i=1}^m \lambda_i (|\tilde{t}_i - \bar{t}_i| + \frac{1}{\beta_i} \max \{ |\tilde{t}_i - r_i(s_{\tau}, a_{\tau}) - \bar{t}_i - r_i(s_{\tau}, a_{\tau})|, 0 \}) \right] \\
& = \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} (|\tilde{t}_0 - \bar{t}_0| + \frac{1}{\beta} |\tilde{t}_0 - \bar{t}_0|) \right] + \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} \sum_{i=1}^m \lambda_i (|\tilde{t}_i - \bar{t}_i| + \frac{1}{\beta_i} |\tilde{t}_i - \bar{t}_i|) \right] \\
& \leq (1 - \gamma)^{-1} (1 + \frac{1}{\beta}) |\tilde{t}_0 - \bar{t}_0| + (1 - \gamma)^{-1} \sum_{i=1}^m \lambda_i (1 + \frac{1}{\beta_i}) |\tilde{t}_i - \bar{t}_i| \\
& \leq (1 - \gamma)^{-1} \|\tilde{t} - \bar{t}\| \sqrt{(1 + \frac{1}{\beta})^2 + \sum_{i=1}^m \lambda_i^2 (1 + \frac{1}{\beta_i})^2}.
\end{aligned}$$

Then, we apply Lemma A.1 to get

$$|\sup_{\pi} \mathcal{L}(\pi, \tilde{t}, \lambda) - \sup_{\pi} \mathcal{L}(\pi, \bar{t}, \lambda)| \leq \sup_{\pi} |\mathcal{L}(\pi, \tilde{t}, \lambda) - \mathcal{L}(\pi, \bar{t}, \lambda)| \quad (20)$$

$$\leq (1 - \gamma)^{-1} \|\tilde{t} - \bar{t}\| \sqrt{(1 + \frac{1}{\beta})^2 + \sum_{i=1}^m \lambda_i^2 (1 + \frac{1}{\beta_i})^2}. \quad (21)$$

Thus the function $\mathcal{L}(\pi, \cdot, \lambda)$ is Lipschitz, with Lipschitz constant

$$c(\gamma, \beta, \lambda) := (1 - \gamma)^{-1} \sqrt{(1 + \frac{1}{\beta})^2 + \sum_{i=1}^m \lambda_i^2 (1 + \frac{1}{\beta_i})^2}. \quad (22)$$

The result then follows since we are interested in multipliers $\lambda \in \Lambda$, with Λ a convex and compact set (and thus, there exists a constant $C \geq c(\gamma, \beta, \lambda)$, independent of λ and t , for which $\mathcal{L}(\pi, \cdot, \lambda)$ is Lipschitz continuous for all $(t, \lambda) \in \mathcal{T} \times \Lambda$). $\square$

A.7 Conditions That Guarantee Assumption 3.8

Let us note that Assumption 3.8 is not particularly restrictive in our setting. In what follows, we discuss some general conditions under which this holds readily, as well as a simple methodology of

enforcing this assumption, if necessary, by a slight algorithmic adjustment. Before we do this, let us first note that $\mathcal{L}(\pi_\theta, \cdot, \lambda)$ is not necessarily differentiable (unless the probability of $t_i - r_i(s, a) = 0$ is 0), because of the plus function $(\cdot)_+$. Nonetheless, adding a small random noise to the reward following a random variable with a smooth mollifier density ensures that the function is infinitely differentiable with respect to both t and λ . Note that, since we assume that the reward is bounded, we cannot (in theory) add Gaussian noise, but any smooth and compactly supported mollifier-based noise is admissible. Thus, we assume (without loss of generality) that $\mathcal{L}(\pi_\theta, \cdot, \lambda)$ is differentiable (in fact, to an arbitrary degree).

Next, we focus on the task of assessing the differentiability properties of $f(t, \lambda)$. Specifically, we identify three general scenarios that guarantee Assumption 3.8 in our case.

1. The first scenario, which is somewhat restrictive, relies on the assumption that, for each $(t, \lambda) \in \mathcal{T} \times \Lambda$, the maximization problem, with respect to θ , admits a unique solution. In that case, we can readily utilize Lemma 2.2 of Shapiro [40], which states that $\mathcal{L}(\pi_{\theta^*(t, \lambda)}, t, \lambda)$ is twice-continuously differentiable. Since $\mathcal{T} \times \Lambda$ is a compact set, this immediately implies that the function is ℓ -smooth over $\mathcal{T} \times \Lambda$, for some constant $\ell > 0$.
2. The second scenario is significantly more general, and relies on the strong second-order sufficient optimality conditions for the maximization over θ and the assumption that the function $\mathcal{L}(\pi_{\theta^*(t, \lambda)}, t, \lambda)$ satisfies the Lojasiewicz inequality with some constant uniform over (t, λ) , for any $\pi_{\theta^*(t, \lambda)} \in \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$, that is, for each $(t, \lambda) \in \mathcal{T} \times \Lambda$, there exist constants $C > 0$ and $\eta > 0$ such that

$$\text{dist} \left(\theta, \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda) \right) \leq C \left(\max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda) - \mathcal{L}(\pi_\theta, t, \lambda) \right)^\eta.$$

We note that the previous holds in several cases (typically with uniform exponent η). For example, this is true when $\max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, \cdot, \cdot)$ is sub-analytic (which, for example, is implied if $\mathcal{L}(\pi_\theta, t, \lambda)$ is analytic, noting that this can easily be enforced in our setting). Since $\mathcal{T} \times \Lambda$ is a compact and convex set, we can readily utilize Theorem 11 from [16] to deduce that under this general assumption, the function $\mathcal{L}(\pi_{\theta^*(t, \lambda)}, t, \lambda)$ is ℓ -smooth over $\mathcal{T} \times \Lambda$, irrespectively of the selections $\theta^*(t, \lambda) \in \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$.

3. Under certain qualification conditions, laid out in Section 4 of Shapiro [40] (that do not require the second-order strong sufficient conditions for the maximization problem over θ), ℓ -smoothness of $\max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$ is also guaranteed. We refer the reader to Shapiro [40] for additional details, since this analysis directly applicable in our setting.

If any of the above conditions holds, then Assumption 3.8 is true, and we can readily compute the gradient of f as

$$\nabla_{t, \lambda} f(t, \lambda) = -\nabla_{t, \lambda} \mathcal{L}(\pi_\theta, t, \lambda)|_{\pi_\theta = \pi_{\theta^*(t, \lambda)}},$$

where $\theta^*(t, \lambda) \in \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$ is an arbitrary selection.

If none of the above is true, we can still guarantee that Assumption 3.8 holds by slightly altering Algorithm 1. Specifically, for some small $\mu > 0$, we can define the surrogate function

$$\mathcal{L}_\mu(\pi_{\theta^*(t, \lambda)}, t, \lambda) := \mathbb{E}_{U_1, U_2} \left\{ \mathcal{L}(\pi_{\theta^*(t + \mu U_1, \lambda + \mu U_2)}, t + \mu U_1, \lambda + \mu U_2) \right\}, \quad (23)$$

where U_1, U_2 follow a uniform distribution over the unit ball (of appropriate dimension in each case). Then, this new surrogate function can be made arbitrarily close to $\max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t, \lambda)$ (where the proximity is uniform in the constant μ), while at the same time being ℓ_μ -smooth, with $\ell_\mu = \mathcal{O}(1/\mu)$.

Then, we can substitute the original Lagrangian with the smoothed Lagrangian given in (23) and run Algorithm 1 on the surrogate problem. Following standard zeroth-order optimization techniques (e.g., see Pougkakiotis and Kalogerias [34]), we note that, in this case, the only algorithmic adjustments that need to be made relate to the computation of the sample gradients of f , which would require two (instead of one) evaluations of the policy oracle. In light of this discussion, we conclude that Assumption 3.8 can be utilized (almost) without loss of generality.

A.8 Proof of Main Theorem

In this section, we provide the proof of Theorem 3.12. To prove the main theorem we will need the following auxiliary lemmas. The proof extends the analysis given in [25], since the presented algorithm only assumes having access to biased and inexact (sub)gradient samples.

Denote by $b_1(\theta^*, \theta^\dagger, t, \lambda)$, $b_2(\theta^*, \theta^\dagger, t, \lambda)$ the components of $b(\theta^*, \theta^\dagger, t, \lambda)$ corresponding to t and λ from Assumption 3.6.

Lemma A.2. $\widehat{\nabla}_t \widehat{\mathcal{L}}(\pi_{\theta^\dagger(t, \lambda)}, t, \lambda)$ and $\widehat{\nabla}_\lambda \widehat{\mathcal{L}}(\pi_{\theta^\dagger(t, \lambda)}, t, \lambda)$ have bounded bias and variance.

Proof. By Assumption 3.6 and by linearity of expectation, we have

$$\begin{aligned} \mathbb{E} \left[\widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^\dagger(t, \lambda)}, t, \lambda) \right] &= \mathbb{E} \left[\widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^*(t, \lambda)}, t, \lambda) + b(\theta^*, \theta^\dagger, t, \lambda) \right] \\ &= -\nabla f(t, \lambda) + \mathbb{E}[b(\theta^*, \theta^\dagger, t, \lambda)]. \quad (\text{by Assumptions 3.6, 3.10}) \end{aligned}$$

The bias of the gradient is controlled by $b(\theta^*, \theta^\dagger, t, \lambda)$, whose norm is at most δ . Furthermore,

$$\begin{aligned} &\mathbb{E} \left[\left\| \widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^\dagger(t, \lambda)}, t, \lambda) \right\|^2 \right] \\ &= \mathbb{E} \left[\left\| \widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^\dagger(t, \lambda)}, t, \lambda) - \widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^*(t, \lambda)}, t, \lambda) + \widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^*(t, \lambda)}, t, \lambda) \right\|^2 \right] \\ &\leq 2\mathbb{E} \left[\left\| b(\theta^*, \theta^\dagger, t, \lambda) \right\|^2 \right] + 2\mathbb{E} \left[\left\| \widehat{\nabla}_{t, \lambda} \widehat{\mathcal{L}}(\pi_{\theta^*(t, \lambda)}, t, \lambda) \right\|^2 \right] \\ &\leq 2 \left\| \nabla f(t, \lambda) \right\|^2 + 2(\sigma^2 + \delta^2). \quad (\text{by Lemma A.2 of [25] with } M = 1 \text{ and Assumptions 3.6, 3.10}) \end{aligned}$$

□

Lemma A.3. Let $\Delta^{(k)} = \mathbb{E} [\Phi(t^{(k)}) - f(t^{(k)}, \lambda^{(k)})]$. The following holds for Algorithm 1:

$$\begin{aligned} \mathbb{E} [\Phi_{1/2\ell}(t^{(k)})] &\leq \mathbb{E} [\Phi_{1/2\ell}(t^{(k-1)})] + 2\eta_t \ell \Delta^{(k-1)} - \frac{\eta_t}{4} \mathbb{E} \left[\left\| \nabla \Phi_{1/2\ell}(t^{(k-1)}) \right\|^2 \right] \\ &\quad + 2\eta_t \delta \ell \cdot \text{diam}(\mathcal{T}) + 3\eta_t^2 \ell \cdot (C^2 + \sigma^2 + \delta^2). \end{aligned}$$

Proof. Let $\widehat{t}^{(k-1)} = \text{prox}_{\Phi/2\ell}(t^{(k-1)})$ and fix $\pi_{\theta^*} = \pi_{\theta^*(t^{(k-1)}, \lambda^{(k-1)})}$ where $\theta^* \in \arg \max_{\theta \in \Theta} \mathcal{L}(\pi_\theta, t^{(k-1)}, \lambda^{(k-1)})$, and $\pi_{\theta^\dagger} = \pi_{\theta^\dagger(t^{(k-1)}, \lambda^{(k-1)})}$ satisfying Assumption 3.6. Then,

$$\begin{aligned} \left\| \widehat{t}^{(k-1)} - t^{(k)} \right\|^2 &= \left\| \widehat{t}^{(k-1)} - \Pi_{\mathcal{T}} \left(t^{(k-1)} - \eta_t \cdot \widehat{\nabla}_t \widehat{\mathcal{L}}(\pi_{\theta^\dagger}, t^{(k-1)}, \lambda^{(k-1)}) \right) \right\|^2 \\ &\leq \left\| \widehat{t}^{(k-1)} - t^{(k-1)} + \eta_t \cdot \widehat{\nabla}_t \widehat{\mathcal{L}}(\pi_{\theta^\dagger}, t^{(k-1)}, \lambda^{(k-1)}) \right\|^2 \\ &\quad (\text{projection is non-expansive}) \\ &= \left\| \widehat{t}^{(k-1)} - t^{(k-1)} \right\|^2 + \eta_t^2 \cdot \left\| \widehat{\nabla}_t \widehat{\mathcal{L}}(\pi_{\theta^\dagger}, t^{(k-1)}, \lambda^{(k-1)}) \right\|^2 \\ &\quad + 2\eta_t \left\langle \widehat{t}^{(k-1)} - t^{(k-1)}, \widehat{\nabla}_t \widehat{\mathcal{L}}(\pi_{\theta^\dagger}, t^{(k-1)}, \lambda^{(k-1)}) \right\rangle. \end{aligned}$$

Taking expectations, conditioned on (t_{k-1}, λ_{k-1}) , and using Assumption 3.6 we have

$$\begin{aligned} &\mathbb{E} \left[\left\| \widehat{t}^{(k-1)} - t^{(k)} \right\|^2 \mid (t^{(k-1)}, \lambda^{(k-1)}) \right] \\ &\leq \mathbb{E} \left[\left\| \widehat{t}^{(k-1)} - t^{(k-1)} \right\|^2 \mid (t^{(k-1)}, \lambda^{(k-1)}) \right] \\ &\quad + 2\eta_t \cdot \left\langle t^{(k-1)} - \widehat{t}^{(k-1)}, \nabla_t f(t^{(k-1)}, \lambda^{(k-1)}) - b(\theta^*, \theta^\dagger, t^{(k-1)}, \lambda^{(k-1)}) \right\rangle \\ &\quad + 3\eta_t^2 \left\| b_1(\theta^*, \theta^\dagger, t^{(k-1)}, \lambda^{(k-1)}) \right\|^2 + 3\eta_t^2 \left\| \nabla_t f(t^{(k-1)}, \lambda^{(k-1)}) \right\|^2 \\ &\quad + 3\eta_t^2 \mathbb{E} \left[\left\| \widehat{\nabla}_t \widehat{\mathcal{L}}(\pi_{\theta^*}, t^{(k-1)}, \lambda^{(k-1)}) + \nabla f(t^{(k-1)}, \lambda^{(k-1)}) \right\|^2 \mid (t^{(k-1)}, \lambda^{(k-1)}) \right]. \\ &\quad (\text{using } \|x + y + z\|^2 \leq 3\|x\|^2 + 3\|y\|^2 + 3\|z\|^2) \end{aligned}$$

Taking expectations of both sides and using Lemma A.2 of [25], that $f(\cdot, \lambda)$ is C -Lipschitz (cf. Lemma 3.7), and Assumption 3.6,

$$\begin{aligned} \mathbb{E} \left\| \hat{t}^{(k-1)} - t^{(k)} \right\|^2 &\leq \mathbb{E} \left\| \hat{t}^{(k-1)} - t^{(k-1)} \right\|^2 \\ &\quad + 2\eta_t \left\langle \hat{t}^{(k-1)} - t^{(k-1)}, \nabla f(t^{(k-1)}, \lambda^{(k-1)}) - b(\theta^*, \theta^\dagger, t^{(k-1)}, \lambda^{(k-1)}) \right\rangle \\ &\quad + 3\eta_t^2 (C^2 + \sigma^2 + \delta^2). \end{aligned}$$

Since $\mathcal{T}$ is compact, define $\text{diam}(\mathcal{T}) := \max_{x, y \in \mathcal{T}} \|x - y\|$. Using Cauchy-Schwarz on the term

$$\begin{aligned} 2\eta_t \left\langle t^{(k-1)} - \hat{t}^{(k-1)}, b_1(\theta^*, \theta^\dagger, t^{(k-1)}, \lambda^{(k-1)}) \right\rangle \\ \leq 2\eta_t \left\| \hat{t}^{(k-1)} - t^{(k-1)} \right\| \cdot \left\| b_1(\theta^*, \theta^\dagger, t^{(k-1)}, \lambda^{(k-1)}) \right\|, \end{aligned}$$

and substituting, yields

$$\begin{aligned} \mathbb{E} \left\| \hat{t}^{(k-1)} - t^{(k)} \right\|^2 &\leq \mathbb{E} \left\| \hat{t}^{(k-1)} - t^{(k-1)} \right\|^2 + 2\eta_t \left\langle \hat{t}^{(k-1)} - t^{(k-1)}, \nabla f(t^{(k-1)}, \lambda^{(k-1)}) \right\rangle \\ &\quad + 2\eta_t \delta \cdot \text{diam}(\mathcal{T}) + 3\eta_t^2 (C^2 + \sigma^2 + \delta^2). \end{aligned} \quad (24)$$

From [25, Eq. (20)], we obtain

$$\Phi_{1/2\ell}(t^{(k)}) \leq \Phi_{1/2\ell}(t^{(k-1)}) + \ell \left\| \hat{t}^{(k-1)} - t^{(k)} \right\|^2 \quad (25)$$

while, from [25, Enq. (23), (24)], we obtain

$$\begin{aligned} &\left\langle \hat{t}^{(k-1)} - t^{(k-1)}, \nabla_t f(t^{(k-1)}, \lambda^{(k-1)}) \right\rangle \\ &\leq f(\hat{t}^{(k-1)}, \lambda^{(k-1)}) - f(t^{(k-1)}, \lambda^{(k-1)}) + \frac{\ell}{2} \left\| \hat{t}^{(k-1)} - t^{(k-1)} \right\|^2 \end{aligned} \quad (26)$$

$$\begin{aligned} &f(\hat{t}^{(k-1)}, \lambda^{(k-1)}) - f(t^{(k-1)}, \lambda^{(k-1)}) \\ &\leq \Phi(\hat{t}^{(k-1)}) - f(t^{(k-1)}, \lambda^{(k-1)}) \leq \Delta^{(k-1)} - \frac{\ell}{2} \left\| \hat{t}^{(k-1)} - t^{(k-1)} \right\|^2. \end{aligned} \quad (27)$$

Plugging in (24) to (25), then combining (26), (27), and the fact that $\left\| \hat{t}^{(k-1)} - t^{(k-1)} \right\| = \left\| \nabla \Phi_{1/2\ell}(t^{(k-1)}) \right\| / 2\ell$, finally gives

$$\begin{aligned} \mathbb{E} \left[\Phi_{1/2\ell}(t^{(k)}) \right] &\leq \mathbb{E} \left[\Phi_{1/2\ell}(t^{(k-1)}) \right] + 2\eta_t \ell \Delta^{(k-1)} - \frac{\eta_t}{4} \mathbb{E} \left\| \nabla \Phi_{1/2\ell}(t^{(k-1)}) \right\|^2 \\ &\quad + 2\eta_t \delta \ell \cdot \text{diam}(\mathcal{T}) + 3\eta_t^2 \ell (C^2 + \sigma^2 + \delta^2), \end{aligned}$$

which completes the proof. $\square$

Lemma A.4. Let $\Delta^{(k)} = \mathbb{E} [\Phi(t^{(k)}) - f(t^{(k)}, \lambda^{(k)})]$ and $\lambda^*(t) \in \arg \max_{\lambda \in \Lambda} f(t, \lambda)$. The following holds for all $s \leq k - 1$:

$$\begin{aligned} \Delta^{(k-1)} &\leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} (2t - 2s - 1) \\ &\quad + \frac{1}{2\eta_\lambda} \left(\mathbb{E} \left\| \lambda^{(k-1)} - \lambda^*(t^{(s)}) \right\|^2 - \mathbb{E} \left\| \lambda^{(k)} - \lambda^*(t^{(s)}) \right\|^2 \right) \\ &\quad + \mathbb{E} \left[f(t^{(k)}, \lambda^{(k)}) - f(t^{(k-1)}, \lambda^{(k-1)}) \right] + \delta \cdot \text{diam}(\Lambda) + 2\eta_\lambda (\sigma^2 + \delta^2). \end{aligned}$$

Proof. For any $\lambda \in \Lambda$, the update of $\lambda^{(k)}$ and convexity of Λ imply that

$$(\lambda - \lambda^{(k)})^\top \left(\lambda^{(k)} - \lambda^{(k-1)} - \eta_\lambda \nabla_\lambda f(t^{(k-1)}, \lambda^{(k-1)}) \right) \geq 0.$$

Then, we have

$$\left\| \lambda - \lambda^{(k)} \right\|^2 \leq 2\eta_\lambda (\lambda^{(k-1)} - \lambda)^\top \nabla_\lambda f(t^{(k-1)}, \lambda^{(k-1)})$$

$$+ 2\eta_\lambda \delta \cdot \text{diam}(\Lambda) + 4\eta_\lambda^2(\sigma^2 + \delta^2).$$

Substituting $\lambda = \lambda^*(t^{(s)})$ (where $s \leq k-1$),

$$\begin{aligned} & \mathbb{E} \left[f(t^{(k-1)}, \lambda^*(t^{(s)})) - f(t^{(k-1)}, \lambda^{(k)}) \right] \\ & \leq \frac{1}{2\eta_\lambda} \left(\mathbb{E} \left\| \lambda^{(k-1)} - \lambda^*(t^{(s)}) \right\|^2 - \mathbb{E} \left\| \lambda^{(k)} - \lambda^*(t^{(s)}) \right\|^2 \right) + \delta \cdot \text{diam}(\Lambda) + 2\eta_\lambda(\sigma^2 + \delta^2). \end{aligned}$$

By the definition of $\Delta^{(k-1)}$,

$$\begin{aligned} \Delta^{(k-1)} & \leq \mathbb{E} \left[f(t^{(k-1)}, \lambda^*(t^{(k-1)})) - f(t^{(k-1)}, \lambda^*(t^{(s)})) \right. \\ & \quad \left. + (f(t^{(k)}, \lambda^{(k)}) - f(t^{(k-1)}, \lambda^{(k-1)})) + (f(t^{(k-1)}, \lambda^{(k)}) - f(t^{(k)}, \lambda^{(k)})) \right] \\ & \quad + \frac{1}{2\eta_\lambda} \left(\mathbb{E} \left\| \lambda^{(k-1)} - \lambda^*(t^{(s)}) \right\|^2 - \mathbb{E} \left\| \lambda^{(k)} - \lambda^*(t^{(s)}) \right\|^2 \right) \\ & \quad + \delta \cdot \text{diam}(\Lambda) + 2\eta_\lambda(\sigma^2 + \delta^2). \end{aligned}$$

Following the steps in [25, Lemma D.4], using that $f(\cdot, \lambda)$ is C -Lipschitz (by Lemma 3.7) and Lemma A.2, we have

$$\begin{aligned} \mathbb{E} \left[f(t^{(k-1)}, \lambda^*(t^{(k-1)})) - f(t^{(s)}, \lambda^*(t^{(k-1)})) \right] & \leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} (t-1-s) \\ \mathbb{E} \left[f(t^{(s)}, \lambda^*(t^{(s)})) - f(t^{(k-1)}, \lambda^*(t^{(s)})) \right] & \leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} (t-1-s) \\ \mathbb{E} \left[f(t^{(k-1)}, \lambda^{(k)}) - f(t^{(k)}, \lambda^{(k)}) \right] & \leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2}. \end{aligned}$$

Using [25, Eqn. (25)],

$$\begin{aligned} \Delta^{(k-1)} & \leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} \cdot (2t-2s-1) \\ & \quad + \frac{1}{2\eta_\lambda} \left(\mathbb{E} \left\| \lambda^{(k-1)} - \lambda^*(t^{(s)}) \right\|^2 - \mathbb{E} \left\| \lambda^{(k)} - \lambda^*(t^{(s)}) \right\|^2 \right) \\ & \quad + \mathbb{E} \left[f(t^{(k)}, \lambda^{(k)}) - f(t^{(k-1)}, \lambda^{(k-1)}) \right] + \delta \cdot \text{diam}(\Lambda) + 2\eta_\lambda(\sigma^2 + \delta^2). \end{aligned}$$

□

Lemma A.5. Let $\Delta^{(k)} = \mathbb{E} [\Phi(t^{(k)}) - f(t^{(k)}, \lambda^{(k)})]$. Let $B \leq J+1$ be such that $(J+1)/B$ is an integer. The following holds

$$\begin{aligned} \frac{1}{J+1} \left(\sum_{k=0}^J \Delta^{(k)} \right) & \leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} (B+1) \\ & \quad + \frac{\text{diam}(\Lambda)^2}{2B\eta_\lambda} + \delta \cdot \text{diam}(\Lambda) + 2\eta_\lambda(\sigma^2 + \delta^2) + \frac{\widehat{\Delta}_0}{J+1}. \end{aligned}$$

Proof. We divide $\{\Delta^{(k)}\}_{k=0}^J$ into blocks where each block contains at most B terms:

$$\left\{ \Delta^{(k)} \right\}_{k=0}^{B-1}, \left\{ \Delta^{(k)} \right\}_{k=B}^{2B-1}, \dots, \left\{ \Delta^{(k)} \right\}_{k=J-B+1}^J.$$

Then,

$$\frac{1}{J+1} \left(\sum_{k=0}^J \Delta^{(k)} \right) \leq \frac{B}{J+1} \left[\sum_{i=0}^{(J+1)/B-1} \left(\frac{1}{B} \sum_{k=iB}^{(i+1)B-1} \Delta^{(k)} \right) \right]. \quad (28)$$

Letting $s = 0$ and applying Lemma A.4,

$$\begin{aligned}
\sum_{k=0}^{B-1} \Delta^{(k)} &\leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} B^2 + \frac{1}{2\eta_\lambda} \mathbb{E} \left\| \lambda^{(0)} - \lambda^*(t^{(0)}) \right\|^2 \\
&\quad + \mathbb{E} \left[f(t^{(B)}, \lambda^{(B)}) - f(t^{(0)}, \lambda^{(0)}) \right] + \delta B \cdot \text{diam}(\Lambda) + 2\eta_\lambda B(\sigma^2 + \delta^2) \\
&\leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} B^2 + \mathbb{E} \left[f(t^{(B)}, \lambda^{(B)}) - f(t^{(0)}, \lambda^{(0)}) \right] \\
&\quad + \frac{\text{diam}(\Lambda)^2}{2\eta_\lambda} + \delta B \cdot \text{diam}(\Lambda) + 2\eta_\lambda B(\sigma^2 + \delta^2).
\end{aligned}$$

Letting $s = iB$ and applying Lemma A.4,

$$\begin{aligned}
\sum_{k=iB}^{(i+1)B-1} \Delta^{(k)} &\leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} B^2 + \mathbb{E} \left[f(t^{(iB+B)}, \lambda^{(iB+B)}) - f(t^{(iB)}, \lambda^{(iB)}) \right] \\
&\quad + \frac{\text{diam}(\Lambda)^2}{2\eta_\lambda} + \delta B \cdot \text{diam}(\Lambda) + 2\eta_\lambda B(\sigma^2 + \delta^2).
\end{aligned}$$

Substituting these into (28):

$$\begin{aligned}
\frac{1}{J+1} \left(\sum_{k=0}^J \Delta^{(k)} \right) &\leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} B + \frac{1}{J+1} \mathbb{E} [f(t_{J+1}, \lambda_{J+1}) - f(t_0, \lambda_0)] \\
&\quad + \frac{\text{diam}(\Lambda)^2}{2B\eta_\lambda} + \delta \cdot \text{diam}(\Lambda) + 2\eta_\lambda(\sigma^2 + \delta^2).
\end{aligned}$$

By the Lipschitzness of $f(\cdot, \lambda)$,

$$f(t_{J+1}, \lambda_{J+1}) - f(t_0, \lambda_0) \leq \eta_t C^2 (J+1) + \widehat{\Delta}_0.$$

This yields

$$\begin{aligned}
\frac{1}{J+1} \left(\sum_{k=0}^J \Delta_k \right) &\leq 2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} (B+1) \\
&\quad + \frac{\text{diam}(\Lambda)^2}{2B\eta_\lambda} + \delta \cdot \text{diam}(\Lambda) + 2\eta_\lambda(\sigma^2 + \delta^2) + \frac{\widehat{\Delta}_0}{J+1}.
\end{aligned}$$

□

Proof of Theorem 3.12

Summing up the inequality from Lemma A.3, over $k = 1, \dots, J+1$, yields

$$\begin{aligned}
\mathbb{E} \left[\Phi_{1/2\ell}(t^{(J+1)}) \right] &\leq \Phi_{1/2\ell}(t^{(0)}) + 2\eta_t \ell \left(\sum_{k=0}^J \Delta^{(k)} \right) - \frac{\eta_t}{4} \left(\sum_{k=0}^J \mathbb{E} \left\| \nabla \Phi_{1/2\ell}(t^{(k)}) \right\|^2 \right) \\
&\quad + (2\eta_t \delta \ell \cdot \text{diam}(\mathcal{T}) + 3\eta_t^2 \ell (C^2 + \sigma^2 + \delta^2)) (J+1).
\end{aligned}$$

Applying Lemma A.5,

$$\begin{aligned}
\mathbb{E} \left[\Phi_{1/2\ell}(t^{(J+1)}) \right] &\leq \Phi_{1/2\ell}(t^{(0)}) \\
&\quad + 2\eta_t \ell (J+1) \left(2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} (B+1) \right. \\
&\quad \left. + \frac{\text{diam}(\Lambda)^2}{2B\eta_\lambda} + \delta \cdot \text{diam}(\Lambda) + 2\eta_\lambda(\sigma^2 + \delta^2) \right)
\end{aligned}$$

$$\begin{aligned}
& + 2\eta_t \ell \widehat{\Delta}_0 - \frac{\eta_t}{4} \left(\sum_{k=0}^J \mathbb{E} \left\| \nabla \Phi_{1/2\ell}(t^{(k)}) \right\|^2 \right) \\
& + (2\eta_t \delta \ell \cdot \text{diam}(\mathcal{T}) + 3\eta_t^2 \ell (C^2 + \sigma^2 + \delta^2)) (J+1).
\end{aligned}$$

By the definition of $\widehat{\Delta}_\Phi$, we obtain

$$\begin{aligned}
& \frac{1}{J+1} \left(\sum_{k=0}^J \mathbb{E} \left\| \nabla \Phi_{1/2\ell}(t^{(k)}) \right\|^2 \right) \\
& \leq \frac{4\widehat{\Delta}_\Phi}{\eta_t(J+1)} + 8\ell \left(2\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2} (B+1) + \frac{\text{diam}(\Lambda)^2}{2B\eta_\lambda} + \delta \text{diam}(\Lambda) + 2\eta_\lambda (\sigma^2 + \delta^2) \right) \\
& \quad + \frac{8\ell \widehat{\Delta}_0}{J+1} + 12\eta_t \ell (C^2 + \sigma^2 + \delta^2) + 8\ell \delta \cdot \text{diam}(\mathcal{T}).
\end{aligned}$$

Letting $B = \frac{\text{diam}(\Lambda)}{2} \sqrt{\frac{1}{\eta_t \eta_\lambda C \sqrt{C^2 + \sigma^2 + \delta^2}}}$, we obtain

$$\begin{aligned}
\frac{1}{J+1} \left(\sum_{k=0}^J \mathbb{E} \left\| \nabla \Phi_{1/2\ell}(t^{(k)}) \right\|^2 \right) & \leq \frac{4\widehat{\Delta}_\Phi}{\eta_t(J+1)} + 24\ell \text{diam}(\Lambda) \sqrt{\frac{\eta_t C \sqrt{C^2 + \sigma^2 + \delta^2}}{\eta_\lambda}} \\
& \quad + 16\eta_\lambda \ell (\sigma^2 + \delta^2) + \frac{8\ell \widehat{\Delta}_0}{J+1} \\
& \quad + 12\eta_t \ell (C^2 + \sigma^2 + \delta^2) + 8\ell \delta \cdot (\text{diam}(\mathcal{T}) + \text{diam}(\Lambda)).
\end{aligned}$$

With the choice of step sizes

$$\eta_\lambda = \min \left\{ \frac{1}{2\ell}, \frac{\epsilon^2}{16\ell(\sigma^2 + \delta^2)} \right\}, \tag{29}$$

$$\eta_t = \min \left\{ \begin{aligned} & \mathcal{O} \left(\frac{\epsilon^2}{\ell(C^2 + \sigma^2 + \delta^2)} \right), \\ & \mathcal{O} \left(\frac{\epsilon^4}{\ell^3 \text{diam}(\Lambda)^2} C \sqrt{C^2 + \sigma^2 + \delta^2} \right), \\ & \mathcal{O} \left(\frac{\epsilon^6}{\ell^3 \text{diam}(\Lambda)^2 (\sigma^2 + \delta^2) C \sqrt{C^2 + \sigma^2 + \delta^2}} \right) \end{aligned} \right\}, \tag{30}$$

we have

$$\begin{aligned}
& \frac{1}{J+1} \left(\sum_{k=0}^J \mathbb{E} \left\| \nabla \Phi_{1/2\ell}(t_k) \right\|^2 \right) \\
& \leq \frac{4\widehat{\Delta}_\Phi}{\eta_t(J+1)} + \frac{8\ell \widehat{\Delta}_0}{J+1} + 8\ell \delta \cdot (\text{diam}(\mathcal{T}) + \text{diam}(\Lambda)) + \mathcal{O}(\epsilon^2).
\end{aligned}$$

Thus, we have an iteration complexity of

$$\mathcal{O} \left(\left(\frac{\ell(C^2 + \sigma^2 + \delta^2) \cdot \widehat{\Delta}_\Phi}{\epsilon^4} + \frac{\ell \widehat{\Delta}_0}{\epsilon^2} \right) \cdot \max \left\{ 1, \frac{\ell^2 \text{diam}(\Lambda)^2}{\epsilon^2}, \frac{\ell^2 \text{diam}(\Lambda)^2 (\sigma^2 + \delta^2)}{\epsilon^4} \right\} \right)$$

for recovering an $\mathcal{O}(\sqrt{\epsilon^2 + \delta \ell (\text{diam}(\mathcal{T}) + \text{diam}(\Lambda))})$ -stationary point (cf. Definition 3.11). By simplifying, we obtain the desired iteration bound:

$$\mathcal{O} \left(\left(\frac{\ell^3 (C^2 + \sigma^2 + \delta^2) (\text{diam}(\Lambda))^2 \cdot \widehat{\Delta}_\Phi}{\epsilon^6} + \frac{\ell^3 (\text{diam}(\Lambda))^2 \cdot \widehat{\Delta}_0}{\epsilon^4} \right) \max \left\{ 1, \frac{\sigma^2 + \delta^2}{\epsilon^2} \right\} \right).$$

□

B Experimental Details

B.1 Simulations

Constraint tasks are a central class of real-world RL problems with applications in robotics, autonomous vehicles, and industrial control. To simulate realistic conditions, we use the locomotion tasks from the Safety-Gymnasium suite [19]. Further details, including visualizations of both constrained and unconstrained agent behaviors, are available on their website⁶.

Velocity Cost The velocity cost is defined as:

$$\text{cost} = \text{bool}(\text{vel}_{\text{current}} > \text{vel}_{\text{threshold}}),$$

where the velocity thresholds $\text{vel}_{\text{threshold}}$ for the tested environments are listed in Table 3. As reported by Towers et al. [47], these thresholds are set to 50% of the maximum velocity achieved by each agent after PPO training for 10^7 steps. At each time step, the agent’s instantaneous velocity is computed as:

$$\text{vel}_{\text{current}} = \sqrt{\text{vel}_{x,\text{current}}^2 + \text{vel}_{y,\text{current}}^2}, \quad (31)$$

where $\text{vel}_{x,\text{current}}$ and $\text{vel}_{y,\text{current}}$ are the agent’s instantaneous velocities along the x- and y-axes, respectively, as provided by the simulator.

Environments We use the latest version (v1) of Safety-Gymnasium without modifying the state, action, or reward space. Constraint violations do not alter the agent’s behavior or environment dynamics. As a result, vanilla PPO converges to high reward solutions with large cumulative costs—effectively ignoring the constraints. All actions are normalized to the range $[-1, 1]$. To isolate the effects of our method, we reset the agent to the same initial state after every termination—a common practice in robotics.

Episode Terminations Episodes terminate either when the time limit is reached or when the agent fails (e.g., by falling). Notably, HalfCheetah, Swimmer, and all safe navigation environments have no failure condition; their episodes always end due to the time limit.

Table 3: Environment-specific parameters used in our experiments.

Variable	Safe Navigation	HalfCheetah	Hopper	Swimmer	Walker2d
Threshold c	0.0	3.2096	0.7402	0.2282	2.3415
Failure condition	X	X	✓	X	✓
Time limit (steps)	500	1000	1000	1000	1000
Initial CVaR variable t_{init}	0.0	-1.3	-0.1	-0.0	-0.975

Stochasticity for Risk Management The environment is fully deterministic—identical actions from the same state always yield the same rewards and transitions—making it difficult to evaluate risk-sensitive behavior. To simulate uncertainty without altering the environment’s internal dynamics, we inject zero-mean Gaussian noise (std. 0.05, i.e., 5% of the action range) into *all* agent actions at every step during both training and evaluation. This controlled perturbation introduces stochasticity in action execution, enabling us to assess how well the agent manages risk under uncertain conditions while maintaining consistent environment behavior.

Evaluation Agents are evaluated every 1000 time steps by averaging the undiscounted sum of rewards over 10 episodes. The PPO agent uses the mean action, ensuring consistency in evaluation. Evaluations are entirely separate from training—no data is stored, and no network updates are performed.

⁶https://safety-gymnasium.readthedocs.io/en/latest/environments/safe_velocity.html

B.2 Proximal Policy Optimization – *solver*

We use Proximal Policy Optimization [39] as a *solver* to learn a policy (i.e., line 4 in Algorithm 1). Our method serves as a *wrapper*, described next, that modifies the agent’s raw reward to incorporate risk measures and constraints.

PPO first collects rollouts of state-action-reward sequences using the current policy, storing them as trajectories. Once sufficient data is gathered, it applies minibatch learning, splitting the rollout data into smaller batches and iteratively updating the policy over multiple epochs. During this process, the dual and CVaR variables remain fixed.

Neural Networks The value function and policy are approximated by neural networks, each with two hidden layers of 64 neurons using the tanh activation function. The value network processes states s and outputs a scalar value. The policy network takes states s as input, extracts hidden features, and passes them to a Gaussian distribution with learnable mean and standard deviation parameters. The action a is then sampled from this distribution.

Table 4: PPO hyperparameters used in the experiments.

Hyperparameter	Value
Optimizer	Adam
Learning rate (all networks)	3×10^{-4}
Linear learning rate decay	✓
Adam ϵ	10^{-6}
Adam α	0.99
# rollout steps	2048
# minibatches per rollout	32
# epochs	10
Discount factor γ	0.99
GAE λ	0.95
Entropy coefficient	0.0
Value loss coefficient	0.5
Maximum gradient norm	0.5
Clip parameter	0.2

Hyperparameters We employ Generalized Advantage Estimation (GAE) [38] to estimate advantages in PPO. The hyperparameters used by the PPO agent is provided in Table 4.

B.3 Reward-Based SGD with Risk Constraints – *wrapper*

We follow the same rollout strategy to optimize the dual and CVaR variables. First, the policy is updated using collected rollouts while keeping λ and t fixed. Then, a new rollout is collected with the updated policy and used to update λ and t , while keeping the policy parameters frozen.

B.3.1 Implementation

To balance reward maximization with constraint handling, we frame the problem as standard risk-neutral reward maximization subject to a constraint that regulates violations through the conditional value-at-risk (CVaR) of the constraint quantity.

Let $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ be the reward function and $v : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ a constraint-quantifying function, e.g., velocity function in (31). We want to solve:

$$\sup_{\pi \in \mathcal{P}(\mathcal{S})} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} r(s_{\tau}, \pi(s_{\tau})) \right] \quad \text{s.t.} \quad \text{CVaR}_{\nu^{\pi}}^{\beta}(v(s, a)) \leq c.$$

This constraint can be equivalently written as

$$\text{CVaR}_{\nu^{\pi}}^{\beta}(v(s, a)) \leq c \iff -\text{CVaR}_{\nu^{\pi}}^{\beta}(-v(s, a)) \geq -c,$$

Table 5: Common hyperparameter values used across all environments.

Variable	Value
VaR level β	0.3
# trajectories used to compute gradients of λ and t	8
Initial dual variable λ_{init}	0.0
Step size η_λ	5×10^{-5}
Learning rate decay on η_λ	$\times$
Step size η_t	5×10^{-5}
Learning rate decay on η_t	$\times$

which aligns with the supremal convolution form of the reflected CVaR. Rewriting the constraint:

$$-\text{CVaR}_{\nu^\pi}^\beta(-v(s, a)) \geq -c \iff \sup_{t \in \mathbb{R}} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau \left(t - \frac{1}{\beta} (t + v(s_\tau, \pi(s_\tau)))_+ \right) \right] \geq -\frac{c}{1-\gamma}.$$

Thus, we solve the following constrained problem:

$$\begin{aligned} \sup_{\pi \in \mathcal{P}(\mathcal{S}), t \in \mathbb{R}} \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau r(s_\tau, \pi(s_\tau)) \right] \\ \text{s.t.} \quad \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau \left(t - \frac{1}{\beta} (t + v(s_\tau, \pi(s_\tau)))_+ \right) \right] \geq -\frac{c}{1-\gamma}. \end{aligned}$$

We implement this by modifying the reward at each time step using the Lagrangian:

$$r(s_\tau, \pi(s_\tau)) + \lambda_i \left(c + t - \frac{1}{\beta} (t + v(s_\tau, \pi(s_\tau)))_+ \right), \quad (32)$$

where PPO is used as a black-box solver for the inner maximization over π , while λ and t are updated using single stochastic gradient descent and ascent steps as in Algorithm 1.

B.3.2 Hyperparameters

All hyperparameters used in our algorithm are listed in Tables 3 and 5.

Setting β The parameter β controls which quantile mean is constrained. Higher values (close to 1) enforce more risk-neutral constraints, while lower values focus on rare events. We chose $\beta = 0.3$ to strike a balance: strong enough risk control to avoid constraint violations, while keeping the problem solvable.

Number of Trajectories for Gradient Computation We tested $n = \{2, 4, 8, 16\}$ (in Algorithm 1) and concluded that $n = 8$ offers the best trade-off between runtime and gradient smoothness.

Step Sizes η_λ and η_t Step sizes were extensively tuned on Hopper and Walker2d. We tested values from 10^{-3} to 10^{-7} . The best-performing configuration— $\eta_\lambda = \eta_t = 5 \times 10^{-5}$ —was selected based on the convergence of λ and t within 15M time steps.

Initial Value of t Since t must take values in the negative real range (due to the supremal form), we expect it to converge to the negative value-at-risk. Thus, we initialized t such that, with step size $\eta_t = 5 \times 10^{-5}$, its magnitude could reach the velocity threshold over training.

Initial Value of λ Because λ scales the penalty term added to the reward in (32), we initialized it neutrally with $\lambda = 0.0$ to avoid overly aggressive penalties at the start.

B.4 Computational Resources

All experiments were performed on a computing system powered by an AMD Ryzen processor with 64 cores and 512 GB of RAM. A single NVIDIA RTX A6000 GPU with 48 GB VRAM was used for neural network training.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The claims we make in the abstract and introduction are that we propose a new framework for risk-aware RL with constraints, that exhibits per-stage robustness, admitting an exact equivalent partial Lagrangian relaxation under appropriate constraint qualifications, convergence analysis, and experiments demonstrating such. These are supported by the related works (Section 1), Section 2.3 and Appendix A.1, Section 3.1-3.2, Section 3.3, and Section 4, respectively.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We include a separate paragraph addressing the limitations at the end of the manuscript. In addition, throughout we qualify our results and claims (e.g., when introducing assumptions like Assumption 3.8). We also clearly state the environments on which we run our experiments, emphasizing that our simulations are not exhaustive and serve as a *proof-of-work*, with computational resources outlined in Appendix B.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best

judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We made all efforts to clearly state all assumptions and claims using the appropriate environments to make them obvious. All proofs are provided either in the main body or in the appendix to save space.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: Our results are reproducible and we provide the hyperparameters, training procedure, and the (anonymized) code to facilitate review and reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The data/environment we use is synthetic and open source. We provide our code in an anonymized repository.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: These experiment details are all provided in Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: This note applies indirectly and exclusively to the histogram shown in Figure 1. Our intention is to estimate the probability distribution of a specific variable, aggregated over 100 trials as noted in the figure caption.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We detail this information in the last subsection of Appendix B (in B.4).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We have read the Code of Ethics and believe our work adheres to it.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Our work is mostly theoretical and the experiments are done on synthetic environments, thus we believe our work as minimal societal impact.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Similar to the previous section, our experiments are done in clean synthetic environments thus this section does not seem applicable to our work. We do list limitations, as mentioned before.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite [19] and include the URL as well as the license in the main text.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our work does not include human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our work does not include human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: There was no usage of LLMs in this work.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.