

PROVABLY EFFICIENT CVaR RL IN LOW-RANK MDPs

Yulai Zhao*

Princeton University
yulaiz@princeton.edu

Wenhao Zhan*

Princeton University
wenhao.zhan@princeton.edu

Xiaoyan Hu*

The Chinese University of Hong Kong
xyhu21@cse.cuhk.edu.hk

Ho-fung Leung

Independent Researcher
ho-fung.leung@outlook.com

Farzan Farnia

The Chinese University of Hong Kong
farnia@cse.cuhk.edu.hk

Wen Sun

Cornell University
ws455@cornell.edu

Jason D. Lee

Princeton University
jasonlee@princeton.edu

ABSTRACT

We study risk-sensitive Reinforcement Learning (RL), where we aim to maximize the Conditional Value at Risk (CVaR) with a fixed risk tolerance τ . Prior theoretical work studying risk-sensitive RL focuses on the tabular Markov Decision Processes (MDPs) setting. To extend CVaR RL to settings where state space is large, function approximation must be deployed. We study CVaR RL in low-rank MDPs with nonlinear function approximation. Low-rank MDPs assume the underlying transition kernel admits a low-rank decomposition, but unlike prior linear models, low-rank MDPs do not assume the feature or state-action representation is known. We propose a novel Upper Confidence Bound (UCB) bonus-driven algorithm to carefully balance the interplay between exploration, exploitation, and representation learning in CVaR RL. We prove that our algorithm achieves a sample complexity of $\tilde{O}\left(\frac{H^7 A^2 d^4}{\tau^2 \epsilon^2}\right)$ to yield an ϵ -optimal CVaR, where H is the length of each episode, A is the capacity of action space, and d is the dimension of representations. Computational-wise, we design a novel discretized Least-Squares Value Iteration (LSVI) algorithm for the CVaR objective as the planning oracle and show that we can find the near-optimal policy in a polynomial running time with a Maximum Likelihood Estimation oracle. To our knowledge, this is the first provably efficient CVaR RL algorithm in low-rank MDPs.

1 INTRODUCTION

As a widely adopted framework to deal with challenging sequential decision-making problems, Reinforcement learning (RL) (Sutton & Barto, 2018) has demonstrated many empirical successes, e.g., popular strategy games (Silver et al., 2016; 2017; Vinyals et al., 2019). However, the classical RL framework often prioritizes the maximization of the expected cumulative rewards solely, which renders it not suitable for many real-world applications that face possible failure or safety concerns. In high-stake scenarios such as autonomous driving (Isele et al., 2018; Wen et al., 2020), finance (Davis & Lleo, 2008; Wang et al., 2021; Filippi et al., 2020) and healthcare (Ernst et al., 2006; Coronato et al., 2020), optimizing only the expected cumulative rewards can produce policies that underestimate the risks associated with rare but catastrophic events. To mitigate this limitation, risk-sensitive decision-making has recently grown more popular (Hu & Leung, 2023).

*Equal contribution. Listing order is random.

To this end, risk measures like Conditional Value-at-Risk (CVaR) (Rockafellar et al., 2000) have been integrated into the RL-based systems as a performance criterion, yielding a more balanced approach that encourages policies to avoid high-risk outcomes. As a popular tool for managing risk, the CVaR metric is widely used in robust RL (Hiraoka et al., 2019), distributional RL (Achab et al., 2022), and portfolio optimization (Krokhmal et al., 2002). CVaR quantifies the expected return in the worst-case scenarios.¹ For a random variable X with distribution P and a confidence level $\tau \in (0, 1)$, the CVaR at confidence level τ is defined as

$$\text{CVaR}_\tau(X) = \sup_{b \in \mathbb{R}} [b - \tau \mathbb{E}_{X \sim P}(b - X)^+]$$

where $x^+ = \max(x, 0)$. In this paper, we consider the random variable X as the cumulative reward of a specific policy where randomness comes from the MDP transitions and the policy itself. We aim to maximize this objective to capture the average worst values of the returns distribution at a certain risk tolerance level τ .

The most related works to this paper are Bastani et al. (2022) and Wang et al. (2023). Specifically, Bastani et al. (2022) proved the first regret bounds for risk-sensitive RL with CVaR metric (CVaR RL), and Wang et al. (2023) improved the results to achieve the minimax optimal regret. However, their algorithms are restricted to the tabular setting, so they are inefficient when the state space is large. To overcome such limitation, we introduce function approximation to the MDP structure and consider **low-rank MDPs** (Jiang et al., 2017) in this paper.

In low-rank MDPs, the key idea is to exploit the (high-dimensional) structure in the model transitions of the MDP by representing them with a lower-dimensional structure. This is often implemented by assuming the MDP transition kernels admit a low-rank factorization, i.e., there exist two *unknown* mappings $\psi(s')$, $\phi(s, a)$, such that the probability of transitioning to the next state s' under the current state and action (s, a) is $P(s'|s, a) = \psi(s')^\top \phi(s, a)$. Low-rank MDPs generalize popular RL models including tabular MDPs, linear MDPs (Jin et al., 2020) and block MDPs (Du et al., 2019; Efroni et al., 2021), and are widely applied in practice (Zhang et al., 2020; 2021; Sodhani et al., 2021; 2022).

Unlike linear MDPs, since the underlying ψ and ϕ are unknown, we need to carefully balance representation learning, exploration, and worst-case failure stakes in low-rank MDPs. The application of non-linear function approximation further complicates the problem, for controlling model errors state and action-wise becomes much harder (Wang et al., 2023). Moreover, even within a given MDP, planning in CVaR RL is not a straightforward task as the CVaR metric exhibits a non-linear structure² (in stark contrast to standard risk-neutral RL), which results in extra computational overhead. All of these challenges call for a more comprehensive algorithm design for the risk-sensitive RL in low-rank MDPs.

In this paper, we aim to fill a gap in the existing body of knowledge by exploring the interplay between risk-averse RL and low-rank MDPs. We summarize our contributions as follows.

Contributions. First, we design an oracle-efficient representation learning algorithm called *ELA* (*REpresentation Learning for CVAR*), which optimizes the CVaR metric in low-rank MDPs. *ELA* leverages an MLE oracle to learn the model dynamics while simultaneously constructing Upper Confidence Bound (UCB)-type bonuses to encourage exploration of the unknown environment. We provide a comprehensive theoretical analysis of the algorithm, demonstrating that *ELA* would provide an ϵ -optimal CVaR with $\tilde{O}(1/\epsilon^2)$ samples. To the best of our knowledge, *ELA* is the first provably sample-efficient algorithm for CVaR RL in low-rank MDPs.

Second, to improve the computational complexity of *ELA* planning, we introduce a computationally efficient planning oracle, which, when combined with *ELA*, leads to the *ELLA* (*REpresentation Learning with LSVI for CVAR*) algorithm. This algorithm leverages least-squares value iteration with discretized rewards to find near-optimal policies via optimistic simulations within the learned model. Importantly, the computational cost solely depends on the dimension of representations rather than the state space size. We show that *ELLA* requires a polynomial running time in addition to polynomial calls to the MLE oracle.

¹Standard CVaR definition considers average worst-case *loss*, thus a lower value is more desirable. However, in this work, we aim to obtain a higher CVaR in the RL context as we are maximizing rewards.

²In RL, planning refers to determining the optimal policy and optimal value function within a given MDP.

2 RELATED WORK

Low-rank MDPs. Theoretical benefits of low-rank structure in MDPs have been broadly explored in various works (Jiang et al., 2017; Sun et al., 2019; Du et al., 2021; Sekhari et al., 2021; Huang et al., 2023). In a contextual decision process (a generic RL model with rich observations and function approximation) with a low Bellman rank, OLIVE (Jiang et al., 2017) yielded a near-optimal policy with polynomial samples. Additionally, Sun et al. (2019) introduced provably efficient algorithms based on a structural parameter called the witness rank, demonstrating that the witness rank is never larger than the Bellman rank (Jiang et al., 2017). Despite their provable efficiency, these algorithms lack computational efficiency.

Leveraging Maximum Likelihood Estimation (MLE) as its computation oracle, Flambe (Agarwal et al., 2020) proposed the first computationally efficient algorithm that was also provably efficient in low-rank MDPs. Following a similar setup, Rep-UCB (Uehara et al., 2022) improves the sample complexity dependencies of Flambe. The key to the improvements is a careful tradeoff between exploration and exploitation by combining the reward signal and exploration bonus.

CVaR RL. There is a long line of works studying (static) CVaR in RL, which refers to the CVaR of *accumulative reward* beyond a certain risk threshold (Chow & Ghavamzadeh, 2014; Chow et al., 2015; Tamar et al., 2015; Bastani et al., 2022; Wang et al., 2023). For tabular MDPs, much work has been done on value-based algorithms (Chow et al., 2015; Stanko & Macek, 2019). However, these results often require the planner to know the model transitions of the MDPs, which is generally infeasible. Recently, there has been a growing interest in the more general setting where the unknown model transitions are learned through online interactions (Yu et al., 2018; Bastani et al., 2022; Wang et al., 2023). However, their results are restricted to the tabular setting and cannot be combined with function approximation, where the state space is often enormous.

Prior works have also studied risk-sensitive RL in the context of other risk measures. Specifically, Du et al. (2022) proposed Iterated CVaR RL, also known as dynamic CVaR, and Chen et al. (2023) demonstrates regret guarantees for Iterated CVaR RL with general function approximations. However, iterated CVaR is intrinsically different from the static CVaR in our setting. Iterated CVaR quantifies the worst τ -percent performance **at each step** of decision-making. Such a definition allows the agent to control the risk throughout the decision process tightly, whereas our setting aims to maximize the CVaR of the **total** reward. Therefore, their algorithm designs and analysis techniques do not apply to our tasks.

Our work also focuses on the static CVaR in RL. Specifically, we present the first sample-efficient algorithm for optimizing the static CVaR metric that carefully balances the interplay between risk-averse RL and low-rank MDP structure. Furthermore, we design a computationally efficient planning oracle that makes our algorithm only require polynomial running time with an MLE oracle.

3 PRELIMINARIES

Notations. We will frequently use $[H]$ to denote the set $\{1, \dots, H\}$. Denote $\Delta(\mathcal{S})$ as the distribution over space \mathcal{S} . Let $U(\mathcal{A})$ be the uniform distribution over the action space. In this work, we use the standard $O(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$ notations to hide universal constant factors and use $\tilde{O}(\cdot)$ to hide logarithmic factors. Please see Table 1 for a comprehensive list of notations used in this work.

3.1 LOW-RANK EPISODIC MDP

We consider an episodic MDP \mathcal{M} with episode length $H \in \mathcal{N}$, state space \mathcal{S} , and a finite action space \mathcal{A} . At each episode $k \in [K]$, a trajectory $\tau = (s_1, a_1, s_2, a_2, \dots, s_H, a_H)$ is generated by an agent, where (a) $s_1 \in \mathcal{S}$ is a *fixed* starting state,³ (b) at step h , the agent chooses action according to a history-dependent policy $a_h \sim \pi(\cdot | \tau_{h-1}, s_h)$ where $\tau_{h-1} := (s_1, a_1, \dots, a_{h-1})$ denotes the history and (c) the model transits to the next state $s_{h+1} \sim P_h^*(\cdot | s_h, a_h)$. The agent repeats these steps till the end of the episode. For each time step, operators $P_h^* : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ denote the (non-stationary) transition dynamics, whereas $r_h : \mathcal{S} \times \mathcal{A} \rightarrow \Delta([0, 1])$ is the (immediate) reward distribution the

³Note that any H -length episodic MDP with a stochastic initial state is equivalent to an $(H + 1)$ -length MDP with a dummy initial state s_0 .

agent could receive from deploying a certain action at a specific state. We use Π to denote the class of all history-dependent policies. Below, we proceed to the low-rank MDP definition.

Definition 3.1 (Low-rank episodic MDP). The transition kernel $\mathcal{P}^* = \{P_h^* : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})\}_{h \in [H]}$ admits a low-rank decomposition with rank d if there exist two embedding functions $\phi^* := \{\phi_h^* : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}^d\}_{h \in [H]}$ and $\psi^* := \{\psi_h^* : \mathcal{S} \mapsto \mathbb{R}^d\}_{h \in [H]}$ such that

$$P_h^*(s'|s, a) = \langle \psi_h^*(s'), \phi_h^*(s, a) \rangle \quad (1)$$

where $\|\phi_h^*(s, a)\|_2 \leq 1$ for all $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$, and for any function $g : \mathcal{S} \mapsto [0, 1]$ and $h \in [H]$, it holds that $\|\int_{s \in \mathcal{S}} \psi_h^*(s)g(s)ds\|_2 \leq \sqrt{d}$. Low-rank episodic MDP admits such a decomposition of \mathcal{P}^* .

We study the function approximation setting where the state spaces \mathcal{S} can be enormous and even infinite. To generalize across states, assume access to two embedding classes $\Psi \subset \{\mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d\}$ and $\Phi \subset \{\mathcal{S} \rightarrow \mathbb{R}^d\}$, which are used to identify the true embeddings (ψ^*, ϕ^*) . Formally, we need the following realizability assumption, which is proven essential for obtaining performance guarantees independent of the state space size in low-rank MDPs (Agarwal et al., 2020).

Assumption 3.2. There exists a model class $\mathcal{F} = (\Psi, \Phi)$ such that $\psi_h^* \in \Psi, \phi_h^* \in \Phi, \forall h \in [H]$.

We expect to obtain sample complexity that scales logarithmically with the (finite) cardinality of the model class \mathcal{F} . Extensions to the infinite function classes with complexity measures are possible. See more discussions in (Agarwal et al., 2020).

3.2 RISK-SENSITIVE RL AND AUGMENTED MDP

We study risk-sensitive RL with the CVaR metric. Throughout the paper, let $\tau \in (0, 1]$ be a fixed risk tolerance. First, we recall the classical definition: for a random variable $X \in [0, 1]$ is from distribution P , the *conditional-value-at-risk* (CVaR) corresponding to the risk tolerance τ is defined as

$$\text{CVaR}_\tau(X) := \sup_{c \in [0, 1]} \{c - \tau^{-1} \cdot \mathbb{E}_{X \sim P}[(c - X)^+]\} \quad (2)$$

where $(x)^+ := \max(x, 0)$ for any $x \in \mathbb{R}$. Interestingly, the supremum in the expression is achieved when c is set as the τ -th percentile (unknown before planning), also known as *value-at-risk* (VaR), i.e., $x_\tau = \inf\{x \in \mathbb{R} : P(X \leq x) \geq \tau\}$.

In risk-sensitive RL, X represents the stochastic cumulative reward accumulated over H successive actions and state transitions: $X = \sum_{h=1}^H r_h$. This multi-step nature of risk-sensitive RL brings the dynamic planning structure to the setting and makes it difficult. We may make an intuitive insight of c by considering it as a *initial budget*, which affects the agent’s action selection and needs to be carefully managed during the planning. Particularly, after receiving a random reward r_h at each timestep $h \in [H]$, the learner deducts it from the current budget, i.e., $c_{h+1} = c_h - r_h$ where c_h is the remaining budget and $c_1 = c$. The main task of risk-sensitive RL is to interplay carefully between policy planning, exploration, and budget control. Inspired by this observation, we incorporate the available budget as an additional state variable that could impact the agent’s choice of action. Formal descriptions of the augmented MDP are introduced below.

Augmented MDP. We introduce the augmented MDP framework (Bauerle & Ott, 2011) to study CVaR RL, which augments the state space \mathcal{S} in classic episodic MDP to $\mathcal{S}_{\text{Aug}} = \mathcal{S} \times [0, H]$ that includes the budget variable as an additional state. The transition kernel of the state and the immediate reward distribution are the same as the original MDP \mathcal{M} . Inspired by the observation above, we consider policies defined on the augmented state space $\Pi_{\text{Aug}} := \{\pi : \mathcal{S}_{\text{Aug}} \rightarrow \Delta(\mathcal{A})\}$.

In the augmented MDP, the agent rolls out an augmented policy $\pi \in \Pi_{\text{Aug}}$ with initial budget $c_1 \in [0, H]$ as follows. At the beginning of an episode, the agent observes the augmented state (s_1, c_1) , selects action $a_1 \sim \pi(\cdot | s_1, c_1)$, receives reward $r_1 \sim r_1(s_1, a_1)$, and transits to $s_2 \sim P_1^*(\cdot | s_1, a_1)$. Most importantly, the available budget is also updated: $c_2 = c_1 - r_1$. The agent then chooses an action based on the (s_2, c_2) pair. The procedure repeats for H times until the end of the episode. For any $\pi \in \Pi_{\text{Aug}}$, the (augmented) Q -function is defined as

$$Q_{h, \mathcal{P}^*}^\pi(s, c, a) := \mathbb{E}_{\pi, \mathcal{P}^*} \left[\left(c_h - \sum_{t=h}^H r_t(s_t, a_t) \right)^+ \middle| s_h = s, c_h = c, a_h = a \right].$$

for any $(h, s, c, a) \in [H] \times \mathcal{S} \times [0, H] \times \mathcal{A}$ and the (augmented) value function is defined as

$$V_{h, \mathcal{P}^*}^\pi(s, c) := \mathbb{E}_{a \sim \pi_h(\cdot | s, c)} [Q_{h, \mathcal{P}^*}^\pi(s, c, a)] \quad (3)$$

for any $(h, s, c) \in [H] \times \mathcal{S} \times [0, H]$. The Bellman equation is given by

$$Q_{h, \mathcal{P}^*}^\pi(s, c, a) = \mathbb{E}_{s' \sim P_h^*(\cdot | s, a), r \sim r_h(s, a)} V_{h+1}^\pi(s', c - r).$$

Goal metric. In this paper, we aim to find the optimal history-dependent policy to maximize the CVaR objective, i.e.,

$$\text{CVaR}_\tau^* := \max_{\pi \in \Pi} \text{CVaR}_\tau(R(\pi)) = \sup_{c \in [0, H]} \left\{ c - \tau^{-1} \cdot \min_{\pi \in \Pi} \mathbb{E}[(c - R(\pi))^+] \right\}$$

where $R(\pi)$ is the random cumulative reward of policy π in \mathcal{M} . Nevertheless, it is known that $\min_{\pi \in \Pi} \{ \tau^{-1} \cdot \mathbb{E}[(c - R(\pi))^+] \}$ can be attained by an augmented policy $\pi^* \in \Pi_{\text{Aug}}$ with initial budget c (Wang et al., 2023). Thus we can indeed focus on searching within $\pi \in \Pi_{\text{Aug}}$:

$$\text{CVaR}_\tau^* = \max_{c \in [0, H]} \left\{ c - \tau^{-1} \cdot \min_{\pi \in \Pi_{\text{Aug}}} V_{1, \mathcal{P}^*}^\pi(s_1, c) \right\} := \text{CVaR}_\tau(R(\pi^*, c^*)) \quad (4)$$

where $\pi^* \in \Pi_{\text{Aug}}$ is the optimal augmented policy and $c^* \in [0, H]$ is the optimal initial budget, and we overload $R(\pi, c)$ to denote the stochastic cumulative reward when the agent rolls out the augmented policy π with initial budget c in augmented MDP.

4 ALGORITHM

In this section, we present the *ELA* algorithm for risk-sensitive RL in low-rank MDPs. The pseudocode is listed in Algorithm 1. The algorithm is iterative in nature, where the k -th episode proceeds in three folds: (1) we collect new data by rolling in with the exploration policy π^{k-1} starting from the initial budget c^{k-1} . (2) Then, all transitions collected so far are used in two aspects. First, we pass all transition tuples to the MLE oracle (Line 11). The MLE oracle returns embedding functions $(\hat{\psi}_h, \hat{\phi}_h)$ for each h , which determine the model. Second, we compute the exploration bonus using the latest learned representation $\hat{\phi}$. (3) The algorithm performs VI on the learned model with the bonus-enhanced reward signal to obtain the exploration policy-budget pair (π^k, c^k) we use in the next iteration. After K iterations, we output the current model and all policy-budget pairs. Next, we illustrate more details about these steps.

4.1 DATA COLLECTION

During training, we collect two (disjoint) sets of transition tuples to compute the bonus terms and estimate the transition kernels. These two datasets are different in their (marginalized) distributions and facilitate the regret analysis in Section 4.3. Next, we clarify how data are collected and added to the two sets.

To collect a new transition tuple for each dataset \mathcal{D}_h and $\tilde{\mathcal{D}}_h$ ($\forall h \in [H - 1]$), at the k -th iteration, the algorithm roll outs policy π^{k-1} with initial budget c^{k-1} to obtain two trajectories. The difference occurs at the $(h - 1)$ -th timestep. Particularly, to obtain $(s_h, a_h, \tilde{s}_{h+1})$ for \mathcal{D}_h , the algorithm *keeps on* to execute policy π^{k-1} for one more step and observes state s_h . Then, a uniform action $a_h \sim U(\mathcal{A})$ is taken to receive the next-state $\tilde{s}_{h+1} \sim P_h^*(\cdot | s_h, a_h)$. However, to obtain $(\tilde{s}_h, \tilde{a}_h, s'_{h+1})$ for $\tilde{\mathcal{D}}_h$, the algorithm takes *two consecutive* uniform actions at timesteps $h - 1$ and h , i.e., $a_{h-1} \sim U(\mathcal{A})$, $\tilde{s}_h \sim P_{h-1}^*(\cdot | s_{h-1}, a_{h-1})$, $\tilde{a}_h \sim U(\mathcal{A})$, and receives state s'_{h+1} . Intuitively, dataset \mathcal{D}_h reveals the behavior of the exploration policy-budget pair (π^k, c^k) up to the h -th timestep, which facilitates the design of bonus terms (cf. Lemma C.2). Meanwhile, dataset $\tilde{\mathcal{D}}_h$ enhances the exploration and leads to improved estimates of transition kernels (cf. Lemma E.3).

At first sight, the data collection procedure requires $2(H - 1)$ trajectories per iteration (i.e., one trajectory for \mathcal{D}_h and one trajectory for $\tilde{\mathcal{D}}_h$). To save sample (trajectory) complexity, we collect

transition tuples for \mathcal{D}_h and $\tilde{\mathcal{D}}_h$ at the $h, (h-1)$ steps in one go, respectively (Lines 8-10). Therefore, the algorithm only requires H trajectories per iteration. For both sets, new transition tuples are concatenated with the existing data to perform representation learning, i.e., learning a factorization and a representation by MLE (Line 11).

4.2 REPRESENTATION LEARNING AND BONUS-DRIVEN VALUE ITERATION

MLE oracle. In the function approximation setting, the agent must estimate the model structure as accurately as possible. Collected transition tuples are used to compute model transitions through the MLE oracle. As a general approach used to estimate the parameters of a probability model, MLE is also gaining more focus in low-rank MDPs (Agarwal et al., 2020; Uehara et al., 2022).

Value Iteration. Based on the learned model, the algorithm runs Value-Iteration (VI) with the exploration bonus term. In risk-sensitive RL, for any $\pi \in \Pi_{Aug}$ and $(h, s, c) \in [H] \times \mathcal{S} \times [0, H]$, we define value function enhanced with exploration bonus $b_h : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ as

$$V_{h,\mathcal{P},b}^\pi(s, c) := \mathbb{E}_{\pi, \mathcal{P}} \left[\left(c_h - \sum_{h'=h}^H r_{h'}(s_{h'}, a_{h'}) \right)^+ - \sum_{h'=h}^H b_{h'}(s_{h'}, a_{h'}) \middle| s_h = s, c_h = c \right] \quad (5)$$

where we deduct the exploration bonus term because the agent desires to *minimize* the value function (4) in risk-sensitive RL. Such a value function is used to perform VI and update policy on the learned model $\hat{\mathcal{P}}$, since the learner has no prior knowledge of the *real* model transitions. Therefore, obtaining an accurate estimation of the model determines the quality of the output policy.

In Algorithm 1, we assume the learner has access to exact VI (Line 15). However, we remark that this step is not computationally efficient due to the continuity of c and potentially large state space \mathcal{S} . To overcome such a computational barrier arising from the nature of controlling risk while planning, in Section 5, we provide a computationally efficient planning oracle that performs LSVI with discretized reward function (with sufficiently high precision). We rigorously prove that we only need polynomial running time with the MLE oracle to output a near-optimal CVaR.

4.3 MAIN RESULTS

In this subsection, we are ready to demonstrate our theoretical guarantees that characterize the regret and sample complexity bounds.

Theorem 4.1. Fix $\delta \in (0, 1)$. Set the parameters in Algorithm 1 as:

$$\alpha^k = O \left(\sqrt{H^2(|\mathcal{A}| + d^2)} \log \left(\frac{|\mathcal{F}|Hk}{\delta} \right) \right), \quad \lambda^k = O \left(d \log \left(\frac{|\mathcal{F}|Hk}{\delta} \right) \right).$$

We have two equivalent interpretations of the theoretical results. In terms of PAC bound, with probability at least $1 - \delta$, the regret is bounded by

$$\sum_{k=1}^K \text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\pi^k, c^k)) = \tilde{O} \left(\tau^{-1} H^3 A d^2 \sqrt{K} \cdot \sqrt{\log(|\mathcal{F}|/\delta)} \right).$$

Alternatively, we can interpret in terms of sample complexity: w.p. at least $1 - \delta$, to present an ϵ -optimal policy and budget pair s.t. $\text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\hat{\pi}, \hat{c})) \leq \epsilon$. The total number of trajectories required is upper bounded by

$$\tilde{O} \left(\frac{H^7 A^2 d^4 \log(|\mathcal{F}|/\delta)}{\tau^2 \epsilon^2} \right).$$

Proof. The proof can be found in Appendix C. □

In Theorem 4.1, we present the first regret/sample complexity bounds for CVaR RL with function approximation, in which exploring the unknown action/space spaces posits extra difficulty. We explicitly characterize the number of samples required to output an ϵ -optimal CVaR in terms of the length of the episode H , the action space size A , the representation dimension d , and confidence level τ . The theorem has no explicit dependence on state space \mathcal{S} , proving nice guarantees even in

Algorithm 1 *ELA*

Require: Risk tolerance $\tau \in (0, 1]$, number of iterations K , parameters $\{\lambda^k\}_{k \in [K]}$ and $\{\alpha^k\}_{k \in [K]}$, models $\mathcal{F} = \{\Psi, \Phi\}$, failure probability $\delta \in (0, 1)$.

- 1: Set datasets $\mathcal{D}_h, \tilde{\mathcal{D}}_h \leftarrow \emptyset$ for each $h \in [H - 1]$.
- 2: Initialize the exploration policy $\pi^0 \leftarrow \{\pi_h^0(s, c) = U(\mathcal{A}), \text{ for any } (s, c) \in \mathcal{S} \times [0, H]\}_{h \in [H]}$.
- 3: Initialize the budget $c^0 \leftarrow 1$.
- 4: **for** iteration $k = 1, \dots, K$ **do**
- 5: Collect a tuple $(\tilde{s}_1, \tilde{a}_1, s'_1)$ by taking $\tilde{a}_1 \sim U(\mathcal{A}), s'_1 \sim P_1^*(\cdot | \tilde{s}_1, \tilde{a}_1)$.
- 6: Update $\tilde{\mathcal{D}}_1 \leftarrow \tilde{\mathcal{D}}_1 \cup \{(\tilde{s}_1, \tilde{a}_1, s'_1)\}$.
- 7: **for** $h = 1, \dots, H - 1$ **do**
- 8: Collect two transition tuples $(s_h, a_h, \tilde{s}_{h+1})$ and $(\tilde{s}_{h+1}, \tilde{a}_{h+1}, s'_{h+2})$ by first rolling out π^{k-1} starting from (s_1, c^{k-1}) into state s_h , taking $a_h \sim U(\mathcal{A})$, and receiving $\tilde{s}_{h+1} \sim P_h^*(\cdot | s_h, a_h)$, then taking $\tilde{a}_{h+1} \sim U(\mathcal{A})$ and receiving $s'_{h+2} \sim P_{h+1}^*(\cdot | \tilde{s}_{h+1}, \tilde{a}_{h+1})$.
- 9: Update $\mathcal{D}_h \leftarrow \mathcal{D}_h \cup \{(s_h, a_h, \tilde{s}_{h+1})\}$.
- 10: Update $\tilde{\mathcal{D}}_{h+1} \leftarrow \tilde{\mathcal{D}}_{h+1} \cup \{(\tilde{s}_{h+1}, \tilde{a}_{h+1}, s'_{h+2})\}$ if $h \leq H - 2$.
- 11: Learn representations via MLE

$$\hat{P}_h := (\hat{\psi}_h, \hat{\phi}_h) \leftarrow \arg \max_{(\psi, \phi) \in \mathcal{F}} \sum_{(s_h, a_h, s_{h+1}) \in \{\mathcal{D}_h + \tilde{\mathcal{D}}_h\}} \log \langle \psi(s_{h+1}), \phi(s_h, a_h) \rangle$$

- 12: Update empirical covariance matrix $\hat{\Sigma}_h = \sum_{(s, a) \in \mathcal{D}_h} \hat{\phi}_h(s, a) \hat{\phi}_h(s, a)^\top + \lambda^k I_d$.
- 13: Set the exploration bonus:

$$\hat{b}_h(s, a) \leftarrow \begin{cases} \min \left(\alpha^k \sqrt{\hat{\phi}_h(s, a) \hat{\Sigma}_h^{-1} \hat{\phi}_h(s, a)^\top}, 2 \right) & h \leq H - 2 \\ 0 & h = H - 1 \end{cases}$$

- 14: **end for**
 - 15: Run Value-Iteration (VI) and obtain $c^k \leftarrow \arg \max_{c \in [0, H]} \left\{ c - \tau^{-1} \min_{\pi} V_{1, \hat{P}, \hat{b}}^\pi(s_1, c) \right\}$.
 - 16: Set $\pi^k \leftarrow \arg \min_{\pi} V_{1, \hat{P}, \hat{b}}^\pi(s_1, c^k)$.
 - 17: **end for**
- Ensure:** Uniformly sample k from $[K]$, return $(\hat{\pi}, \hat{c}) = (\pi^k, c^k)$.

the infinite-state setting. As far as we are concerned, the only existing theoretical guarantees in the field of risk-averse CVaR RL were provided in the tabular setting (Bastani et al., 2022; Wang et al., 2023), where representations are known. Moreover, these results explicitly depend on $|\mathcal{S}|$, thus can not apply to the function approximation setting with enormous state space. Given the above, our work accomplishes a great leap by incorporating exploration in the comprehensive function approximation setting, which evidently better aligns with real-world applications than the tabular and/or linear MDP settings.

As for the sample complexity, the dependencies on A and d match the same rates as the analysis of risk-neutral RL in low-rank MDP (Uehara et al., 2022), showing our algorithm overcomes the extra hardness of balancing exploration and budget control even with a large action space. Our sample complexity matches the dependency on τ with the rates in the CVaR-UCBVI algorithm with the Hoeffding bonus (Wang et al., 2023) and UCB algorithm (Bastani et al., 2022). On the other hand, the sample complexity enjoys an $\Omega\left(\frac{1}{\tau \epsilon^2}\right)$ lower bound (Wang et al., 2023, Theorem 3.1). Tightening the dependency on τ is left as future work.

Our theoretical guarantees have a slightly worse dependency on H^7 while Rep-UCB (Uehara et al., 2022) scales as H^5 , since we do not assume the trajectory reward is normalized to 1. If we adopt this reward normalization, our PAC bound can be improved to $\tilde{O}(H^5)$, and thus matching the rate. In addition, we point out that it might be possible to reduce another H^2 dependency by utilizing a Bernstein-type bonus in the algorithm instead of the Hoeffding-type bonus, which is left for future work.

5 PLANNING ORACLE: CVaR-LSVI

In Algorithm 1, we need to calculate $c^k \leftarrow \arg \max_{c \in [0, H]} \left\{ c - \tau^{-1} \min_{\pi} V_{1, \hat{\mathcal{P}}, \hat{b}}^{\pi}(s_1, c) \right\}$ (Line 15), which is not naively computationally efficient since the objective $c - \tau^{-1} \min_{\pi} V_{1, \hat{\mathcal{P}}, \hat{b}}^{\pi}(s_1, c)$ is not concave. In this section, we introduce a feasible planning oracle for this step and provide the corresponding theoretical guarantees. For simplicity, we assume the reward distribution r is discrete and $r_h(s, a)$ only takes value in iv where $v > 0$ and $0 \leq i \leq \lceil 1/v \rceil$. For a continuous r , we can discretize it with sufficiently high precision so that all the analysis still applies. The details are deferred to Appendix B.

Since the supremum of the CVaR objective (2) is attained at τ -th quantile of the cumulative reward distribution, which is also discrete when the reward r_h is discrete, we can only search c^k within the grid in Line 15 of Algorithm 1:

$$c^k \leftarrow v \cdot \arg \max_{0 \leq i \leq \lceil H/v \rceil} \left\{ iv - \tau^{-1} \min_{\pi \in \Pi_{Aug}} V_{1, \hat{\mathcal{P}}, \hat{b}}^{\pi}(s_1, iv) \right\}.$$

Nevertheless, if we run standard value iteration to calculate $\min_{\pi \in \Pi_{Aug}} V_{1, \hat{\mathcal{P}}, \hat{b}}^{\pi}(s_1, iv)$, our sample complexity will scale with the size of the state space $|\mathcal{S}|$, which can be infinite in low-rank MDPs. To circumvent such dependency on $|\mathcal{S}|$, we introduce a novel LSVI-UCB algorithm for the CVaR objective in this subsection, called *CVaR-LSVI*. In the following discussion, we fix i_1 where $0 \leq i_1 \leq \lceil H/v \rceil$ and aim at calculating $\min_{\pi \in \Pi_{Aug}} V_{1, \hat{\mathcal{P}}, \hat{b}}^{\pi}(s_1, i_1 v)$. We will drop the subscript $\hat{\mathcal{P}}$ and \hat{b} when it is clear from the context. We use V_h^* to denote $\min_{\pi \in \Pi_{Aug}} V_{h, \hat{\mathcal{P}}, \hat{b}}^{\pi}$ and $(\hat{\mathcal{P}}, r)$ to denote the MDP model whose transition is $\hat{\mathcal{P}}$ and reward distribution is r .

First, recall that the discrete reward distribution $r_h(s, a)$ only takes values of iv where $0 \leq i \leq \lceil 1/v \rceil$. This means that it is always linear with respect to a $(\lceil 1/v \rceil + 1)$ -dimension vector:

$$r_h(iv|s, a) = \langle \phi_{h,r}(s, a), \psi_{h,r}(iv) \rangle,$$

where $(\phi_{h,r}(s, a))_i = r_h(iv|s, a)$ and $\psi_{h,r}(iv) = e_i$ for all $0 \leq i \leq \lceil 1/v \rceil$ and $s \in \mathcal{S}, a \in \mathcal{A}, h \in [H]$. Since $\hat{P}_h(s'|s, a) = \langle \hat{\phi}_h(s, a), \hat{\psi}_h(s') \rangle$, this implies that for all $s, s' \in \mathcal{S}, a \in \mathcal{A}, h \in [H], 0 \leq i \leq \lceil 1/v \rceil$, we have

$$\hat{P}_h(s'|s, a) r_h(iv|s, a) = \langle \bar{\phi}_h(s, a), \bar{\psi}_h(s', iv) \rangle,$$

where $\bar{\phi}_h(s, a) = \hat{\phi}_h(s, a) \otimes \phi_{h,r}(s, a)$ and $\bar{\psi}_h(s', iv) = \hat{\psi}_h(s', a) \otimes \psi_{h,r}(iv)$.

The linearity of transition and reward implies that we can utilize LSVI to compute $Q_h^{\pi}(s, c, a)$. More specifically, we propose an iterative algorithm consisting of the following steps:

- **Step 1: Ridge Regression.** In the t -th iteration, denote the trajectories collected before t -th iteration by $\{(s_h^j, a_h^j, r_h^j)_{h=1}^H\}_{j=1}^{t-1}$. Let $V_{H+1}^t(s, iv) = iv$ for all $s \in \mathcal{S}$ and $0 \leq i \leq \lceil H/v \rceil$. From $h = H$ to $h = 1$, for each $0 \leq i \leq \lceil H/v \rceil$, we first compute:

$$w_h^t(iv) \leftarrow (\Lambda_h^t)^{-1} \sum_{j=1}^{t-1} \bar{\phi}_h(s_h^j, a_h^j) \cdot \left[V_{h+1}^t(s_{h+1}^j, iv - r_h^j) \right],$$

where $\Lambda_h^t = \lambda I + \sum_{j=1}^{t-1} \bar{\phi}_h(s_h^j, a_h^j) (\bar{\phi}_h(s_h^j, a_h^j))^{\top}$.

Then for any $j \in [t-1], a \in \mathcal{A}$ and $0 \leq i \leq \lceil H/v \rceil$, we can estimate the value function $V_h^t(s_h^j, iv)$ as:

$$Q_h^t(s_h^j, iv, a) = \text{Clip}_{[-H, H]} \left(-\hat{b}_h(s_h^j, a) + (\bar{\phi}_h(s_h^j, a))^{\top} w_h^t(iv) - \beta \|\bar{\phi}_h(s_h^j, a)\|_{(\Lambda_h^t)^{-1}} \right),$$

$$V_h^t(s_h^j, iv) = \min_{a \in \mathcal{A}} Q_h^t(s_h^j, iv, a).$$

Note that although we do not compute $Q_h^t(s, iv, a)$ for all $s \in \mathcal{S}$ (which will incur computation cost scaling with $|\mathcal{S}|$), they can be implicitly expressed via $w_h^t(iv)$, i.e., for all $s \in \mathcal{S}, a \in \mathcal{A}, 0 \leq i \leq \lceil H/v \rceil$, we know

$$Q_h^t(s, iv, a) = \text{Clip}_{[-H, H]} \left(-\hat{b}_h(s, a) + (\bar{\phi}_h(s, a))^{\top} w_h^t(iv) - \beta \|\bar{\phi}_h(s, a)\|_{(\Lambda_h^t)^{-1}} \right). \quad (6)$$

- **Step 2: Sample Collection.** In the t -th iteration, simulate the greedy policy $\tilde{\pi}^t$ (w.r.t. the estimated Q function $Q_h^t(s, iv, a)$) with the initial budget $c_1 = i_1 v$ in the MDP model $(\hat{\mathcal{P}}, r)$ and collect a trajectory $(s_h^t, a_h^t, r_h^t)_{h=1}^H$. Then, go back to the first step.
- **Step 3: Policy Evaluation.** After repeating the above two steps for T_1 iterations, we simulate each policy $\tilde{\pi}^t$ with initial budget $c_1 = i_1 v$ in $(\hat{\mathcal{P}}, r)$ for T_2 episodes. Suppose the collected trajectories are $\left\{ \left(s_h^{t,j}, a_h^{t,j}, r_h^{t,j} \right)_{h=1}^H \right\}_{j=1}^{T_2}$ and we estimate the empirical value function of $\tilde{\pi}^t$ as follows:

$$\hat{V}_1^{\tilde{\pi}^t}(s_1, i_1 v) = \frac{1}{T_2} \sum_{j=1}^{T_2} \left(i_1 v - \sum_{h=1}^H r_h^{t,j}(s_h^{t,j}, a_h^{t,j}) \right)^+ - \sum_{h=1}^H \hat{b}_h(s_h^{t,j}, a_h^{t,j}).$$

Then we simply use $\max_{t \in [T_1]} \hat{V}_1^{\tilde{\pi}^t}(s_1, i_1 v)$ as a surrogate for $V_1^*(s_1, i_1 v)$.

The details of *CVaR-LSVI* are stated in Algorithm 2 (cf. Appendix B). Note that in Line 16 of Algorithm 1, we can also use the above *CVaR-LSVI* algorithm to compute π^k . Combining Algorithm 1 and 2, we can derive a computationally efficient algorithm, called *ELLA*, for CVaR objective in low-rank MDPs, which is shown in Algorithm 3 (cf. Appendix B). Now, we are ready present the computational complexity of *ELLA*. Particularly, the following theorem characterizes the computational cost for finding an ϵ -optimal policy:

Theorem 5.1 (Informal). *Let the parameters in Algorithm 1 and 2 take appropriate values, then we have with probability at least $1 - \delta$ that $\text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\hat{\pi}, \hat{c})) \leq \epsilon$ where $(\hat{\pi}, \hat{c})$ is the returned policy and initial budget by Algorithm 3. In total, the sample complexity is upper bounded by $\tilde{O}\left(\frac{H^7 A^2 d^4 \log \frac{|F|}{\delta}}{\tau^2 \epsilon^2}\right)$. The MLE oracle is called $\tilde{O}\left(\frac{H^7 A^2 d^4 \log \frac{|F|}{\delta}}{\tau^2 \epsilon^2}\right)$ times and the rest of the computation cost is $\tilde{O}\left(\frac{H^{19} A^3 d^{12} \log \frac{|F|}{\delta}}{v^{10} \tau^6 \epsilon^6}\right)$.*

Theorem 5.1 is a special case of the continuous reward setting, whose formal statement is in Theorem B.3 and the proof is deferred to Appendix D. Theorem 5.1 indicates that Algorithm 3 is able to find a near-optimal policy with polynomial sample complexity and polynomial computational complexity given an MLE oracle. Note that calling *CVaR-LSVI* in Line 17 and 20 of Algorithm 3 will not increase the sample complexity because we are only simulating with a known model $(\hat{\mathcal{P}}, \bar{r})$ and do not need to interact with the ground-truth environment.

6 CONCLUDING REMARKS

The paper proposes *ELA*, the first provably efficient algorithm for risk-sensitive reinforcement learning with the CVaR objective in low-rank MDPs. *ELA* achieves a sample complexity of $\tilde{O}(1/\epsilon^2)$ to find an ϵ -optimal policy. To improve computational efficiency, we propose the *ELLA* algorithm, which leverages least-squares value iteration upon discretized reward as the planning oracle. Given the MLE oracle, we prove this algorithm only requires a polynomial computational complexity.

Regarding future work, we believe establishing lower bounds for CVaR RL in low-rank MDPs is an exciting and challenging direction due to the complexity of the risk landscape and the non-linearity in function approximation methods used. We point out that even for standard episodic RL, which is a well-studied area, we have only recently seen some results on sample lower bounds in the context of low-rank MDPs (Cheng et al., 2023; Zhao et al., 2023). However, these results do not directly apply to the CVaR risk landscape. We suppose the sample complexity lower bound may take the form of $\Omega\left(\frac{HAd}{\tau \epsilon^2}\right)$ according to the lower bounds in low-rank MDPs (Cheng et al., 2023; Zhao et al., 2023) and tabular CVaR RL (Wang et al., 2023), which implies that our algorithm has a loose dependency on H and τ . Such a gap might be alleviated if we use the Hoeffding-type bonus instead of the Bernstein-type bonus in the bonus-driven exploration. We leave it as future work to tighten the dependencies.

REFERENCES

- Mastane Achab, Reda Alami, Yasser Abdelaziz Dahou Djilali, Kirill Fedyanin, Eric Moulines, and Maxim Panov. Distributional deep q-learning with cvar regression. In *Deep Reinforcement Learning Workshop NeurIPS 2022*, 2022.
- Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity and representation learning of low rank MDPs. *Advances in neural information processing systems*, 33:20095–20107, 2020.
- Osbert Bastani, Jason Yecheng Ma, Estelle Shen, and Wanqiao Xu. Regret bounds for risk-sensitive reinforcement learning. *Advances in Neural Information Processing Systems*, 35:36259–36269, 2022.
- Nicole Bäuerle and Jonathan Ott. Markov decision processes with average-value-at-risk criteria. *Mathematical Methods of Operations Research*, 74:361–379, 2011.
- Yu Chen, Yihan Du, Pihe Hu, Siwei Wang, Desheng Wu, and Longbo Huang. Provably efficient iterated cvar reinforcement learning with function approximation. *arXiv preprint arXiv:2307.02842*, 2023.
- Yuan Cheng, Ruiquan Huang, Jing Yang, and Yingbin Liang. Improved sample complexity for reward-free reinforcement learning under low-rank mdps. *arXiv preprint arXiv:2303.10859*, 2023.
- Yinlam Chow and Mohammad Ghavamzadeh. Algorithms for CVaR optimization in MDPs. *Advances in neural information processing systems*, 27, 2014.
- Yinlam Chow, Aviv Tamar, Shie Mannor, and Marco Pavone. Risk-sensitive and robust decision-making: a CVaR optimization approach. *Advances in neural information processing systems*, 28, 2015.
- Antonio Coronato, Muddasar Naeem, Giuseppe De Pietro, and Giovanni Paragliola. Reinforcement learning for intelligent healthcare applications: A survey. *Artificial Intelligence in Medicine*, 109: 101964, 2020.
- Mark Davis and Sébastien Lleo. Risk-sensitive benchmarked asset management. *Quantitative Finance*, 8(4):415–426, 2008.
- Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudik, and John Langford. Provably efficient RL with rich observations via latent state decoding. In *International Conference on Machine Learning*, pp. 1665–1674. PMLR, 2019.
- Simon Du, Sham Kakade, Jason Lee, Shachar Lovett, Gaurav Mahajan, Wen Sun, and Ruosong Wang. Bilinear classes: A structural framework for provable generalization in RL. In *International Conference on Machine Learning*, pp. 2826–2836. PMLR, 2021.
- Yihan Du, Siwei Wang, and Longbo Huang. Provably efficient risk-sensitive reinforcement learning: Iterated cvar and worst path. In *The Eleventh International Conference on Learning Representations*, 2022.
- Yonathan Efroni, Dipendra Misra, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Provable rl with exogenous distractors via multistep inverse dynamics. *arXiv preprint arXiv:2110.08847*, 2021.
- Damien Ernst, Guy-Bart Stan, Jorge Goncalves, and Louis Wehenkel. Clinical data based optimal STI strategies for HIV: a reinforcement learning approach. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pp. 667–672. IEEE, 2006.
- Carlo Filippi, Gianfranco Guastaroba, and Maria Grazia Speranza. Conditional value-at-risk beyond finance: a survey. *International Transactions in Operational Research*, 27(3):1277–1319, 2020.
- Takuya Hiraoka, Takahisa Imagawa, Tatsuya Mori, Takashi Onishi, and Yoshimasa Tsuruoka. Learning robust options by conditional value at risk optimization. *Advances in Neural Information Processing Systems*, 32, 2019.

- Xiaoyan Hu and Ho-Fung Leung. A tighter problem-dependent regret bound for risk-sensitive reinforcement learning. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 5411–5437. PMLR, 25–27 Apr 2023. URL <https://proceedings.mlr.press/v206/hu23b.html>.
- Audrey Huang, Jinglin Chen, and Nan Jiang. Reinforcement learning in low-rank MDPs with density features. *arXiv preprint arXiv:2302.02252*, 2023.
- David Isele, Alireza Nakhaei, and Kikuo Fujimura. Safe reinforcement learning on autonomous vehicles. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1–6. IEEE, 2018.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are PAC-learnable. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2017.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020.
- Pavlo Krokhmal, Jonas Palmquist, and Stanislav Uryasev. Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of risk*, 4:43–68, 2002.
- R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.
- Ayush Sekhari, Christoph Dann, Mehryar Mohri, Yishay Mansour, and Karthik Sridharan. Agnostic reinforcement learning with low-rank MDPs and rich observations. *Advances in Neural Information Processing Systems*, 34:19033–19045, 2021.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of Go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Shagun Sodhani, Amy Zhang, and Joelle Pineau. Multi-task reinforcement learning with context-based representations. In *International Conference on Machine Learning*, pp. 9767–9779. PMLR, 2021.
- Shagun Sodhani, Franziska Meier, Joelle Pineau, and Amy Zhang. Block contextual MDPs for continual learning. In *Learning for Dynamics and Control Conference*, pp. 608–623. PMLR, 2022.
- Silvestr Stanko and Karel Macek. Risk-averse distributional reinforcement learning: A CVaR optimization approach. In *IJCCI*, pp. 412–423, 2019.
- Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Model-based RL in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Conference on learning theory*, pp. 2898–2933. PMLR, 2019.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Aviv Tamar, Yonatan Glassner, and Shie Mannor. Optimizing the CVaR via sampling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- Masatoshi Uehara, Xuezhou Zhang, and Wen Sun. Representation learning for online and offline RL in low-rank MDPs. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=J4iSIR9fhY0>.

- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Kaiwen Wang, Nathan Kallus, and Wen Sun. Near-minimax-optimal risk-sensitive reinforcement learning with CVaR. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 35864–35907. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/wang23m.html>.
- Zhicheng Wang, Biwei Huang, Shikui Tu, Kun Zhang, and Lei Xu. Deeptrader: a deep reinforcement learning approach for risk-return balanced portfolio management with market conditions embedding. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 643–650, 2021.
- Lu Wen, Jingliang Duan, Shengbo Eben Li, Shaobing Xu, and Hui Peng. Safe reinforcement learning for autonomous vehicles through parallel constrained policy optimization. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1–7. IEEE, 2020.
- 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.
- Amy Zhang, Clare Lyle, Shagun Sodhani, Angelos Filos, Marta Kwiatkowska, Joelle Pineau, Yarin Gal, and Doina Precup. Invariant causal prediction for block MDPs. In *International Conference on Machine Learning*, pp. 11214–11224. PMLR, 2020.
- Amy Zhang, Rowan Thomas McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=-2FCwDKRREu>.
- Canzhe Zhao, Ruofeng Yang, Baoxiang Wang, Xuezhou Zhang, and Shuai Li. Learning adversarial low-rank markov decision processes with unknown transition and full-information feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

A NOTATIONS

Table 1: List of Notations

Basics	
$A = \mathcal{A} $	cardinality of the action space
$[H] = \{1, \dots, H\}$	
$\Delta(\cdot)$	probability simplex
$Z_{\mathcal{P}, c_1}^\pi$	return distribution of rolling π with initial budget c_1 in an MDP with \mathcal{P}
$V_{h, \mathcal{P}, b}^\pi : \mathcal{S} \times [0, H] \mapsto [0, H]$	defined in (5)
$d_{h, \mathcal{P}}^{(\pi, c_1)}(s), d_{h, \mathcal{P}}^{(\pi, c_1)}(s, a)$	occupancy of s and (s, a) at step h when rolling π with initial budget c_1 in an MDP with \mathcal{P}
$d_h^{(\pi, c_1)}(s) = d_{h, \mathcal{P}^*}^{(\pi, c_1)}(s), d_h^{(\pi, c_1)}(s, a) = d_{h, \mathcal{P}^*}^{(\pi, c_1)}(s, a)$	occupancy of s and (s, a) at step h when rolling π with initial budget c_1 in the (true) environment
In-Algorithm	
Initialization	
$\mathcal{F} = \{\Psi, \Phi\}$	model class
$\lambda^k = O(d \log(\mathcal{F} H k / \delta))$	regularizer at iteration k
$\alpha^k = \sqrt{H^2(A + d^2) \log(\mathcal{F} H k / \delta)}$	parameter at iteration k
Datasets (at the end of the k -th episode)	
$\mathcal{D}_h^k = \{(s_h^i, a_h^i, \tilde{s}_{h+1}^i)\}_{i \in [k]}$	$s_h^i \sim d_{h, c^{i-1}}^{\pi^{i-1}}, a_h^i \sim U(\mathcal{A})$
$\tilde{\mathcal{D}}_h^k = \{(\tilde{s}_h^i, \tilde{a}_h^i, s_{h+1}^i)\}_{i \in [k]}$	$\tilde{s}_h^i \sim d_{h-1, c^{i-1}}^{\pi^{i-1}} \times U(\mathcal{A}) \times P_h^*, \tilde{a}_h^i \sim U(\mathcal{A})$
Estimations (at the end of the k -th episode)	
$\{(\hat{\psi}_h^k, \hat{\phi}_h^k)\}_{h \in [H]}$	learned representations
$\hat{\mathcal{P}}^k = \{\hat{P}_h^k\}_{h \in [H]}$	empirical transition kernel
$\hat{\Sigma}_h^k = \sum_{(s, a) \in \mathcal{D}_h^k} \hat{\phi}_h^k (\hat{\phi}_h^k)^\top + \lambda^k I$	empirical covariance matrix
$\hat{b}^k = \{\hat{b}_h^k\}_{h \in [H]}$	bonus term
In-Analysis	
$\rho_h^k(s) = \frac{1}{k} \sum_{i=0}^{k-1} d_{h, c^i}^{\pi^i}(s)$	occupancy of s in dataset \mathcal{D}_h^k
$\rho_h^k(s, a) = \frac{1}{k} \sum_{i=0}^{k-1} d_{h, c^i}^{\pi^i}(s, a)$	
$\eta_h^k(s) = \sum_{s', a'} \rho_{h-1}^k(s') U(a') P_{h-1}^*(s s', a')$	occupancy of s in dataset $\tilde{\mathcal{D}}_h^k$
$\Sigma_{\rho_h^k \times U(\mathcal{A}), \phi} = k \mathbb{E}_{s \sim \rho_h^k, a \sim U(\mathcal{A})} [\phi \phi^\top] + \lambda^k I$	
$\Sigma_{\rho_h^k, \phi} = k \mathbb{E}_{(s, a) \sim \rho_h^k} [\phi \phi^\top] + \lambda^k I$	
$\hat{\Sigma}_{h, \phi}^k = k \mathbb{E}_{(s, a) \sim \mathcal{D}_h^k} [\phi \phi^\top] + \lambda^k I$	unbiased estimate of $\Sigma_{\rho_h^k \times U(\mathcal{A}), \phi}$
$\zeta^k = \log(\mathcal{F} H k / \delta) / k$	
$f_h^k(s, a) = \ \hat{P}_h^k(\cdot s, a) - P_h^*(\cdot s, a)\ _1$	estimation error in L_1 norm
$\omega_{h, \mathcal{P}}^{(\pi, c_1)}(\cdot s, a)$	the distribution of the remaining budget at h for any (s, a) when rolling out (π, c_1) in an MDP with \mathcal{P}
$\omega_h^{(\pi, c_1)}(\cdot s, a) = \omega_{h, \mathcal{P}^*}^{(\pi, c_1)}(\cdot s, a)$	the distribution of the remaining budget at h for any (s, a) when rolling out (π, c_1) in the true MDP

B ELLA FOR CONTINUOUS REWARD

Now, we extend the analysis in Section 5 to continuous reward distribution.

B.1 DISCRETIZED REWARD

Inspired by Wang et al. (2023), we discretize the reward r_h and the budget c_h at each step. In this way, we only need to plan over a finite grid. More specifically, suppose the precision is $v \geq 0$, then we round up the reward r to $U(r) := \lceil r/v \rceil v$. In the following discussion, we use $\overline{\mathcal{M}}$ to denote the MDP which shares the same transition as \mathcal{M} while its reward \bar{r} is discretized from the original reward distribution r of \mathcal{M} , i.e., $\bar{r} = U(r)$. Since the supremum of the CVaR objective (2) is attained at τ -th quantile of the return distribution, which is also discretized in $(\widehat{\mathcal{P}}, \bar{r})$, we can only search c^k within the grid for $(\widehat{\mathcal{P}}, \bar{r})$ in Line 15 of Algorithm 1:

$$c^k \leftarrow v \cdot \arg \max_{0 \leq i \leq \lceil H/v \rceil} \left\{ iv - \tau^{-1} \min_{\pi \in \overline{\Pi}_{Aug}} \overline{V}_{1, \widehat{\mathcal{P}}, \widehat{b}}^\pi(s_1, iv) \right\}, \quad (7)$$

where $\overline{\Pi}_{Aug}$ is the augmented policy class of augmented $\overline{\mathcal{M}}$ and $\overline{V}_{h, \mathcal{P}, b}^\pi$ is the value function of (\mathcal{P}, \bar{r}) with bonus b :

$$\overline{V}_{h, \mathcal{P}, b}^\pi(s, c) := \mathbb{E}_{\pi, \mathcal{P}} \left[\left(c_h - \sum_{h'=h}^H U(r_{h'}) \right)^+ - \sum_{h'=h}^H b_{h'}(s_{h'}, a_{h'}) \middle| s_h = s, c_h = c \right].$$

Similarly, the Q function of (\mathcal{P}, \bar{r}) with bonus b is defined as:

$$\overline{Q}_{h, \mathcal{P}, b}^\pi(s, c, a) := \mathbb{E}_{\pi, \mathcal{P}} \left[\left(c_h - \sum_{h'=h}^H U(r_{h'}) \right)^+ - \sum_{h'=h}^H b_{h'}(s_{h'}, a_{h'}) \middle| s_h = s, c_h = c, a_h = a \right].$$

To stay consistent with Line 15, we also derive the optimal augmented policy of $(\widehat{\mathcal{P}}, \bar{r})$ with bonus \widehat{b} in Line 16 of Algorithm 1, i.e.,

$$\pi^k \leftarrow \arg \min_{\pi \in \overline{\Pi}_{Aug}} \overline{V}_{1, \widehat{\mathcal{P}}, \widehat{b}}^\pi(s_1, c^k).$$

Note that in Line 8 of Algorithm 1 we need to roll out π^k in \mathcal{M} to collect samples. Since $\pi^k \in \overline{\Pi}_{Aug}$ works only within the grid, we discretize the reward we observe when executing π^k in \mathcal{M} , which is equivalent to playing the following augmented policy $\overline{\pi}^k \in \overline{\Pi}_{Aug}$ in \mathcal{M} :

$$\overline{\pi}_h^k \left(s, c^k - \sum_{t=1}^{h-1} r_t \right) = \pi_h^k \left(s, c^k - \sum_{t=1}^{h-1} U(r_t) \right), \forall h \in [H]. \quad (8)$$

Planning within a discretized grid via (7) will inevitably incur errors compared to the original objective. Nevertheless, we can show that if the precision v is sufficiently small, the discretized MDP $\overline{\mathcal{M}}$ will be an excellent approximation to \mathcal{M} and thus the induced error of planning via (7) will be negligible. Formally, let $\overline{R}(\pi, c)$ denote the return of executing $\pi \in \overline{\Pi}_{Aug}$ with initial budget $c = iv$ ($0 \leq i \leq \lceil H/v \rceil$) in $\overline{\mathcal{M}}$, then we have the following properties from the literature (Wang et al., 2023):

Proposition B.1. *For any $0 < \tau < 1, v > 0$, policy $\pi \in \overline{\Pi}_{Aug}$ and $0 \leq i \leq \lceil H/v \rceil$, we have*

- (1) $\text{CVaR}_\tau^* \leq \overline{\text{CVaR}}_\tau^* := \max_{\pi \in \overline{\Pi}_{Aug}, 0 \leq i \leq \lceil H/v \rceil} \text{CVaR}_\tau(\overline{R}(\pi, iv))$,
- (2) $0 \leq \text{CVaR}_\tau(\overline{R}(\pi, iv)) - \text{CVaR}_\tau(R(\overline{\pi}, iv)) \leq \frac{Hv}{\tau}$.

B.2 CVaR-LSVI

With discretization, we only need to search within the grid of c in each iteration. For each discrete value iv , we apply the CVaR-LSVI algorithm as planning oracle as introduced in Section 5. The only difference is that now we simulate with the discretized reward model \bar{r} . The full algorithm is shown in Algorithm 2. In particular, in Line 12 we can express \bar{Q}_h^t with w_h^t :

$$\bar{Q}_h^t(s, iv, a) = \text{Clip}_{[-H, H]} \left(-\hat{b}_h(s, a) + (\bar{\phi}_h(s, a))^\top w_h^t(iv) - \beta \|\bar{\phi}_h(s, a)\|_{(\Lambda_h^t)^{-1}} \right). \quad (9)$$

Moreover, note that we can bound the norm of the newly-constructed feature vectors as follows:

$$\|\bar{\phi}_h(s, a)\|_2 \leq 1, \quad \left\| \sum_{0 \leq i \leq \lceil 1/v \rceil} \int_{\mathcal{S}} \bar{\psi}_h(s', iv) ds \right\| \leq (1 + \lceil 1/v \rceil) \sqrt{d}.$$

This bound will be useful in our analysis.

Algorithm 2 CVaR-LSVI

Require: MDP transition model and reward ($\hat{\mathcal{P}} = (\hat{\phi}, \hat{\psi}, \bar{r})$), bonus \hat{b} , initial budget $i_1 v$, number of iterations T_1 , parameters λ and β , number of policy evaluation episodes T_2 .

- 1: Compute $\bar{\phi}_h(s, a) = \hat{\phi}_h(s, a) \otimes \phi_{h,r}(s, a)$ for all $h \in [H], s \in \mathcal{S}, a \in \mathcal{A}$ where $\phi_{h,r}(s, a) \in \mathbb{R}^{\lceil 1/v \rceil + 1}$ and $(\phi_{h,r}(s, a))_i = \bar{r}_h(iv|s, a)$ for all $0 \leq i \leq \lceil 1/v \rceil$.
- 2: **for** $t = 1, \dots, T_1$ **do**
- 3: Initialize $\bar{V}_{H+1}^t(s, iv) \leftarrow iv$ for all $s \in \mathcal{S}$ and $0 \leq i \leq \lceil H/v \rceil$.
- 4: **for** $h = H, \dots, 1$ **do**
- 5: Compute $\Lambda_h^t \leftarrow \lambda I + \sum_{j=1}^{t-1} \bar{\phi}_h(s_h^j, a_h^j) (\bar{\phi}_h(s_h^j, a_h^j))^\top$.
- 6: Calculate $w_h^t(iv) \leftarrow (\Lambda_h^t)^{-1} \sum_{j=1}^{t-1} \bar{\phi}_h(s_h^j, a_h^j) \cdot \left[\bar{V}_{h+1}^t(s_{h+1}^j, iv - \bar{r}_h^j) \right], \forall 0 \leq i \leq \lceil H/v \rceil$.
- 7: **for** $j = 1, \dots, t-1$ **do**
- 8: Compute for all $a \in \mathcal{A}, 0 \leq i \leq \lceil H/v \rceil$:

$$\bar{Q}_h^t(s_h^j, iv, a) \leftarrow \text{Clip}_{[-H, H]} \left(-\hat{b}_h(s_h^j, a) + (\bar{\phi}_h(s_h^j, a))^\top w_h^t(iv) - \beta \|\bar{\phi}_h(s_h^j, a)\|_{(\Lambda_h^t)^{-1}} \right).$$

- 9: $\bar{V}_h^t(s_h^j, iv) \leftarrow \min_{a \in \mathcal{A}} \bar{Q}_h^t(s_h^j, iv, a)$.
 - 10: **end for**
 - 11: **end for**
 - 12: Simulate the greedy policy $\tilde{\pi}^t$ (w.r.t. $\bar{Q}_h^t(s, iv, a)$ defined in (9)) with the initial budget $c_1 = i_1 v$ in the MDP $(\hat{\mathcal{P}}, \bar{r})$ and collect a trajectory $(s_h^t, a_h^t, \bar{r}_h^t)_{h=1}^H$.
 - 13: **end for**
 - 14: **for** $t = 1, \dots, T_1$ **do**
 - 15: Simulate $\tilde{\pi}^t$ with initial budget $c_1 = i_1 v$ in $(\hat{\mathcal{P}}, \bar{r})$ for T_2 episodes and collect trajectories $\left\{ \left(s_h^{t,j}, a_h^{t,j}, \bar{r}_h^{t,j} \right)_{h=1}^H \right\}_{j=1}^{T_2}$.
 - 16: Compute $\widehat{V}_1^{\tilde{\pi}^t}(s_1, i_1 v) \leftarrow \frac{1}{T_2} \sum_{j=1}^{T_2} \left(i_1 v - \sum_{h=1}^H \bar{r}_h^{t,j}(s_h^{t,j}, a_h^{t,j}) \right)^+ - \sum_{h=1}^H \hat{b}_h(s_h^{t,j}, a_h^{t,j})$.
 - 17: **end for**
 - 18: **Return:** value estimate $\min_{t \in [T_1]} \widehat{V}_1^{\tilde{\pi}^t}(s_1, i_1 v)$ and policy $\arg \min_{\tilde{\pi}^t} \widehat{V}_1^{\tilde{\pi}^t}(s_1, i_1 v)$
-

Now let $\bar{V}_{h, \hat{\mathcal{P}}, \hat{b}}^*$ denote $\min_{\pi \in \bar{\Pi}_{Avg}} \bar{V}_{h, \hat{\mathcal{P}}, \hat{b}}^\pi$. Then we have the following theorem indicating that Algorithm 2 can do planning in $(\hat{\mathcal{P}}, \bar{r})$ accurately with appropriate T_1 and T_2 :

Theorem B.2. *Let*

$$\lambda = 1, \beta = \tilde{O} \left(\frac{H^{\frac{3}{2}} d \iota^{\frac{1}{4}}}{v} \right), T_1 = \tilde{O} \left(\frac{H^5 d^3 \iota}{v^3 \varepsilon^2} \right), T_2 = \tilde{O} \left(\frac{H^2 \log \frac{T_1}{\delta}}{\varepsilon^2} \right),$$

where $\iota = \log^2 \frac{HdT_1}{v\delta}$, we have with probability at least $1 - \delta$ that

$$(1) \quad \widehat{V}_{1,\widehat{\mathcal{P}},\widehat{b}}^*(s_1, i_1 v) \leq \overline{V}_{1,\widehat{\mathcal{P}},\widehat{b}}^*(s_1, i_1 v) + \frac{3}{4}\varepsilon,$$

$$(2) \quad \left| \widehat{V}_{1,\widehat{\mathcal{P}},\widehat{b}}^*(s_1, i_1 v) - \widehat{V}_{1,\widehat{\mathcal{P}},\widehat{b}}^{\widehat{\pi}}(s_1, i_1 v) \right| \leq \frac{1}{4}\varepsilon.$$

where $\widehat{\pi} = \arg \min_{\pi^t} \widehat{V}_1^{\pi^t}(s_1, i_1 v)$ and $\widehat{V}_{1,\widehat{\mathcal{P}},\widehat{b}}^*(s_1, i_1 v) := \min_{\pi^t} \widehat{V}_1^{\pi^t}(s_1, i_1 v)$ are the returned values of CVaR-LSVI.

The proof is deferred to Appendix D.1.

B.3 COMPUTATIONAL COMPLEXITY

Equipping *ELA* with *CVaR-LSVI*, we can derive *ELLA*, which is shown in Algorithm 3. Based on the above discussions about discretization and *CVaR-LSVI*, the computational complexity of *ELLA* for finding an ϵ -policy can be characterized as follows:

Theorem B.3. *Let*

$$\begin{aligned} \alpha^k &= O\left(\sqrt{H^2(|\mathcal{A}| + d^2)} \log\left(\frac{|\mathcal{F}|Hk}{\delta}\right)\right), \quad \lambda^k = O\left(d \log\left(\frac{|\mathcal{F}|Hk}{\delta}\right)\right), \\ K &= \tilde{O}\left(\frac{H^6 A^2 d^4 \log \frac{|\mathcal{F}|}{\delta}}{\tau^2 \epsilon^2}\right), \quad v = \frac{\epsilon \tau}{3H}, \quad \lambda = 1, \quad \beta = \tilde{O}\left(\frac{H^{\frac{3}{2}} d \iota^{\frac{1}{4}}}{v}\right), \\ T_1 &= \tilde{O}\left(\frac{H^5 d^3 \iota}{v^3 \tau^2 \epsilon^2}\right), \quad T_2 = \tilde{O}\left(\frac{H^2 \log \frac{T_1}{\delta}}{\tau^2 \epsilon^2}\right). \end{aligned}$$

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

$$\text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\widehat{\pi}, \widehat{c})) \leq \epsilon.$$

In total, the sample complexity is upper bounded by $\tilde{O}\left(\frac{H^7 A^2 d^4 \log \frac{|\mathcal{F}|}{\delta}}{\tau^2 \epsilon^2}\right)$. The MLE oracle is called

$$\tilde{O}\left(\frac{H^7 A^2 d^4 \log \frac{|\mathcal{F}|}{\delta}}{\tau^2 \epsilon^2}\right) \text{ times and the rest of the computation cost is } \tilde{O}\left(\frac{H^{29} A^3 d^{12} \log \frac{|\mathcal{F}|}{\delta}}{\tau^{16} \epsilon^{16}}\right).$$

The proof is deferred to Appendix D.2. Theorem B.3 indicates that even when the reward is continuous, Algorithm 3 can still achieve polynomial sample complexity and polynomial computational complexity given an MLE oracle.

C PROOFS FOR SECTION 4

Proof Sketch Recall that π^k and c^k are the exploration policy and initial budget output from Line 16 of Algorithm 1 at the end of the k -th iteration, respectively. In the proof, we show that, with probability at least $1 - \delta$, it holds that

$$\text{Regret}(K) \lesssim \tau^{-1} H^3 A d^2 \sqrt{K} \sqrt{\log\left(1 + \frac{K}{d\lambda^1}\right) \log\left(\frac{|\mathcal{F}|Hk}{\delta}\right)} \quad (10)$$

which is formalized in Lemma C.1. Once it is established, the suboptimality of the uniform mixture of $\{(c^k, \pi^k)\}$ satisfies that

$$\text{CVaR}_\tau^* - \frac{1}{K} \sum_{k=1}^K \text{CVaR}_\tau(R(\pi^k, c^k)) \lesssim \tau^{-1} H^3 A d^2 K^{-\frac{1}{2}} \sqrt{\log\left(1 + \frac{K}{d\lambda^1}\right) \log\left(\frac{|\mathcal{F}|Hk}{\delta}\right)} \quad (11)$$

Let the RHS smaller than ϵ , we have that when

$$K \gtrsim O\left(\frac{H^6 A^2 d^4}{\tau^2 \epsilon^2} \log\left(1 + \frac{K}{d\lambda^1}\right) \log\left(\frac{|\mathcal{F}|Hk}{\delta}\right)\right)$$

Algorithm 3 *ELLA*

Require: Risk tolerance $\tau \in (0, 1]$, number of iterations K , parameters $\{\lambda^k\}_{k \in [K]}$ and $\{\alpha^k\}_{k \in [K]}$, models $\mathcal{F} = \{\Psi, \Phi\}$, failure probability $\delta \in (0, 1)$, discretization precision v , *CVaR-LSVI* parameters λ, β, T_1, T_2 .

- 1: Set datasets $\mathcal{D}_h, \tilde{\mathcal{D}}_h \leftarrow \emptyset$ for each $h \in [H - 1]$.
- 2: Calculate the discretized reward distribution \bar{r} .
- 3: Initialize the exploration policy $\pi^0 \leftarrow \{\pi_h^0(s, iv) = U(\mathcal{A}), \text{ for any } s \in \mathcal{S}, 0 \leq i \leq \lceil H/v \rceil\}_{h \in [H]}$.
- 4: Initialize the budget $i^0 \leftarrow \lceil H/v \rceil$.
- 5: **for** iteration $k = 1, \dots, K$ **do**
- 6: Collect a tuple $(\tilde{s}_1, \tilde{a}_1, s'_2)$ by taking $\tilde{a}_1 \sim U(\mathcal{A}), s'_2 \sim P_1^*(\cdot | \tilde{s}_1, \tilde{a}_1)$.
- 7: Update $\tilde{\mathcal{D}}_1 \leftarrow \tilde{\mathcal{D}}_1 \cup \{(\tilde{s}_1, \tilde{a}_1, s'_2)\}$.
- 8: **for** $h = 1, \dots, H - 1$ **do**
- 9: Collect two transition tuples $(s_h, a_h, \tilde{s}_{h+1})$ and $(\tilde{s}_{h+1}, \tilde{a}_{h+1}, s'_{h+2})$ by first rolling out $\bar{\pi}^{k-1}$ (defined in (8)) starting from $(s_1, i^{k-1}v)$ into state s_h , taking $a_h \sim U(\mathcal{A})$, and receiving $\tilde{s}_{h+1} \sim P_h^*(\cdot | s_h, a_h)$, then taking $\tilde{a}_{h+1} \sim U(\mathcal{A})$ and receiving $s'_{h+2} \sim P_{h+1}^*(\cdot | \tilde{s}_{h+1}, \tilde{a}_{h+1})$.
- 10: Update $\mathcal{D}_h \leftarrow \mathcal{D}_h \cup \{(s_h, a_h, \tilde{s}_{h+1})\}$.
- 11: Update $\tilde{\mathcal{D}}_{h+1} \leftarrow \tilde{\mathcal{D}}_{h+1} \cup \{(\tilde{s}_{h+1}, \tilde{a}_{h+1}, s'_{h+2})\}$ if $h \leq H - 2$.
- 12: Learn representations via MLE

$$\hat{P}_h := (\hat{\psi}_h, \hat{\phi}_h) \leftarrow \arg \max_{(\psi, \phi) \in \mathcal{F}} \sum_{(s_h, a_h, s_{h+1}) \in \{\mathcal{D}_h + \tilde{\mathcal{D}}_h\}} \log \langle \psi(s_{h+1}), \phi(s_h, a_h) \rangle$$

- 13: Update empirical covariance matrix $\hat{\Sigma}_h = \sum_{(s, a) \in \mathcal{D}_h} \hat{\phi}_h(s, a) \hat{\phi}_h(s, a)^\top + \lambda^k I_d$.
- 14: Set the exploration bonus:

$$\hat{b}_h(s, a) \leftarrow \begin{cases} \min \left(\alpha^k \sqrt{\hat{\phi}_h(s, a) \hat{\Sigma}_h^{-1} \hat{\phi}_h(s, a)^\top}, 2 \right) & h \leq H - 2 \\ 0 & h = H - 1 \end{cases}$$

- 15: **end for**
 - 16: **for** $i = 0, 1, \dots, \lceil H/v \rceil$ **do**
 - 17: Run *CVaR-LSVI* (Algorithm 2) with MDP model (\hat{P}, \bar{r}) , bonus \hat{b} , initial budget iv and parameters $(\lambda, \beta, T_1, T_2)$ and let the returned value estimate and policy be $\hat{V}_1^*(s_1, iv)$ and $\hat{\pi}(i)$.
 - 18: **end for**
 - 19: Obtain $i^k \leftarrow \arg \max_{0 \leq i \leq \lceil H/v \rceil} \left\{ iv - \tau^{-1} \hat{V}_1^*(s_1, iv) \right\}$ and $\pi^k \leftarrow \hat{\pi}(i^k)$.
 - 20: **end for**
- Ensure:** Uniformly sample k from $[K]$, return $(\hat{\pi}, \hat{c}) = (\pi^k, i^k v)$.

the uniform mixture of $\{(\pi^k, c^k)\}$ is an ϵ -optimal policy. Finally, noting that H trajectories are collected per iteration, we conclude the proof of Theorem 4.1. To establish (10), we decompose the suboptimality of the k -th iteration into

$$\begin{aligned} & \text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\pi^k, c^k)) \\ &= c^* - \tau^{-1} V_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^*}(s_1, c^*) - \text{CVaR}_\tau(R(\pi^k, c^k)) + \text{CVaR}_\tau(R(\pi^*, c^*)) - \left(c^* - \tau^{-1} V_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^*}(s_1, c^*) \right) \\ &\leq c^k - \tau^{-1} V_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^k}(s_1, c^k) - \text{CVaR}_\tau(R(\pi^k, c^k)) + c^* - \tau^{-1} V_{1, \mathcal{P}^*, 0}^{\pi^*}(s_1, c^*) - \left(c^* - \tau^{-1} V_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^*}(s_1, c^*) \right) \\ &\leq \tau^{-1} \underbrace{\left(V_{1, \mathcal{P}^*, 0}^{\pi^k}(s_1, c^k) - V_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^k}(s_1, c^k) \right)}_{(i)} + \tau^{-1} \underbrace{\left(V_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^*}(s_1, c^*) - V_{1, \mathcal{P}^*, 0}^{\pi^*}(s_1, c^*) \right)}_{(ii)} \end{aligned}$$

where the first inequality holds by the fact that π^k is greedy (Line 16 of Algorithm 1) and the last inequality holds by

$$\text{CVaR}_\tau(R(\pi^k, c^k)) = \sup_{t \in \mathbb{R}} (t - \tau^{-1} \mathbb{E}_{R \sim R(\pi^k, c^k)}[(t - R)^+]) \geq c^k - \tau^{-1} \mathbb{E}_{R \sim R(\pi^k, c^k)}[(c^k - R)^+]$$

Utilizing simulation lemma C.3 for risk-sensitive RL, we further upper bound terms (i) and (ii) by the error in estimating the transition kernel and the bonus terms, i.e.,

$$\begin{aligned} \text{term (i)} &\leq \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} [\widehat{b}_h^k(s, a)] + 3H \cdot \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} [f_h^k(s, a)] \\ \text{term (ii)} &\leq \mathbb{E}_{(s,a) \sim d_{h, \widehat{P}^k}^{(\pi^*, c^*)}} [H \cdot f_h^k(s, a) - \widehat{b}_h^k(s, a)] \end{aligned}$$

where $f_h^k(s, a) := \|\widehat{P}_h^k(\cdot | s, a) - P_h^*(\cdot | s, a)\|_1$. By the analysis in Lemma C.1 and C.2, we prove the inequality (10).

Before proceeding to the detailed proofs, we first introduce the essential regret decomposition.

Lemma C.1 (Regret). *With probability $1 - \delta$, we have that*

$$\sum_{k=1}^K \text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\pi^k, c^k)) \lesssim \tau^{-1} H^3 A d^2 \sqrt{K} \sqrt{\log \left(1 + \frac{K}{d\lambda^1}\right) \log \left(\frac{|\mathcal{F}| H k}{\delta}\right)} \quad (12)$$

Proof. Letting $f_h^k(s, a) = \|\widehat{P}_h^k(\cdot | s, a) - P_h^*(\cdot | s, a)\|$, we condition on the event that for all $(k, h) \in [K] \times [H]$, the following inequalities hold

$$\begin{aligned} \mathbb{E}_{s \sim \rho_h^k, a \sim U(\mathcal{A})} \left[(f_h^k(s, a))^2 \right] &\leq \zeta^k, \quad \mathbb{E}_{s \sim \eta_h^k, a \sim U(\mathcal{A})} \left[(f_h^k(s, a))^2 \right] \leq \zeta^k \\ \|\phi(s, a)\|_{(\widehat{\Sigma}_{h, \phi}^k)^{-1}} &= \Theta(\|\phi(s, a)\|_{(\Sigma_{\rho_h^k \times U(\mathcal{A}), \phi}^k)^{-1}}) \end{aligned}$$

By Lemmas E.1 and E.2, this event happens with probability $1 - \delta$. For any iteration k , we have that For a fixed iteration k , we have that

$$\begin{aligned} &\text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\pi^k, c^k)) \\ &= c^* - \tau^{-1} V_{1, \widehat{P}^k, \widehat{b}^k}^{\pi^*}(s_1, c^*) - \text{CVaR}_\tau(R(\pi^k, c^k)) + \text{CVaR}_\tau(R(\pi^*, c^*)) - \left(c^* - \tau^{-1} V_{1, \widehat{P}^k, \widehat{b}^k}^{\pi^*}(s_1, c^*) \right) \\ &\leq c^k - \tau^{-1} V_{1, \widehat{P}^k, \widehat{b}^k}^{\pi^k}(s_1, c^k) - \text{CVaR}_\tau(R(\pi^k, c^k)) + c^* - \tau^{-1} V_{1, \mathcal{P}^*, 0}^{\pi^*}(s_1, c^*) - \left(c^* - \tau^{-1} V_{1, \widehat{P}^k, \widehat{b}^k}^{\pi^*}(s_1, c^*) \right) \\ &\leq \tau^{-1} \left(V_{1, \mathcal{P}^*, 0}^{\pi^k}(s_1, c^k) - V_{1, \widehat{P}^k, \widehat{b}^k}^{\pi^k}(s_1, c^k) \right) + \tau^{-1} \underbrace{\left(V_{1, \widehat{P}^k, \widehat{b}^k}^{\pi^*}(s_1, c^*) - V_{1, \mathcal{P}^*, 0}^{\pi^*}(s_1, c^*) \right)}_{\leq \sqrt{H^2 A \zeta^k} \text{ by Lemma C.2}} \quad (13) \end{aligned}$$

where the first inequality holds by the fact that π^k is greedy (Line 15 of Algorithm 1) and the last inequality holds by

$$\text{CVaR}_\tau(R(\pi^k, c^k)) = \sup_{t \in \mathbb{R}} (t - \tau^{-1} \mathbb{E}_{R \sim R(\pi^k, c^k)}[(t - R)^+]) \geq c^k - \tau^{-1} \mathbb{E}_{R \sim R(\pi^k, c^k)}[(c^k - R)^+]$$

Therefore, it remains to bound the first term, which by simulation lemma C.3, can be further written as

$$\begin{aligned} &V_{1, \mathcal{P}^*, 0}^{\pi^k}(s_1, c^k) - V_{1, \widehat{P}^k, \widehat{b}^k}^{\pi^k}(s_1, c^k) \\ &= \mathbb{E}_{\pi^k, \mathcal{P}^*} \left[\left(c^k - \sum_{h=1}^H r_h(s_h, a_h) \right)^+ \right] - \mathbb{E}_{\pi^k, \widehat{P}^k} \left[\left(c^k - \sum_{h=1}^H r_h(s_h, a_h) \right)^+ - \sum_{h=1}^H \widehat{b}_h^k(s_h, a_h) \middle| c_1 = c^k \right] \\ &\leq \sum_{h=1}^H \mathbb{E}_{\pi^k, \widehat{P}^k} \left[\widehat{b}_h^k(s_h, a_h) \middle| c_1 = c^k \right] + H \cdot \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} [f_h^k(s, a)] \\ &\leq \underbrace{\sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} [\widehat{b}_h^k(s, a)]}_{\text{(i)}} + 3H \cdot \underbrace{\sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} [f_h^k(s, a)]}_{\text{(ii)}} \quad (14) \end{aligned}$$

where the last inequality holds by the simulation Lemma E.5 for risk-neutral RL and the fact that $\|\widehat{b}_h^k\|_\infty \leq 2$. where the last inequality holds by $\|V_{h, \widehat{\rho}_h^k, r+\widehat{b}_h^k}\|_\infty \leq 2$. We first consider term (i). By Lemma E.4 and noting that the bonus term \widehat{b}_h^k is $O(1)$, we have

$$\begin{aligned}
& \sum_{h=1}^{H-1} \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\widehat{b}_h^k(s, a) \right] \\
& \lesssim \sum_{h=1}^{H-1} \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\min \left\{ \alpha^k \|\widehat{\phi}_h^k(s, a)\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \widehat{\phi}_h^k}^{-1}}, 2 \right\} \right] \\
& \leq \sqrt{A \cdot (\alpha^k)^2 \cdot \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} \left[\|\widehat{\phi}_1^k(s, a)\|_{\Sigma_{\rho_1^k \times U(\mathcal{A}), \widehat{\phi}_1^k}^{-1}}^2 \right]} \\
& \quad + \sum_{h=1}^{H-2} \mathbb{E}_{(s,a) \sim d_{h+1}^{(\pi^k, c^k)}} \left[\alpha^k \|\widehat{\phi}_{h+1}^k(s, a)\|_{\Sigma_{\rho_{h+1}^k \times U(\mathcal{A}), \widehat{\phi}_{h+1}^k}^{-1}} \right] \\
& \leq \sum_{h=1}^{H-2} \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\|\phi_h^*(s, a)\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \sqrt{k(\alpha^k)^2 A \cdot \mathbb{E}_{s \sim \rho_{h+1}^k, a \sim U(\mathcal{A})} \left[\|\widehat{\phi}_{h+1}^k(s, a)\|_{\Sigma_{\rho_{h+1}^k \times U(\mathcal{A}), \widehat{\phi}_{h+1}^k}^{-1}}^2 \right]} + 4\lambda^k d \right] \\
& \quad + \sqrt{A \cdot (\alpha^k)^2 \cdot \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} \left[\|\widehat{\phi}_1^k(s, a)\|_{\Sigma_{\rho_1^k \times U(\mathcal{A}), \widehat{\phi}_1^k}^{-1}}^2 \right]} \\
& \leq \sqrt{\frac{dA(\alpha^k)^2}{k}} + \sqrt{dA(\alpha^k)^2 + 4\lambda^k d} \cdot \sum_{h=1}^{H-2} \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\|\phi_h^*(s, a)\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \right] \tag{15}
\end{aligned}$$

where the last inequality holds by

$$k \cdot \mathbb{E}_{s \sim \rho_h^k, a \sim U(\mathcal{A})} \left[\|\widehat{\phi}_h^k(s, a)\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \widehat{\phi}_h^k}^{-1}}^2 \right] = k \text{Tr} \left(\mathbb{E}_{\rho_h^k \times U(\mathcal{A})} [\widehat{\phi}_h^k (\widehat{\phi}_h^k)^\top] \left\{ k \mathbb{E}_{\rho_h^k \times U(\mathcal{A})} [\widehat{\phi}_h^k (\widehat{\phi}_h^k)^\top] + \lambda^k \right\}^{-1} \right) \leq d$$

Similarly, for term (ii), we have that

$$\begin{aligned}
& \sum_{h=1}^{H-1} \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} [f_h^k(s, a)] \\
& \leq \sqrt{A \cdot \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} \left[(f_1^k(s, a))^2 \right]} \\
& \quad + \sum_{h=1}^{H-2} \mathbb{E}_{(s,a) \sim d_{h+1}^{(\pi^k, c^k)}} \left[\|\phi_{h+1}^*(s, a)\|_{\Sigma_{\rho_{h+1}^k, \phi^*}^{-1}} \sqrt{kA \cdot \mathbb{E}_{s \sim \rho_{h+1}^k, a \sim U(\mathcal{A})} \left[(f_{h+1}^k(s, a))^2 \right]} + 4\lambda^k d \right] \\
& \leq \sqrt{A\zeta_1^k} + \alpha^k \cdot \sum_{h=1}^{H-2} \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\|\phi_h^*(s, a)\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \right], \tag{16}
\end{aligned}$$

where the last step uses

$$\alpha^k \lesssim \sqrt{H^2 A d^2 \log \left(\frac{|\mathcal{F}| H k}{\delta} \right)}.$$

Combining (13), (14), (15), and (16) we obtain

$$\begin{aligned}
& \tau \left(\sum_{k=1}^K (\text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\pi^k, c^k))) \right) \\
& \lesssim \sum_{k=1}^K \left(\sqrt{\frac{dA(\alpha^k)^2}{k}} + \sqrt{H^2 A \zeta^k} \right) + \sum_{k=1}^K \sum_{h=1}^{H-2} \left(2H \cdot \alpha^k \cdot \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\|\phi_h^*(s, a)\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \right] \right)
\end{aligned}$$

$$\begin{aligned}
& + \sum_{k=1}^K \sqrt{dA(\alpha^k)^2 + 4\lambda^k d} \sum_{h=1}^{H-2} \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\|\phi_h^*(s, a)\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \right] \\
& \lesssim H^2 A d^{\frac{3}{2}} \sqrt{\log \left(\frac{|\mathcal{F}|Hk}{\delta} \right)} \sum_{k=1}^K \sum_{h=1}^{H-2} \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\|\phi_h^*(s, a)\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \right] \\
& \lesssim H^3 A d^2 \sqrt{K} \sqrt{\log \left(1 + \frac{K}{d\lambda^1} \right)} \log \left(\frac{|\mathcal{F}|Hk}{\delta} \right)
\end{aligned}$$

where the last inequality holds by (Uehara et al., 2022, Lemma 18), i.e.,

$$\sum_{k=1}^K \mathbb{E}_{(s,a) \sim d_h^{(\pi^k, c^k)}} \left[\|\phi_h^*(s, a)\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \right] \leq \sqrt{dK \log \left(1 + \frac{K}{d\lambda^1} \right)}$$

Therefore, we conclude the proof. \square

Lemma C.2 (Almost Optimism at the Initial Distribution). *Consider an episode $k \in [K]$ and set*

$$\alpha^k = \sqrt{H^2(A + d^2) \log \left(\frac{|\mathcal{F}|Hk}{\delta} \right)}, \lambda^k = O \left(d \log \left(\frac{|\mathcal{F}|Hk}{\delta} \right) \right), \zeta^k = O \left(\frac{1}{k} \log \left(\frac{|\mathcal{F}|Hk}{\delta} \right) \right) \quad (17)$$

with probability $1 - \delta$, we have that

$$V_{1, \widehat{\mathcal{P}}^k, \widehat{b}^k}^{\pi^*}(s_1, c^*) - V_{1, \mathcal{P}^*, 0}^{\pi^*}(s_1, c^*) \leq \sqrt{H^2 A \zeta^k} \quad (18)$$

Proof. Similar to the proof of Lemma C.1, letting $f_h^k(s, a) = \|\widehat{P}_h^k(\cdot | s, a) - P_h^*(\cdot | s, a)\|_1$, we condition on the event that for all $(k, h) \in [K] \times [H]$, the following inequalities hold

$$\begin{aligned}
\mathbb{E}_{s \sim \rho_h^k, a \sim U(\mathcal{A})} \left[(f_h^k(s, a))^2 \right] & \leq \zeta^k, \quad \mathbb{E}_{s \sim \eta_h^k, a \sim U(\mathcal{A})} \left[(f_h^k(s, a))^2 \right] \leq \zeta^k \\
\|\phi(s, a)\|_{(\widehat{\Sigma}_{h, \phi}^k)^{-1}} & = \Theta(\|\phi(s, a)\|_{(\Sigma_{\rho_h^k \times U(\mathcal{A}), \phi}^k)^{-1}})
\end{aligned}$$

From Lemmas E.1 and E.2, this event happens with probability $1 - \delta$. Note that $b_H^k(s, a) = f_H^k(s, a) := 0$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Then, for any policy π , from the simulation Lemma C.3, we have that

$$\begin{aligned}
& V_{1, \widehat{\mathcal{P}}^k, \widehat{b}^k}^{\pi^*}(s_1, c^*) - V_{1, \mathcal{P}^*, 0}^{\pi^*}(s_1, c^*) \\
& = \mathbb{E}_{\pi^*, \widehat{\mathcal{P}}^k} \left[\left(c^* - \sum_{h=1}^H r_h(s_h, a_h) \right)^+ - \sum_{h=1}^H \widehat{b}_h^k(s_h, a_h) \Big| c_1 = c^* \right] - \mathbb{E}_{\pi^*, \mathcal{P}^*} \left[\left(c^* - \sum_{h=1}^H r_h(s_h, a_h) \right)^+ \right] \\
& \leq \sum_{h=1}^{H-1} \mathbb{E}_{(s,a) \sim d_{h, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} \left[H \cdot f_h^k(s, a) - \widehat{b}_h^k(s, a) \right] \quad (19)
\end{aligned}$$

For any $h+1 \in \{2, \dots, H-1\}$, by Lemma E.3 and noting that $\|f_{h+1}\|_\infty \leq 2$, we have that

$$\begin{aligned}
& \left| \mathbb{E}_{(s', a') \sim d_{h+1, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} [f_{h+1}^k(s', a')] \right| \\
& \leq \mathbb{E}_{(s,a) \sim d_{h, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} \left[\|\widehat{\phi}_h^k(s, a)\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \widehat{\phi}_h^k}^{-1}} \cdot \sqrt{kA \cdot \mathbb{E}_{s' \sim \eta_{h+1}^k, a' \sim U(\mathcal{A})} \left[(f_{h+1}^k(s', a'))^2 \right] + 4\lambda^k d + 4k\zeta^k} \right]
\end{aligned}$$

Hence,

$$- \mathbb{E}_{(s', a') \sim d_{h+1, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} [f_{h+1}^k(s', a')] \geq -\sqrt{\beta^k} \cdot \mathbb{E}_{(s,a) \sim d_{h, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} \left[\|\widehat{\phi}_h^k(s, a)\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \widehat{\phi}_h^k}^{-1}} \right] \quad (20)$$

where

$$\beta^k := kA\zeta^k + \lambda^k d + k\zeta^k \lesssim (A + d^2) \log(|\mathcal{F}|Hk/\delta)$$

Combining (19) and (20), we further derive

$$\begin{aligned}
& V_{1, \widehat{\mathcal{P}}^k, \widehat{b}^k}^{\pi^*}(s_1, c^*) - V_{1, \mathcal{P}^*, 0}^{\pi^*}(s_1, c^*) \\
& \leq - \sum_{h=1}^{H-1} \mathbb{E}_{(s,a) \sim d_{h, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} \left[\widehat{b}_h^k(s, a) \right] + \sqrt{H^2 A \zeta^k} + H \cdot \sum_{i=1}^{H-2} \mathbb{E}_{(s,a) \sim d_{i+1, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} \left[f_{i+1}^k(s, a) \right] \\
& \leq - \sum_{h=1}^{H-1} \mathbb{E}_{(s,a) \sim d_{h, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} \left[\widehat{b}_h^k(s, a) \right] + \sum_{h=1}^{H-2} \mathbb{E}_{(s,a) \sim d_{h, \widehat{\mathcal{P}}^k}^{(\pi^*, c^*)}} \left[\min \left\{ \alpha^k \|\widehat{\phi}_h^k(s, a)\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \widehat{\phi}_h^k}^{-1}}, 2 \right\} \right] + \sqrt{H^2 A \zeta^k} \\
& \lesssim \sqrt{H^2 A \zeta^k}
\end{aligned}$$

Therefore, we conclude the proof. \square

Lemma C.3 (Simulation lemma for risk-sensitive RL). *Given two episodic MDPs $(H, \mathcal{P} = \{P_h\}_{h \in [H]}, r)$ and $(H, \widehat{\mathcal{P}} = \{\widehat{P}_h\}_{h \in [H]}, r)$, for any fixed $c \in [0, 1]$ and policy $\pi = \{\pi_h : \mathcal{S} \times [0, H] \mapsto \Delta(\mathcal{A})\}_{h \in [H]}$, we have that*

$$V_{1, \mathcal{P}, 0}^{\pi}(s_1, c) - V_{1, \widehat{\mathcal{P}}, 0}^{\pi}(s_1, c) \leq H \cdot \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_{h, \mathcal{P}}^{(\pi, c)}} [f_h(s, a)] \quad (21)$$

where $f_h(s, a) := \|P_h(\cdot | s, a) - \widehat{P}_h(\cdot | s, a)\|_1$ for any $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$.

Proof. By definition, we derive that

$$\begin{aligned}
& V_{1, \mathcal{P}, 0}^{\pi}(s_1, c) - V_{1, \widehat{\mathcal{P}}, 0}^{\pi}(s_1, c) \\
& = \mathbb{E}_{a_1 \sim \pi_1(\cdot | s_1, c), r_1} \left\{ \mathbb{E}_{s' \sim P_1(\cdot | s_1, a_1)} [V_{2, \mathcal{P}, 0}^{\pi}(s', c - r_1)] - \mathbb{E}_{s' \sim \widehat{P}_1(\cdot | s_1, a_1)} [V_{2, \widehat{\mathcal{P}}, 0}^{\pi}(s', c - r_1)] \right\} \\
& = \mathbb{E}_{a_1 \sim \pi_1(\cdot | s_1, c), r_1} \left\{ \mathbb{E}_{s' \sim \widehat{P}_1(\cdot | s_1, a_1)} \left[V_{2, \mathcal{P}, 0}^{\pi}(s', c - r_1) - V_{2, \widehat{\mathcal{P}}, 0}^{\pi}(s', c - r_1) \right] \right. \\
& \quad \left. + \mathbb{E}_{s' \sim P_1(\cdot | s_1, a_1)} [V_{2, \mathcal{P}, 0}^{\pi}(s', c - r_1)] - \mathbb{E}_{s' \sim \widehat{P}_1(\cdot | s_1, a_1)} [V_{2, \mathcal{P}, 0}^{\pi}(s', c - r_1)] \right\} \\
& \leq \mathbb{E}_{a_1 \sim \pi_1(\cdot | s_1, c), r_1} \mathbb{E}_{s' \sim \widehat{P}_1(\cdot | s_1, a_1)} \left[V_{2, \mathcal{P}, 0}^{\pi}(s', c - r_1) - V_{2, \widehat{\mathcal{P}}, 0}^{\pi}(s', c - r_1) \right] + H \cdot \mathbb{E}_{(s,a) \sim d_{1, \mathcal{P}}^{\pi}} [f_1(s, a)] \\
& \leq \dots \leq H \cdot \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_{h, \mathcal{P}}^{(\pi, c)}} [f_h(s, a)]
\end{aligned}$$

which concludes the proof. \square

D PROOFS FOR APPENDIX B

D.1 PROOF OF THEOREM B.2

In the following discussion we use $\bar{\phi}_h^j$ to denote $\bar{\phi}_h(s_h^j, a_h^j)$ and $(\widehat{\mathcal{P}}_h \bar{V})(s, iv, a)$ to denote

$$\mathbb{E}_{r \sim \bar{r}_h(\cdot | s, a), s' \sim \widehat{P}_h(\cdot | s, a)} [\bar{V}(s', iv - r)].$$

We also use $\mathcal{G}_{t, h}$ denote the filtration generated by $\{(s_{h'}^j, a_{h'}^j, \bar{r}_{h'}^j)_{h'=1}^H\}_{j=1}^{t-1} \cup (s_{h'}^t, a_{h'}^t, \bar{r}_{h'}^t)_{h'=1}^{h-1} \cup (s_h^t, a_h^t)$.

First, note that we have the following concentration lemma:

Lemma D.1. *For all $h \in [H]$, $t \in [T_1]$, $0 \leq i \leq \lceil H/v \rceil$, with probability at $1 - \delta/4$, we have*

$$\left\| \sum_{j=1}^{t-1} \bar{\phi}_h^j \left[\bar{V}_{h+1}^t(s_{h+1}^j, iv - \bar{r}_h^j) - \widehat{\mathcal{P}}_h \bar{V}_{h+1}^t(s_h^j, iv, a_h^j) \right] \right\|_{(\Lambda_h^t)^{-1}} \leq \tilde{O} \left(\frac{H^{\frac{3}{2}} d}{v} \sqrt{\log \frac{H d T_1}{v \delta}} \right).$$

The proof is deferred to Appendix D.3. Let \mathcal{E}_1 denote the event in Lemma D.1.

On the other hand, we can further bound the difference between $\overline{Q}_h^\pi(s, iv, a)$ and $-\widehat{b}_h(s, a) + (\overline{\phi}_h(s, a))^\top w_h^t(iv)$ as follows:

Lemma D.2. *For any policy π , conditioned on \mathcal{E}_1 , we have for all $s \in \mathcal{S}, 0 \leq i \leq \lceil H/v \rceil, a \in \mathcal{A}, h \in [H], t \in [T_1]$ that*

$$-\widehat{b}_h(s, a) + (\overline{\phi}_h(s, a))^\top w_h^t(iv) - \overline{Q}_h^\pi(s, iv, a) = \left(\widehat{\mathcal{P}}_h \left(\overline{V}_{h+1}^t - \overline{V}_{h+1}^\pi \right) \right) (s, iv, a) + \xi_h^t(s, iv, a),$$

where $|\xi_h^t(s, iv, a)| \leq \beta \|\overline{\phi}_h(s, a)\|_{(\Lambda_h^t)^{-1}}$.

The proof of Lemma D.2 follows the same arguments in the proof of Jin et al. (2020)[Lemma B.4] and thus is omitted here.

With Lemma D.2, we can prove that the estimated Q function \overline{Q}^t is optimistic:

Lemma D.3. *Conditioned on event \mathcal{E}_1 , we have for all $s \in \mathcal{S}, 0 \leq i \leq \lceil H/v \rceil, h \in [H], t \in [T_1]$ that $\overline{V}_h^t(s, iv) \leq \overline{V}_h^*(s, iv)$ where $\overline{V}_h^*(s, iv) = \sup_\pi \overline{V}_h^\pi(s, iv)$.*

The proof is deferred to Appendix D.4.

Combining Lemma D.2 and Lemma D.3, we can bound the regret of Algorithm 2 as follows:

Lemma D.4. *With probability at least $1 - \delta/2$, we have*

$$\sum_{t=1}^{T_1} \overline{V}_1^{\tilde{\pi}^t}(s_1, i_1 v) - \overline{V}_1^*(s_1, i_1 v) \leq \tilde{O} \left(\sqrt{\frac{H^5 d^3 T_1 \iota}{v^3}} \right),$$

where $\iota = \log^2 \frac{H d T_1}{v \delta}$.

The proof is deferred to Appendix D.5. Denote the event in Lemma D.4 by \mathcal{E}_2 .

Lemma D.4 implies that by setting $T_1 = \tilde{O} \left(\frac{H^5 d^3 \iota}{v^3 \varepsilon^2} \right)$, we have conditioned on event \mathcal{E}_2 that

$$\min_{t \in [T_1]} \overline{V}_1^{\tilde{\pi}^t}(s_1, i_1 v) - \overline{V}_1^*(s_1, i_1 v) \leq \varepsilon/2.$$

Let t_1 denote $\arg \min_{t \in [T_1]} \overline{V}_1^{\tilde{\pi}^t}(s_1, i_1 v)$. Then on the other hand, by setting $T_2 = \tilde{O} \left(\frac{H^2 \log \frac{T_1}{\delta}}{\varepsilon^2} \right)$, with probability at least $1 - \delta/2$ we have that for all $t \in [T_1]$,

$$\left| \widehat{\overline{V}}_1^{\tilde{\pi}^t}(s_1, i_1 v) - \overline{V}_1^{\tilde{\pi}^t}(s_1, i_1 v) \right| \leq \varepsilon/4.$$

Denote the above event by \mathcal{E}_3 and let $\widehat{\tilde{\pi}}$ denote $\arg \min_{\tilde{\pi}^t} \widehat{\overline{V}}_1^{\tilde{\pi}^t}(s_1, i_1 v)$. Then conditioned on event $\mathcal{E}_2 \cap \mathcal{E}_3$, we have

$$\widehat{\overline{V}}_1^{\widehat{\tilde{\pi}}}(s_1, i_1 v) - \overline{V}_1^*(s_1, i_1 v) \leq \widehat{\overline{V}}_1^{\tilde{\pi}^{t_1}}(s_1, i_1 v) - \overline{V}_1^*(s_1, i_1 v) \leq \overline{V}_1^{\tilde{\pi}^{t_1}}(s_1, i_1 v) - \overline{V}_1^*(s_1, i_1 v) + \frac{\varepsilon}{4} \leq \frac{3}{4} \varepsilon.$$

D.2 PROOF OF THEOREM B.3

First, we show that Algorithm 3 can achieve sublinear regret in the discretized MDP $\overline{\mathcal{M}}$. Let $(\pi_{\text{dis}}^*, i^*) := \arg \max_{\pi \in \overline{\Pi}_{\text{Aug}}, 0 \leq i \leq \lceil H/v \rceil} \text{CVaR}_\tau(\overline{R}(\pi, iv))$. Note that from (4) we have

$$\overline{\text{CVaR}}_\tau^* = i^* v - \tau^{-1} \overline{V}_{1, \mathcal{P}^*, 0}^{\pi_{\text{dis}}^*}(s_1, i^* v).$$

This implies that

$$\overline{\text{CVaR}}_\tau^* - \text{CVaR}_\tau(\overline{R}(\pi^k, i^k v))$$

$$\begin{aligned}
&= \left(i^* v - \tau^{-1} \overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^* \text{dis}}(s_1, i^* v) \right) - \text{CVaR}_\tau(\overline{R}(\pi^k, i^k v)) + \left(i^* v - \tau^{-1} \overline{V}_{1, \mathcal{P}^*, 0}^{\pi^* \text{dis}}(s_1, i^* v) \right) \\
&\quad - \left(i^* v - \tau^{-1} \overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^* \text{dis}}(s_1, i^* v) \right) \\
&= \left(i^* v - \tau^{-1} \overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^* \text{dis}}(s_1, i^* v) \right) - \text{CVaR}_\tau(\overline{R}(\pi^k, i^k v)) + \tau^{-1} \left(\overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^* \text{dis}}(s_1, i^* v) - \overline{V}_{1, \mathcal{P}^*, 0}^{\pi^* \text{dis}}(s_1, i^* v) \right)
\end{aligned} \tag{22}$$

Suppose in k -th iteration, *CVaR-LSVI* returns $\widehat{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^*$ for the initial budget $i^* v$. Then from Theorem B.2, we know

$$\overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^* \text{dis}}(s_1, i^* v) \geq \widehat{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^* - \frac{\tau}{8} \epsilon.$$

This implies that

$$\begin{aligned}
i^* v - \tau^{-1} \overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^* \text{dis}}(s_1, i^* v) &\leq i^* v - \tau^{-1} \widehat{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^* + \frac{\epsilon}{8} \\
&\leq i^k - \tau^{-1} \widehat{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^* + \frac{\epsilon}{8} \\
&\leq i^k - \tau^{-1} \overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^k} + \frac{\epsilon}{6},
\end{aligned} \tag{23}$$

where the last step is due to Theorem 5.1.

On the other hand, from (2) we know

$$\text{CVaR}_\tau(\overline{R}(\pi^k, i^k v)) \geq i^k - \overline{V}_{1, \mathcal{P}^*, 0}^{\pi^k}(s_1, i^k v). \tag{24}$$

Plug Equation (23) and (24) into (22), we have

$$\begin{aligned}
\overline{\text{CVaR}}_\tau^* - \text{CVaR}_\tau(\overline{R}(\pi^k, i^k v)) &\leq \tau^{-1} \left(\overline{V}_{1, \mathcal{P}^*, 0}^{\pi^k}(s_1, i^k v) - \overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^k}(s_1, i^k v) \right) \\
&\quad + \tau^{-1} \left(\overline{V}_{1, \hat{\mathcal{P}}^k, \hat{b}^k}^{\pi^* \text{dis}}(s_1, i^* v) - \overline{V}_{1, \mathcal{P}^*, 0}^{\pi^* \text{dis}}(s_1, i^* v) \right) + \frac{\epsilon}{6}
\end{aligned}$$

Then, following the same arguments in the proof of Theorem 4.1, we can obtain

$$\sum_{k=1}^K \overline{\text{CVaR}}_\tau^* - \text{CVaR}_\tau(\overline{R}(\pi^k, i^k v)) \leq \tilde{O} \left(\tau^{-1} H^3 A d^2 \sqrt{K} \cdot \sqrt{\log(|\mathcal{F}|/\delta)} \right) + \frac{K\epsilon}{6}. \tag{25}$$

Next we bridge $\overline{\text{CVaR}}_\tau^*$ and CVaR_τ via Proposition B.1. Note that from Proposition B.1 we have

$$\sum_{k=1}^K \text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\overline{\pi}^k, i^k v)) \leq \sum_{k=1}^K \overline{\text{CVaR}}_\tau^* - \text{CVaR}_\tau(\overline{R}(\pi^k, i^k v)) + \frac{KHv}{\tau}. \tag{26}$$

Combining (25) and (26), we have

$$\sum_{k=1}^K \text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\overline{\pi}^k, i^k v)) \leq \tilde{O} \left(\tau^{-1} H^3 A d^2 \sqrt{K} \cdot \sqrt{\log(|\mathcal{F}|/\delta)} \right) + \frac{K\epsilon}{6} + \frac{KHv}{\tau}.$$

Substituting the values of the parameters, we will obtain

$$\frac{1}{K} \sum_{k=1}^K \text{CVaR}_\tau^* - \text{CVaR}_\tau(R(\overline{\pi}^k, i^k v)) \leq \epsilon.$$

This implies that the uniformly mixed policy of $\{(\overline{\pi}^k, i^k v)\}_{k=1}^K$ is ϵ -optimal. The total number of collected trajectories is $KH = \tilde{O} \left(\frac{H^7 A^2 d^4 \log \frac{|\mathcal{F}|}{\delta}}{\tau^2 \epsilon^2} \right)$.

Now, we discuss the computational complexity. The MLE oracle is called KH times. For the rest of the computation, the cost is dominated by running $CVaR$ - $LSVI$ in Line 17 of Algorithm 3. In $CVaR$ - $LSVI$, we compute $(\Lambda_h^t)^{-1}$ by the Sherman-Morrison formula, then the computational complexity of $CVaR$ - $LSVI$ is dominated by Line 6 of Algorithm 2, which requires $O(\frac{d^2 AT_1}{v^2})$ time per step and leads to a total computational complexity of $O(\frac{d^2 AH^2 T_1^2}{v^3})$ for each call of $CVaR$ - $LSVI$. Note that $CVaR$ - $LSVI$ is called totally $\frac{KH}{v}$ times in Algorithm 3 and thus the total computation cost of Algorithm 3 is

$$O\left(\frac{Kd^2 AH^3 T_1^2}{v^4}\right) = \tilde{O}\left(\frac{H^{29} A^3 d^{12} \log \frac{|\mathcal{F}|}{\delta}}{\tau^{16} \epsilon^{16}}\right).$$

D.3 PROOF OF LEMMA D.1

The proof largely follows the proof of Jin et al. (2020)[Lemma B.3]. We first bound $w_h^t(iv)$. Note that we have for any vector $v \in \mathbb{R}^{\bar{d}}$ where $\bar{d} := d(1 + \lceil 1/v \rceil)$ that

$$\begin{aligned} |v^\top w_h^t(iv)| &\leq \sum_{j=1}^{t-1} \left| v^\top (\Lambda_h^t)^{-1} \bar{\phi}_h^j \right| \cdot H \\ &\leq \sqrt{\left(\sum_{j=1}^{t-1} v^\top (\Lambda_h^t)^{-1} v \right) \left(\sum_{j=1}^{t-1} (\bar{\phi}_h^j)^\top (\Lambda_h^t)^{-1} \bar{\phi}_h^j \right)} \cdot H \\ &\leq H \|v\| \sqrt{\frac{t}{\lambda}} \cdot \sqrt{\sum_{j=1}^{t-1} (\bar{\phi}_h^j)^\top (\Lambda_h^t)^{-1} \bar{\phi}_h^j}. \end{aligned}$$

Note that with Lemma E.6 we have $\sum_{j=1}^{t-1} (\bar{\phi}_h^j)^\top (\Lambda_h^t)^{-1} \bar{\phi}_h^j \leq \bar{d}$ and therefore we can obtain that for all $v \in \mathbb{R}^{\bar{d}}$

$$|v^\top w_h^t(iv)| \leq H \|v\| \sqrt{\frac{t\bar{d}}{\lambda}},$$

which implies that $\|w_h^t(iv)\| \leq H \sqrt{\frac{t\bar{d}}{\lambda}}$.

To utilize Lemma E.7, we further need to bound the covering number of the class of estimated value functions. Fix $h \in [H]$, it can be observed that all $\bar{V}_h^t(\cdot, \cdot)$ where $t \in [T_1]$ belongs to the following function class \mathcal{V} parametrized by w and A :

$$\left\{ V \mid V(s, iv) = \min_a \text{Clip}_{[-H, H]} \left(\bar{\phi}_h^\top(s, a) w(iv) - \hat{b}_h(s, a) - \|\bar{\phi}_h(s, a)\|_A \right) \right\},$$

where $\|w(iv)\| \leq H \sqrt{\frac{t\bar{d}}{\lambda}}$ and $\|A\| \leq \beta^2/\lambda$. Then for any V_1 (parametrized by w_1, A_1) and V_2 (parametrized by w_2, A_2), we know

$$\begin{aligned} &\sup_{s \in \mathcal{S}, 0 \leq i \leq \lceil H/v \rceil} |V_1(s, iv) - V_2(s, iv)| \\ &\leq \sup_{s \in \mathcal{S}, a \in \mathcal{A}, 0 \leq i \leq \lceil H/v \rceil} \left| \bar{\phi}_h^\top(s, a) (w_1(iv) - w_2(iv)) - (\|\bar{\phi}_h(s, a)\|_{A_1} - \|\bar{\phi}_h(s, a)\|_{A_2}) \right| \\ &\leq \sup_{\|\bar{\phi}\| \leq 1, 0 \leq i \leq \lceil H/v \rceil} \left| \bar{\phi}^\top (w_1(iv) - w_2(iv)) \right| + \sup_{\|\bar{\phi}\| \leq 1} \|\bar{\phi}\|_{A_1 - A_2} \\ &\leq \sup_{0 \leq i \leq \lceil H/v \rceil} \|w_1(iv) - w_2(iv)\| + \sqrt{\|A_1 - A_2\|_F}. \end{aligned}$$

This indicates that the covering number of \mathcal{V} with respect to ℓ_∞ norm can be upper bounded by the covering number of w and A . More specifically, let $\mathcal{N}(\epsilon)$ denote the covering number of \mathcal{V} with

respect to ℓ_∞ norm, then we have

$$\log \mathcal{N}(\epsilon) \leq O\left(\bar{d}\left(\bar{d} + \frac{H}{v}\right) \log \frac{HT_1\bar{d}}{\epsilon\lambda}\right) \leq O\left(\frac{Hd^2}{v^2} \log \frac{HT_1d}{\epsilon\lambda v}\right).$$

Substituting the above bound into Lemma E.7 concludes the proof.

D.4 PROOF OF LEMMA D.3

The proof is conducted via induction. For the base case, when $h = H + 1$, we have $\bar{V}_{H+1}^t(s, iv) = \bar{V}_{H+1}^*(s, iv) = iv$ for all $s \in \mathcal{S}$ and $0 \leq i \leq \lceil H/v \rceil$.

Then suppose we have $\bar{V}_{h+1}^t(s, iv) \leq \bar{V}_{h+1}^*(s, iv)$ for all $s \in \mathcal{S}$ and $0 \leq i \leq \lceil H/v \rceil$. For any $s \in \mathcal{S}, a \in \mathcal{A}$ and $0 \leq i \leq \lceil H/v \rceil$, if $\bar{Q}_h^t(s, iv, a) = -H$, then $\bar{Q}_h^t(s, iv, a) \leq \bar{Q}_h^*(s, iv, a)$ naturally holds. Otherwise, from Lemma D.2 we know

$$\begin{aligned} \bar{Q}_h^t(s, iv, a) - \bar{Q}_h^*(s, iv, a) &\leq -\hat{b}_h(s, a) + (\bar{\phi}_h(s, a))^\top w_h^t(iv) - \beta \|\bar{\phi}_h(s, a)\|_{(\Lambda_h^t)^{-1}} - \bar{Q}_h^*(s, iv, a) \\ &\leq \left(\hat{\mathcal{P}}_h \left(\bar{V}_{h+1}^t - \bar{V}_{h+1}^\pi\right)\right)(s, iv, a) \leq 0. \end{aligned}$$

Therefore we have $\bar{Q}_h^t(s, iv, a) \leq \bar{Q}_h^*(s, iv, a)$ for all $s \in \mathcal{S}, a \in \mathcal{A}$ and $0 \leq i \leq \lceil H/v \rceil$, which leads to $\bar{V}_h^t(s, iv) \leq \bar{V}_h^*(s, iv)$ for all $s \in \mathcal{S}$ and $0 \leq i \leq \lceil H/v \rceil$ as well. This concludes our proof.

D.5 PROOF OF LEMMA D.4

First note that from Lemma D.3 we have $\bar{V}_1^*(s_1, i_1v) \geq \bar{V}_1^t(s_1, i_1v)$ for all $t \in [T_1]$. This indicates that conditioned on \mathcal{E}_1 ,

$$\sum_{t=1}^{T_1} \bar{V}_1^{\tilde{\pi}^t}(s_1, i_1v) - \bar{V}_1^*(s_1, i_1v) \leq \sum_{t=1}^{T_1} \bar{V}_1^{\tilde{\pi}^t}(s_1, i_1v) - \bar{V}_1^t(s_1, i_1v). \quad (27)$$

Fix $t \in [T_1]$ and then we know

$$\bar{V}_1^{\tilde{\pi}^t}(s_1, i_1v) - \bar{V}_1^t(s_1, i_1v) = \bar{Q}_1^{\tilde{\pi}^t}(s_1^t, i_1v, a_1^t) - \bar{Q}_1^t(s_1, i_1v, a_1^t).$$

If $\bar{Q}_1^t(s_1^t, i_1v, a_1^t) = H$, then we know $\bar{Q}_1^{\tilde{\pi}^t}(s_1^t, i_1v, a_1^t) - \bar{Q}_1^t(s_1, i_1v, a_1^t) \leq 0$. Otherwise, we have $\bar{Q}_1^{\tilde{\pi}^t}(s_1^t, i_1v, a_1^t) - \bar{Q}_1^t(s_1, i_1v, a_1^t) \leq \bar{Q}_1^{\tilde{\pi}^t}(s_1^t, i_1v, a_1^t) + \hat{b}_h(s_1^t, a_1^t) - \langle \bar{\phi}_1^t, w_h^t(i_1v) \rangle + \beta \|\bar{\phi}_1^t\|_{(\Lambda_1^t)^{-1}}$.

Then with Lemma D.2, we know

$$\bar{Q}_1^{\tilde{\pi}^t}(s_1^t, i_1v, a_1^t) - \bar{Q}_1^t(s_1, i_1v, a_1^t) \leq \left(\hat{\mathcal{P}}_h \left(\bar{V}_2^{\tilde{\pi}^t} - \bar{V}_2^t\right)\right)(s_1^t, i_1v, a_1^t) + 2\beta \|\bar{\phi}_1^t\|_{(\Lambda_1^t)^{-1}}. \quad (28)$$

Let $i_h^t v$ denote $i_1v - \sum_{h'=1}^{h-1} \bar{r}_{h'}^t$ and let $\{(\zeta_h^t)_{h=1}^H\}_{t=1}^{T_1}$ denote the following random variable:

$$\zeta_h^t := \begin{cases} 0, & \text{if there exists some } h' \leq h \text{ s.t. } \bar{Q}_{h'}^t(s_{h'}^t, i_{h'}^t v, a_{h'}^t) = H, \\ \left(\hat{\mathcal{P}}_h \left(\bar{V}_{h+1}^{\tilde{\pi}^t} - \bar{V}_{h+1}^t\right)\right)(s_h^t, i_h^t v, a_h^t) - \left(\bar{V}_{h+1}^{\tilde{\pi}^t}(s_{h+1}^t, i_{h+1}^t v) - \bar{V}_2^t(s_{h+1}^t, i_{h+1}^t v)\right), & \text{otherwise.} \end{cases}$$

It can be easily observed that $\zeta_h^t \in \mathcal{G}_{t, h+1}$ and $\mathbb{E}[\zeta_h^t | \mathcal{G}_{t, h}] = 0$, which means that $\{(\zeta_h^t)_{h=1}^H\}_{t=1}^{T_1}$ is a martingale with respect to $\{\mathcal{G}_{t, h}\}$. Then when $\bar{Q}_1^t(s_1^t, i_1v, a_1^t) \neq H$, Equation (28) is equivalent to

$$\bar{Q}_1^{\tilde{\pi}^t}(s_1^t, i_1v, a_1^t) - \bar{Q}_1^t(s_1, i_1v, a_1^t) \leq \bar{V}_2^{\tilde{\pi}^t}(s_2^t, i_2^t v) - \bar{V}_2^t(s_2^t, i_2^t v) + \zeta_h^t + 2\beta \|\bar{\phi}_1^t\|_{(\Lambda_1^t)^{-1}}.$$

Repeat such expansion till a step h_t such that $\bar{Q}_{h_t}^t(s_{h_t}^t, i_{h_t}^t v, a_{h_t}^t) = H$ or $h_t = H + 1$. Then we have

$$\bar{V}_1^{\tilde{\pi}^t}(s_1, i_1v) - \bar{V}_1^t(s_1, i_1v) \leq \sum_{h=1}^H \zeta_h^t + 2\beta \sum_{h=1}^{h_t-1} \|\bar{\phi}_h^t\|_{(\Lambda_h^t)^{-1}} \leq \sum_{h=1}^H \zeta_h^t + 2\beta \sum_{h=1}^H \|\bar{\phi}_h^t\|_{(\Lambda_h^t)^{-1}}. \quad (29)$$

Therefore combining (28) and (29), we know

$$\sum_{t=1}^{T_1} \bar{V}_1^{\bar{\pi}^t}(s_1, i_1 v) - \bar{V}_1^*(s_1, i_1 v) \leq \sum_{t=1}^{T_1} \sum_{h=1}^H \zeta_h^t + 2\beta \sum_{t=1}^{T_1} \sum_{h=1}^H \|\bar{\phi}_h^t\|_{(\Lambda_h^t)^{-1}}. \quad (30)$$

For the first term of the RHS of (30), from Azuma-Hoeffding's inequality, we have with probability at least $1 - \delta/2$ that

$$\sum_{t=1}^{T_1} \sum_{h=1}^H \zeta_h^t \leq 2H \sqrt{T_1 H \ell^{\frac{1}{2}}}.$$

For the second term, we can utilize the elliptical potential lemma (Lemma E.8 in Appendix E), which yields

$$\sum_{t=1}^{T_1} \sum_{h=1}^H \|\bar{\phi}_h^t\|_{(\Lambda_h^t)^{-1}} \leq \sum_{h=1}^H \sqrt{T_1} \cdot \sqrt{\sum_{t=1}^{T_1} \sum_{h=1}^H (\bar{\phi}_h^t)^\top (\Lambda_h^t)^{-1} \bar{\phi}_h^t} \leq H \sqrt{\frac{2dT_1 \ell^{\frac{1}{2}}}{v}}.$$

Plug the above two inequalities into (30) leads to the result in Lemma D.4.

E AUXILIARY LEMMAS

Lemma E.1 (MLE guarantees). *For any fixed $(h, k) \in [H] \times [K]$, with probability at least $1 - \delta$, we have that*

$$\mathbb{E}_{s \sim \{0.5\rho_h^k + 0.5\eta_h^k\}, a \sim U(\mathcal{A})} [\|\hat{P}_h^k(\cdot|s, a) - P_h^*(\cdot|s, a)\|_1^2] \leq \zeta, \quad \zeta := \frac{\log(|\mathcal{F}|/\delta)}{k}. \quad (31)$$

Using union bound, a direct corollary is: with probability at least $1 - \delta$ the following holds for all $h \in [H], k \in [K]$

$$\mathbb{E}_{s \sim \{0.5\rho_h^k + 0.5\eta_h^k\}, a \sim U(\mathcal{A})} [\|\hat{P}_h^k(\cdot|s, a) - P_h^*(\cdot|s, a)\|_1^2] \leq 0.5\zeta^k, \quad \zeta^k := \frac{\log(|\mathcal{F}|Hk/\delta)}{k}. \quad (32)$$

Lemma E.2 (Concentration of the bonus term). *Set $\lambda^k = \Theta(d \log(kH|\mathcal{F}|/\delta))$ at the k -th episode. Define*

$$\Sigma_{\rho_h^k, \phi} = k \mathbb{E}_{s \sim \rho_h^k, a \sim U(\mathcal{A})} [\phi(s, a) \phi^\top(s, a)] + \lambda^k I, \quad \hat{\Sigma}_{k, \phi} = \sum_{s, a \in \mathcal{D}_h} \phi(s, a) \phi^\top(s, a) + \lambda^k I.$$

With probability $1 - \delta$, we have for any $k \in \mathbb{N}^+, h \in [H], \phi \in \Phi$,

$$\|\phi(s, a)\|_{(\hat{\Sigma}_{k, \phi})^{-1}} = \Theta \left(\|\phi(s, a)\|_{\Sigma_{\rho_h^k, \phi}^{-1}} \right)$$

Lemma E.3 (One-step back inequality for the learned model). *Let $\pi = \{\pi_h : \mathcal{S} \times [0, H] \mapsto \Delta(\mathcal{A})\}_{h \in [H]}$ denote any policy on the augmented state. Fix an initial budget $c_1 \in [0, 1]$. Take any $g : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ such that $\|g\|_\infty \leq B$. We condition on the event where the MLE guarantee (32):*

$$\mathbb{E}_{s \sim \rho_h^k, a \sim U(\mathcal{A})} [f_h^k(s, a)] \lesssim \zeta^k$$

holds for any $h \in [H]$. For $h = 1$, we have that

$$\left| \mathbb{E}_{(s, a) \sim d_{1, \hat{\mathcal{P}}^k}^{(\pi, c_1)}} [g(s, a)] \right| \leq \sqrt{A \cdot \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} [g^2(s, a)]}$$

For any $h + 1 \in \{2, \dots, H\}$ and policy π , we have that

$$\begin{aligned} & \left| \mathbb{E}_{(s', a') \sim d_{h+1, \hat{\mathcal{P}}^k}^{(\pi, c_1)}} [g(s', a')] \right| \\ & \leq \mathbb{E}_{(s, a) \sim d_{h, \hat{\mathcal{P}}^k}^{(\pi, c_1)}} \left[\left\| \hat{\phi}_h^k(s, a) \right\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \hat{\phi}_h^k}^{-1}} \sqrt{kA \cdot \mathbb{E}_{s' \sim \eta_{h+1}^k, a' \sim U(\mathcal{A})} [g^2(s', a')] + B^2 \lambda^k d + kB^2 \zeta^k} \right] \end{aligned} \quad (33)$$

where $\Sigma_{\rho_h^k \times U(\mathcal{A}), \hat{\phi}_h^k} = k \mathbb{E}_{s \sim \rho_h^k, a \sim U(\mathcal{A})} [\hat{\phi}_h^k (\hat{\phi}_h^k)^\top] + \lambda^k I$.

Proof. For $h = 1$, we have that

$$\begin{aligned} \mathbb{E}_{(s,a) \sim d_{1,\hat{\rho}^k}^{(\pi,c_1)}} [g(s,a)] &\leq \sqrt{\max_{s,a} \frac{d_1(s)\pi_1(a|s,c_1)}{\rho_1^k(s)U(a)} \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} [g^2(s,a)]} \\ &\leq \sqrt{A \cdot \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} [g^2(s,a)]} \end{aligned}$$

where we use the fact that $d_1 = \rho_1^k$. Recall that $\rho_h^k(s) = \frac{1}{k} \sum_{i=0}^{k-1} d_{h,c^i}^{\pi^i}(s)$ is the (expected) occupancy of any $s \in \mathcal{S}$ in dataset \mathcal{D}_h . Let $\omega_{h,\mathcal{P}}^{(\pi,c_1)} : \mathcal{S} \times \mathcal{A} \mapsto \Delta([0,1])$ denote the distribution of the remaining budget c_h at timestep h conditioned on any $(s,a) \in \mathcal{S} \times \mathcal{A}$ when rolling policy π in an MDP with transition kernel \mathcal{P} and the initial budget c_1 . For any $h \in \{1, \dots, H-1\}$, we have that

$$\begin{aligned} &\mathbb{E}_{(s',a') \sim d_{h+1,\hat{\rho}^k}^{(\pi,c_1)}} [g(s',a')] \\ &= \mathbb{E}_{(s,a) \sim d_{h,\hat{\rho}^k}^{(\pi,c_1)}} \mathbb{E}_{s' \sim \hat{P}_h^k(\cdot|s,a)} \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] \\ &= \mathbb{E}_{(s,a) \sim d_{h,\hat{\rho}^k}^{(\pi,c_1)}} \left[\hat{\phi}_h^k(s,a)^\top \int_{\mathcal{S}} \hat{\psi}_h^k(s') \cdot \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] ds' \right] \\ &\leq \mathbb{E}_{(s,a) \sim d_{h,\hat{\rho}^k}^{(\pi,c_1)}} \left[\left\| \hat{\phi}_h^k(s,a) \right\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \hat{\phi}_h^k}^{-1}} \right. \\ &\quad \left. \cdot \left\| \int_{\mathcal{S}} \hat{\psi}_h^k(s') \cdot \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] ds' \right\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \hat{\phi}_h^k}} \right] \end{aligned}$$

where the last inequality is because of CS inequality. For any $(s,a) \in \mathcal{S} \times \mathcal{A}$, we have that

$$\begin{aligned} &\left\| \int_{\mathcal{S}} \hat{\psi}_h^k(s') \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] ds' \right\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \hat{\phi}_h^k}}^2 \\ &\leq \left(\int_{\mathcal{S}} \hat{\psi}_h^k(s') \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] ds' \right)^\top \cdot \left\{ k \mathbb{E}_{\tilde{s} \sim \rho_h^k, \tilde{a} \sim U(\mathcal{A})} [\hat{\phi}_h^k(\tilde{s}, \tilde{a}) (\hat{\phi}_h^k)^\top] + \lambda^k I \right\} \\ &\quad \cdot \left(\int_{\mathcal{S}} \hat{\psi}_h^k(s') \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] ds' \right) \\ &\leq k \cdot \mathbb{E}_{\tilde{s} \sim \rho_h^k, \tilde{a} \sim U(\mathcal{A})} \left[\left(\int_{\mathcal{S}} \hat{\psi}_h^k(s')^\top \hat{\phi}_h^k(\tilde{s}, \tilde{a}) \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] ds' \right)^2 \right] + B^2 \lambda^k d \end{aligned}$$

where we use the assumption

$$\left\| \sum_{a'} \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] \right\| \leq B,$$

and

$$\int_{\mathcal{S}} \|\hat{\psi}_h^k(s') y(s')\|_2 ds' \leq \sqrt{d},$$

for any $y : \mathcal{S} \mapsto [0,1]$. Note that $\hat{\psi}_h^k(s') \hat{\phi}_h^k(\tilde{s}, \tilde{a}) = \hat{P}_h^k(s'|\tilde{s}, \tilde{a})$. Further, we derive that

$$\begin{aligned} &k \cdot \mathbb{E}_{\tilde{s} \sim \rho_h^k, \tilde{a} \sim U(\mathcal{A})} \left[\left(\int_{\mathcal{S}} \hat{P}_h^k(s'|\tilde{s}, \tilde{a}) \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] ds' \right)^2 \right] + B^2 \lambda^k d \\ &\leq k \cdot \mathbb{E}_{\tilde{s} \sim \rho_h^k, \tilde{a} \sim U(\mathcal{A})} \left[\left(\mathbb{E}_{s' \sim P_h^*(\cdot|\tilde{s}, \tilde{a})} \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')] \right)^2 \right] + B^2 \lambda^k d + kB^2 \zeta^k \\ &\leq k \cdot \mathbb{E}_{\tilde{s} \sim \rho_h^k, \tilde{a} \sim U(\mathcal{A}), s' \sim P_h^*(\cdot|\tilde{s}, \tilde{a})} \mathbb{E}_{c \sim \omega_{h,\hat{\rho}^k}^{(\pi,c_1)}(\cdot|s,a), r, a' \sim \pi_{h+1}(\cdot|s',c-r)} [g(s',a')]^2 + B^2 \lambda^k d + kB^2 \zeta^k \\ &\leq kA \cdot \mathbb{E}_{\tilde{s} \sim \rho_h^k, \tilde{a} \sim U(\mathcal{A}), s' \sim P_h^*(\cdot|\tilde{s}, \tilde{a}), a' \sim U(\mathcal{A})} [g^2(s',a')] + B^2 \lambda^k d + kB^2 \zeta^k \end{aligned}$$

Note that $\tilde{s} \sim \rho_h^k$, $\tilde{a} \sim U(\mathcal{A})$, $s' \sim P_h^*(\cdot|\tilde{s}, \tilde{a})$ is equivalent to $s' \sim \eta_{h+1}^k$. Therefore,

$$\begin{aligned} & \mathbb{E}_{(s', a') \sim d_{h+1, \hat{P}^k}^{(\pi, c_1)}} [g(s', a')] \\ & \leq \mathbb{E}_{(s, a) \sim d_{h, \hat{P}^k}^{(\pi, c_1)}} \left[\left\| \hat{\phi}_h^k(s, a) \right\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \hat{\phi}_h^k}^{-1}} \left\| \int_{\mathcal{S}} \sum_{a'} \hat{\psi}_h^k(s') \mathbb{E}_{c, r_h} [\pi_{h+1}(a'|s', c - r_h)] g(s', a') ds' \right\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \hat{\phi}_h^k}} \right] \\ & \leq \mathbb{E}_{(s, a) \sim d_{h, \hat{P}^k}^{(\pi, c_1)}} \left[\left\| \hat{\phi}_h^k(s, a) \right\|_{\Sigma_{\rho_h^k \times U(\mathcal{A}), \hat{\phi}_h^k}^{-1}} \sqrt{kA \cdot \mathbb{E}_{s' \sim \eta_{h+1}^k, a' \sim U(\mathcal{A})} [g^2(s', a')] + B^2 \lambda^k d + kB^2 \zeta^k} \right] \end{aligned}$$

which concludes the proof. \square

Lemma E.4 (One-step back inequality for the true model). *Let $\pi = \{\pi_h : \mathcal{S} \times [0, H] \mapsto \Delta(\mathcal{A})\}_{h \in [H]}$ denote any policy on the augmented state. Fix an initial budget $c_1 \in [0, 1]$. Take any $g : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ such that $\|g\|_\infty \leq B$. Then for $h = 1$, we have that*

$$\mathbb{E}_{(s, a) \sim d_1^{(\pi, c_1)}} [g(s, a)] \leq \sqrt{A \cdot \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} [g^2(s, a)]}$$

For $h + 1 \in \{2, \dots, H\}$, we have that

$$\mathbb{E}_{(s', a') \sim d_{h+1}^{(\pi, c_1)}} [g(s', a')] \leq \mathbb{E}_{(s, a) \sim d_h^{(\pi, c_1)}} \left[\left\| \phi_h^*(s, a) \right\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \sqrt{kA \cdot \mathbb{E}_{s' \sim \rho_{h+1}^k, a' \sim U(\mathcal{A})} [g^2(s', a')] + \lambda^k dB^2} \right]$$

Proof. For $h = 1$, we have that

$$\mathbb{E}_{(s, a) \sim d_1^{(\pi, c_1)}} [g(s, a)] \leq \sqrt{\frac{d_1(s) \pi_1(a|s, c_1)}{\rho_1^k(s) U(a)} \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} [g^2(s, a)]} \leq \sqrt{A \cdot \mathbb{E}_{s \sim \rho_1^k, a \sim U(\mathcal{A})} [g^2(s, a)]}$$

Let $\omega_h^{(\pi, c_1)} := \omega_{h, \hat{P}^*}^{(\pi, c_1)}$ denote the distribution of the remaining budget c_h at timestep h conditioned on any $(s, a) \in \mathcal{S} \times \mathcal{A}$ when rolling policy π in the (true) environment and the initial budget c_1 . For $h + 1 \in \{2, \dots, H\}$, by CS inequality, we derive that

$$\begin{aligned} & \mathbb{E}_{(s', a') \sim d_{h+1}^{(\pi, c_1)}} [g(s', a')] \\ & = \mathbb{E}_{(s, a) \sim d_h^{(\pi, c_1)}, s' \sim P_h^*(\cdot|s, a)} \mathbb{E}_{c \sim \omega_h^{(\pi, c_1)}, r, a' \sim \pi_{h+1}(\cdot|s', c-r)} [g(s', a')] \\ & \leq \mathbb{E}_{(s, a) \sim d_h^{(\pi, c_1)}} \left[\left\| \phi_h^*(s, a) \right\|_{\Sigma_{\rho_h^k, \phi^*}^{-1}} \left\| \int_{\mathcal{S}} \psi_h^*(s') \mathbb{E}_{c \sim \omega_h^{(\pi, c_1)}, r, a' \sim \pi_{h+1}(\cdot|s', c-r)} [g(s', a')] ds' \right\|_{\Sigma_{\rho_h^k, \phi^*}} \right] \end{aligned}$$

For any $(s, a) \in \mathcal{S} \times \mathcal{A}$, we obtain that

$$\begin{aligned} & \left\| \int_{\mathcal{S}} \psi_h^*(s') \mathbb{E}_{c \sim \omega_h^{(\pi, c_1)}, r, a' \sim \pi_{h+1}(\cdot|s', c-r)} [g(s', a')] ds' \right\|_{\Sigma_{\rho_h^k, \phi^*}}^2 \\ & \leq \left\{ \int_{\mathcal{S}} \psi_h^*(s') \mathbb{E}_{c \sim \omega_h^{(\pi, c_1)}, r, a' \sim \pi_{h+1}(\cdot|s', c-r)} [g(s', a')] ds' \right\}^\top \left\{ k \mathbb{E}_{(\tilde{s}, \tilde{a}) \sim \rho_h^k} [\phi_h^*(\phi_h^*)^\top] + \lambda^k I \right\} \\ & \quad \cdot \left\{ \int_{\mathcal{S}} \psi_h^*(s') \mathbb{E}_{c \sim \omega_h^{(\pi, c_1)}, r, a' \sim \pi_{h+1}(\cdot|s', c-r)} [g(s', a')] ds' \right\} \\ & \leq k \mathbb{E}_{(\tilde{s}, \tilde{a}) \sim \rho_h^k, s' \sim P_h^*(\cdot|\tilde{s}, \tilde{a})} \mathbb{E}_{c \sim \omega_h^{(\pi, c_1)}, r, a' \sim \pi_{h+1}(a', s', c-r)} [g^2(s', a')] + \lambda^k dB^2 \\ & \leq kA \cdot \mathbb{E}_{s' \sim \rho_{h+1}^k, a' \sim U(\mathcal{A})} [g^2(s', a')] + \lambda^k dB^2 \end{aligned}$$

where the last inequality holds by the fact that $(\tilde{s}, \tilde{a}) \sim \rho_h^k$, $s' \sim P_h^*(\cdot|\tilde{s}, \tilde{a})$ is equivalent to $s' \sim \rho_{h+1}^k$ and importance sampling. Therefore, we conclude the proof. \square

Lemma E.5 (Simulation lemma for risk-neutral RL). *Let $\widehat{\mathcal{M}}$ and \mathcal{M} be to MDPs with the same state and action spaces, but they have different reward and transition functions, i.e., $(\widehat{r}_h, \widehat{P}_h)_{h=1}^H$ and*

$(r_h, P_h)_{h=1}^H$, respectively. Consider any policy $\pi = \{\pi_h : \mathcal{S} \rightarrow \Delta(\mathcal{A})\}_{h \in [H]}$, denote $\{V_{h, \widehat{\mathcal{M}}}^\pi\}_{h=1}^H$ and $\{V_h^\pi\}_{h=1}^H$ be the value functions. Then we have that

$$V_{1, \widehat{\mathcal{M}}}^\pi(s_1) - V_1^\pi(s_1) = \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_{h, \widehat{\mathcal{M}}}^\pi} \left(\widehat{r}_h(s, a) - r_h(s, a) + \left\langle \widehat{P}_h(\cdot|s, a) - P_h(\cdot|s, a), V_{h+1}^\pi(\cdot) \right\rangle \right)$$

Equivalently, we have that

$$V_{1, \widehat{\mathcal{M}}}^\pi(s_1) - V_1^\pi(s_1) = \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_h^\pi} \left(\widehat{r}_h(s, a) - r_h(s, a) + \left\langle \widehat{P}_h(\cdot|s, a) - P_h(\cdot|s, a), V_{h+1, \widehat{\mathcal{M}}}^\pi(\cdot) \right\rangle \right)$$

where $V_{H+1}^\pi = \widehat{V}_{H+1, \widehat{\mathcal{M}}} = 0$ (as the episode ends at H).

Proof. Recall that $d_h^\pi(s, a)$ and $d_{h, \widehat{\mathcal{M}}}^\pi(s, a)$ are the occupancy measures of any $(s, a) \in \mathcal{S} \times \mathcal{A}$ at timestep $h \in [H]$ when executing policy π in \mathcal{M} and $\widehat{\mathcal{M}}$, respectively. By definition, we have that

$$\begin{aligned} & V_{1, \widehat{\mathcal{M}}}^\pi(s_1) - V_1^\pi(s_1) \\ &= \sum_a \pi_1(s_1, a) \left(\widehat{r}_1(s_1, a) - r_1(s_1, a) + \sum_{s'} \left(\widehat{P}_1(s'|s_1, a) \widehat{V}_{2, \widehat{\mathcal{M}}}^\pi(s') - P_1(s'|s_1, a) V_2^\pi(s') \right) \right) \\ &= \mathbb{E}_{(s,a) \sim d_{1, \widehat{\mathcal{M}}}^\pi} \left[\widehat{r}_1(s, a) - r_1(s, a) + \left\langle \widehat{P}_1(\cdot|s, a), \widehat{V}_{2, \widehat{\mathcal{M}}}^\pi(\cdot) - V_2^\pi(\cdot) \right\rangle + \left\langle \widehat{P}_1(\cdot|s, a) - P_1(\cdot|s, a), V_2^\pi(\cdot) \right\rangle \right] \\ &= \dots = \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_{h, \widehat{\mathcal{M}}}^\pi} \left[\widehat{r}_h(s, a) - r_h(s, a) + \left\langle \widehat{P}_h(\cdot|s, a) - P_h(\cdot|s, a), V_{h+1}^\pi(\cdot) \right\rangle \right] \end{aligned}$$

Equivalently, we have that

$$\begin{aligned} & V_{1, \widehat{\mathcal{M}}}^\pi(s_1) - V_1^\pi(s_1) \\ &= \mathbb{E}_{(s,a) \sim d_1^\pi} \left[\widehat{r}_1(s, a) - r_1(s, a) + \left\langle P_1(\cdot|s, a), \widehat{V}_{2, \widehat{\mathcal{M}}}^\pi(\cdot) - V_2^\pi(\cdot) \right\rangle + \left\langle \widehat{P}_1(\cdot|s, a) - P_1(\cdot|s, a), \widehat{V}_{2, \widehat{\mathcal{M}}}^\pi(\cdot) \right\rangle \right] \\ &= \dots = \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_h^\pi} \left[\widehat{r}_h(s, a) - r_h(s, a) + \left\langle \widehat{P}_h(\cdot|s, a) - P_h(\cdot|s, a), \widehat{V}_{h+1, \widehat{\mathcal{M}}}^\pi(\cdot) \right\rangle \right] \end{aligned}$$

which concludes the proof. \square

Lemma E.6. Let $\Lambda_t = \lambda I + \sum_{j=1}^t \bar{\phi}^j (\bar{\phi}^j)^\top$ where $\bar{\phi}_i \in \mathbb{R}^d, \lambda > 0, d > 0$, then we have

$$\sum_{j=1}^t (\bar{\phi}^j)^\top \Lambda_t^{-1} \bar{\phi}^j \leq d.$$

Proof. Please refer to (Jin et al., 2020, Lemma D.1). \square

Lemma E.7. Let $\{x_j\}_{j=1}^\infty$ be a stochastic process on the space $(\mathcal{S}, \{i_v\}_{i=0}^{\lceil H/v \rceil})$ with corresponding filtration $\{\mathcal{G}_j\}_{j=0}^\infty$. Let $\{\bar{\phi}_j\}_{j=0}^\infty$ be an $\mathbb{R}^{\bar{d}}$ -valued stochastic process where $\bar{\phi}_j \in \mathcal{G}_{j-1}$, and $\|\bar{\phi}_j\| \leq 1$. Let $\Lambda_t = \lambda I + \sum_{j=1}^{t-1} \bar{\phi}_j \bar{\phi}_j^\top$. Then for any $\delta > 0$, with probability at least $1 - \delta$, for all $t \geq 0$, and any $V \in \mathcal{V}$ so that $\sup_x |V(x)| \leq H$, we have:

$$\left\| \sum_{j=1}^{t-1} \bar{\phi}_j \{V(x_j) - \mathbb{E}[V(x_j) | \mathcal{G}_{j-1}]\} \right\|_{\Lambda_t^{-1}}^2 \leq 4H^2 \left[\frac{\bar{d}}{2} \log \left(\frac{t + \lambda}{\lambda} \right) + \log \frac{\mathcal{N}_\epsilon}{\delta} \right] + \frac{8t^2 \epsilon^2}{\lambda},$$

where \mathcal{N}_ϵ is the ϵ -covering number of \mathcal{V} with respect to the ℓ_∞ -norm $\sup_x |V(x) - V'(x)|$.

Proof. Please refer to (Jin et al., 2020, Lemma D.4). \square

Lemma E.8. Suppose $\bar{\phi}_t \in \mathbb{R}^{\bar{d}}$ and $\|\bar{\phi}_t\| \leq 1$ for all $t \geq 0$. For any $t \geq 0$, we define $\Lambda_t = I + \sum_{j=1}^t \bar{\phi}_j \bar{\phi}_j^\top$. Then we have

$$\sum_{j=1}^t \bar{\phi}_j^\top \Lambda_{j-1}^{-1} \bar{\phi}_j \leq 2 \log \left[\frac{\det(\Lambda_t)}{\det(\Lambda_0)} \right].$$

Proof. Please refer to (Jin et al., 2020, Lemma D.2). □